

INTERACTION DESIGN



beyond human-computer interaction
5th Edition

Interaction Design continues to be the standard textbook in the field. Seasoned practitioners will find it useful when they need a reference to best practices or to explain a concept to a colleague. Students can turn to *Interaction Design* for an easy-to-understand description of the basics or in-depth how-tos. From personas and disabilities to the design of UX organizations and working in Agile, if you're going to pick one book to bring into the office, it should be this one.

Jofish Kaye, Principal Research Scientist, Mozilla, USA

This is the perfect textbook for a wide range of User interface/User experience design courses. For an undergraduate, it provides a variety of compelling examples that illustrate best practice in Interaction Design. For a graduate student, it provides a foundational overview of advanced topics. This book is also essential for the professional who wants to know the state of the art in Interaction design. I use this textbook and recommend it widely.

*Rosa I. Arriaga, Ph.D., Senior Research Scientist, School of Interactive Computing
Georgia Institute of Technology, USA*

The *Interaction Design* book has immensely contributed to a growing Namibian HCI skilled community over the last decade. Exposing students, academics and practitioners to the basic principles and theories as well as most recent trends and technologies, with global and local case studies, in the latest edition, allows for reflective applications within very specific contexts. This book remains our number one reference in the education of future generations of interaction designers in Namibia, promoting the creation of thoughtful user experiences for responsible citizens.

*Heike Winschiers-Theophilus, Professor, Faculty of Computing and Informatics,
Namibia University of Science and Technology, Africa*

Throughout my teaching of user experience and interaction design, the book by Rogers, Preece and Sharp has been an absolute cornerstone textbook for students. The authors bring together their own wealth of knowledge of academic HCI with a deep understanding of industry practice to provide what must be the most comprehensive introduction to the key areas of interaction design and user experience work, now an established field of practice. I put this book in the “essential reading” section of many of the reading lists I give to students.

Simon Attfield, Associate Professor in Human Centred Technology, Middlesex University, UK

Interaction design has gone through tremendous changes in the last few years—for example the rising importance of “big” data streams to design, and the growing prevalence of everyday ubiquitous computing issues of sensing and blending gracefully and ethically into peoples’ daily lives. This is an important and timely update to a text that’s long been considered gold standard in our field. I’m looking forward to using it with my students to help prepare them for the design challenges they will face in today’s industrial practice.

Katherine Isbister, Professor, Computational Media, University of California Santa Cruz, USA

More than ever, designing effective human-computer interactions is crucial for modern technological systems. As digital devices become smaller, faster and smarter, the interface and interaction challenges become ever more complex. Vast quantities of data are often accessed on handheld screens, or no screens at all through voice commands; and AI systems have interfaces that “bite back” with sophisticated dialogue structures. What are the best interaction metaphors for these technologies? What are the best tools for creating interfaces that are enjoyable and universally accessible? How do we ensure emerging technologies remain relevant and respectful of human values? In this book, you’ll find detailed analysis of these questions and much more. (It is a valuable resource for both the mature student and the reflective professional.)

*Frank Vetere, Professor of Interaction Design, School of Computing and Information Systems,
University of Melbourne, Australia*

This is at the top of my recommended reading list for undergraduate and master’s students as well as professionals looking to change career paths. Core issues to interaction design are brought to life through compelling vignettes and contemporary case examples from leading experts. What has long been a comprehensive resource for interaction design now incorporates timely topics in computing, such as data at scale, artificial intelligence, and ethics, making it essential reading for anyone entering the field of interaction design.

*Anne-Marie Piper, PhD, Associate Professor, Departments of Communication Studies,
Electrical Engineering and Computer Science, Northwestern University, USA*

I have been using *Interaction Design* as a textbook since its first edition for both my undergraduate and graduate introductory HCI courses. This is a must-read seminal book which provides a thorough coverage of the discipline of HCI and the practice of user-centered design. The fifth edition lives up to its phenomenal reputation by including updated content on the process of interaction design, the practice of interaction design (e.g., technical debt in UX, Lean UX), design ethics, new types of interfaces, etc. I always recommend *Interaction Design* to students and practitioners who want to gain a comprehensive overview of the fields of HCI and UX.

Olivier St-Cyr, Assistant Professor, Teaching Stream, University of Toronto, Canada

Interaction design is a practice that spans many domains. The authors acknowledge this by providing a tremendous amount of information across a wide spectrum of disciplines. This book has evolved from a simple textbook for HCI students, to an encyclopedia of design practices, examples, discussions of related topics, suggestions for further reading, exercises, interviews with practitioners, and even a bit of interesting history here and there. I see it as one of the few sources effectively bridging the gulf between theory and practice. A copy has persistently occupied my desk since the first edition, and I regularly find myself revisiting various sections for inspiration on how to communicate the reasoning behind my own decisions to colleagues and peers.

William R. Hazlewood, PhD, Principal Design Technologist, Retail Experience Design Concept Lab, Amazon, USA

For years *Interaction Design* has been my favourite book not only for supporting my classes but also as my primary source for preparing UX studies to industrial and academic settings. The chapters engage readers with easy-to-read content while presenting, harmonically, theories, examples and case studies which touch in multidisciplinary aspects of construction and evaluation of interactive products. The fifth edition again maintains the tradition of being an up-to-date book on HCI, and includes new discussions on Lean UX, emotional interaction, social and cognitive aspects, and ethics in human studies, which are certainly contemporary topics of utmost relevance for practitioners and academics in interaction design.

Luciana Zaina, Senior Lecturer, Federal University of São Carlos, Brazil

This book is always my primary recommendation for newcomers to human-computer interaction. It addresses the subject from several perspectives: understanding of human behaviour in context, the challenges of ever-changing technology, and the practical processes involved in interaction design and evaluation. The new edition again shows the authors' dedication to keeping both the primary content and representative examples up to date.

Robert Biddle, Professor of Human-Computer Interaction, Carleton University, Ottawa, Canada

This fifth edition provides a timely update to one of the must-have classics on interaction design. The changes in our field, including how to deal with emerging sensing technology and the volumes of data it provides, are well addressed in this volume. This is a book for those new to and experienced in interaction design.

Jodi Forlizzi, Professor and Geschke Director, Human-Computer Interaction Institute, The School of Computer Science, CMU, USA

The milieu of digital life surrounds us. However, how we choose to design and create our experiences and interactions with these emerging technologies remains a significant challenge. This book provides both a road-map of essential skills and methodologies to tackle these designs confidently as well as the critical deeper history, literature, and poetry of interaction design. You will return to this book throughout your career to operationalize, ground and inspire your creative practice of interaction design.

Eric Paulos, Professor, Electrical Engineering and Computer Sciences, UC Berkeley, USA

Preece, Sharp and Rogers offer once again an engaging excursion through the world of interaction design. This series is always up-to-date and offers a fresh view on a broad range of topics needed for students in the field of interaction design, human-computer interaction, information design, web design or ubiquitous computing. The book should be the book every student should have in their backpack. It is a "survival guide"! It guides one through the jungle of information and the dark technological forests of our digital age. It also helps to develop a critical view on developing novel technologies as our computing research community needs to confront much more seriously the negative impacts of our innovations. The online resources are a great help for me to create good classes and remove some weight from the backpacks of my students.

Johannes Schöning, Professor of Computer Science, Bremen University, Germany

Nearly 20 years have passed since the release of the first edition of *Interaction Design*, with massive changes to technology and thus the science and practice of interaction design. The new edition combines the brilliance of the first book with the wisdom of the lessons learned in the meantime, and the excitement of new technological frontiers. Complex concepts are elegantly and beautifully explained, and the reader is left with little doubt as to how to put them into practice. The book is an excellent resource for those new to interaction design, or as a guidebook or reference to practitioners.

Dana McKay, UX Researcher, Practitioner and Academic, University of Melbourne, Australia

Computers are ubiquitous and embedded in virtually every new device and system, ranging from the omnipresent cellphone to the complex web of sociotechnical systems that envelop most every sphere of personal and professional life. They connect our activities to ever-expanding information resources with previously unimaginable computational power. To ensure interface design respects human needs and augments our abilities is an intellectual challenge of singular importance. It involves not only complex theoretical and methodological issues of how to design effective representations and mechanisms of interaction but also confronts complex social, cultural, and political issues such as those of privacy, control of attention, and ownership of information. The new edition of *Interaction Design* continues to be the introductory book I recommend to my students and to anyone interested in this crucially important area.

Jim Hollan, Distinguished Professor of Cognitive Science, University of California San Diego, USA

Interaction Design continues to be my favorite textbook on HCI. Even named our undergraduate and postgraduate programmes at Aalborg University after it. In its fifth edition, it captures the newest developments in the field's cumulative body of knowledge, and continues to be the most updated and accessible work available. As always, it serves as a clear pointer to emerging trends in interactive technology design and use.

Jesper Kjeldskov, Professor and Head of Department of Computer Science, Aalborg University, Denmark

I got to learn about the field of HCI and *interaction design* when I came across the first edition of this book at the library in my junior year of college. As an HCI researcher and educator, I have been having the pleasure of introducing the subject to undergraduates and professional master's students using the previous editions. I thank the authors for their studious efforts to update and add new contents that are relevant for students, academics, and professionals to help them learn this ever-evolving field of HCI and *interaction design* in a delightful manner.

Eun Kyoung Choe, Professor of Human-Computer Interaction, College of Information Studies, University of Maryland, USA

This new edition is, without competition, the most comprehensive and authoritative source in the field when it comes to modern interaction design. It is highly accessible and it is a pleasure to read. The authors of this book have once again delivered what the field needs!

Erik Stolterman, Professor in Informatics, School of Informatics and Computing, Indiana University, Bloomington, USA

This book illuminates the interaction design field like no other. Interaction design is such a vast, multidisciplinary field that you might think it would be impossible to synthesize the most relevant knowledge in one book. This book does not only that, but goes even further: it eloquently brings contemporary examples and diverse voices to make the knowledge concrete and actionable, so it is useful for students, researchers, and practitioners alike. This new edition includes invaluable discussions about the current challenges we now face with data at scale, embracing the ethical design concerns our society needs so much in this era.

Simone D. J. Barbosa, Professor of Computer Science, PUC-Rio, and Co-Editor-in-Chief of ACM Interactions, Brazil

My students like this book a lot! It provides a comprehensive coverage of the essential aspects of HCI/UX, which is key to the success of any software applications. I also like many aspects of the book, particularly the examples and videos (some of which are provided as hyperlinks) because they not only help to illustrate the HCI/UX concepts and principles, but also relate very well to readers. I highly recommend this book to anyone who wants to learn more about HCI/UX.

Fiona Fui-Hoon Nah, Editor-in-Chief of AIS Transactions on Human-Computer Interaction, Professor of Business and Information Technology, Missouri University of Science and Technology, Rolla, Missouri, USA

I have been using the book for several years in my Human-Computer Interaction class. It helps me, not only for teaching, but also for these supervision. I really appreciate the authors regarding their efforts in maintaining the relevance and up-to-dateness of the Interaction Design book. For example, they put Data At Scale and AgileUX in the new edition. Really love the book!

Harry B. Santoso, PhD, Instructor of Interaction System (HCI) course at Faculty of Computer Science, Universitas Indonesia, Indonesia

During my PhD already the first edition of *Interaction Design: beyond human-computer interaction* in 2002 quickly became my preferred reference book. Seventeen years later, and now in its fifth edition, I commend the authors for their meticulous and consistent effort in updating and enriching what has become the field's standard introductory textbook. Not just about objects and artefact, design today is increasingly recognized as a sophisticated and holistic approach for systems thinking. Similarly, Preece, Sharp, and Rogers have kept the book's coverage with the times by providing a comprehensive, compelling, and accessible coverage of concepts, methods and cases of interaction design across many domains such as experience design, ubiquitous computing, and urban informatics.

Marcus Foth, Professor of Urban Informatics, QUT Design Lab, Brisbane, Australia

"Interaction Design" has long been my textbook of choice for general HCI courses. The latest edition has introduced a stronger practitioner focus that should add value for students transitioning into practice, for practitioners, and also for others interested in interaction design and its role in product development. It manages to be an engaging read while also being "snackable", to cover the basics and also inspire. I still find it a great read, and believe others will too."

Ann Blandford, Professor of Human – Computer Interaction, University College London

Very clear style, with plenty of active learning material and pointers to further reading. I found that it works very well with engineering students.

Albert Ali Salah, Professor at Utrecht University, the Netherlands

INTERACTION DESIGN

beyond human-computer
interaction

Fifth Edition



WILEY

Interaction Design: beyond human-computer interaction, Fifth Edition

Published by
John Wiley & Sons, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2019 by John Wiley & Sons, Inc., Indianapolis, Indiana
Published simultaneously in Canada

ISBN: 978-1-119-54725-9

ISBN: 978-1-119-54735-8 (ebk)

ISBN: 978-1-119-54730-3 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or website may provide or recommendations it may make. Further, readers should be aware that Internet websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2019932998

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

About the Authors

The authors are senior academics with a background in teaching, researching, and consulting in the United Kingdom, United States, Canada, India, Australia, South Africa, and Europe. Having worked together on four previous editions of this book, as well as an earlier textbook on human-computer interaction, they bring considerable experience in curriculum development using a variety of media for online learning as well as face-to-face teaching. They have considerable knowledge in creating learning texts and websites that motivate and support learning for a range of students. All three authors are specialists in interaction design and human-computer interaction (HCI). In addition, they bring skills from other disciplines; for instance, Yvonne Rogers started off as a cognitive scientist, Helen Sharp is a software engineer, and Jenny Preece works in information systems. Their complementary knowledge and skills enable them to cover the breadth of concepts in interaction design and HCI to produce an interdisciplinary text and website.

Helen Sharp is a Professor of Software Engineering and Associate Dean in the Faculty of Science, Technology, Engineering, and Mathematics at the Open University. Originally trained as a software engineer, it was by watching the frustration of users and the clever “work-arounds” they developed that inspired her to investigate HCI, user-centered design, and the other related disciplines that now underpin the field of interaction design. Her research focuses on the study of professional software practice and the effect of human and social aspects on software development, leveraging her expertise in the intersection between interaction design and software engineering and working closely with practitioners to support practical impact. She is active in both the software engineering and CHI communities, and she has had a long association with practitioner-related conferences. Helen is on the editorial board of several software engineering journals, and she is a regular invited speaker at academic and practitioner venues.

Yvonne Rogers is the Director of the Interaction Centre at University College London, a Professor of Interaction Design, and a deputy head of department for Computer Science. She is internationally renowned for her work in HCI and ubiquitous computing and, in particular, for her pioneering approach to innovation and ubiquitous learning. Yvonne is widely published, and she is the author of two recent books: *Research in the Wild* (2017, co-authored with Paul Marshall) and *The Secrets of Creative People* (2014). She is also a regular keynote speaker at computing and HCI conferences worldwide. Former positions include Professor of Interaction Design at the Open University (2006–2011), Professor of Human-Computer Interaction at the School of Informatics and Computing at Indiana University (2003–2006), and Professor in the former School of Cognitive and Computing Sciences at Sussex University (1992–2003). She has also been a Visiting Professor at UCSC, University of Cape Town, Melbourne University, Stanford, Apple, Queensland University, and UCSD. She has been elected as a Fellow of the ACM, the British Computer Society, and the ACM’s CHI Academy.

Jennifer Preece is Professor and Dean Emerita in the College of Information Studies—Maryland’s iSchool—at the University of Maryland. Jenny’s research focuses on the intersection of information, community, and technology. She is particularly interested in community participation online and offline. She has researched ways to support empathy and social support online, patterns of online participation, reasons for not participating (for example, lurking and infrequent participation), strategies for supporting online communication, development of norms, and the attributes of successful technology-supported communities. Currently, Jenny focuses on how technology can be used to educate and motivate citizens to engage and contribute quality data to citizen science projects. This research contributes to the broader need for the collection of data about the world’s flora and fauna at a time when many species are in rapid decline due to habitat loss, pollution, and climate change. She was author of one of the first books on online communities—*Online Communities: Designing Usability, Supporting Sociability* (2000) published by John Wiley & Sons Ltd and several other HCI texts. Jenny is also widely published, a regular keynote speaker, and a member of the ACM’s CHI Academy.

Credits

Associate Publisher

Jim Minatel

Editorial Manager

Pete Gaughan

Production Manager

Katie Wisor

Project Editor

Gary Schwartz

Production Editor

Barath Kumar Rajasekaran

Technical Editors

Danelle Bailey

Jill L. H. Reed

Copy Editor

Kim Wimpsett

Proofreader

Nancy Bell

Indexer

Johnna VanHoose Dinse

Cover Designer

Wiley

Cover Image

© Wiley; Jennifer Preece photo courtesy of

Craig Allan Taylor

Acknowledgments

Many people have helped us over the years in writing the four previous editions of this book. We have benefited from the advice and support of our many professional colleagues across the world and from our students, friends, and families. We especially would like to thank everyone who generously contributed their ideas and time to help make all of the editions of this book successful.

These include our colleagues and students at the College of Information Studies—“Maryland’s iSchool”—at University of Maryland, the Human-Computer Interaction Laboratory (HCIL) and Center for the Advanced Study of Communities and Information (CASCI), the Open University, and University College London. We would especially like to thank (in alphabetical first name order) all of the following individuals who have helped us over the years:

Alex Quinn, Alice Robbin, Alice Siempelkamp, Alina Goldman, Allison Druin, Ana Javornik, Anijo Mathew, Ann Blandford, Ann Jones, Anne Adams, Ben Bederson, Ben Shneiderman, Blaine Price, Carol Boston, Cathy Holloway, Clarisse Sieckenius de Souza, Connie Golsteijn, Dan Green, Dana Rotman, danah boyd, Debbie Stone, Derek Hansen, Duncan Brown, Edwin Blake, Eva Hornecker, Fiona Nah, Gill Clough, Godwin Egbeyi, Harry Brignull, Janet van der Linden, Jeff Rick, Jennifer Ferreira, Jennifer Golbeck, Jeremy Mayes, Joh Hunt, Johannes Schöning, Jon Bird, Jonathan Lazar, Judith Segal, Julia Galliers, Fiona Nah, Kent Norman, Laura Plonka, Leeann Brumby, Leon Reicherts, Mark Woodroffe, Michael Wood, Nadia Pantidi, Nick Dalton, Nicolai Marquardt, Paul Cairns, Paul Marshall, Philip “Fei” Wu, Rachael Bradley, Rafael Cronin, Richard Morris, Richie Hazlewood, Rob Jacob, Rose Johnson, Stefan Kreitmayer, Steve Hodges, Stephanie Wilson, Tamara Clegg, Tammy Toscos, Tina Fuchs, Tom Hume, Tom Ventsias, Toni Robertson, and Youn-Kyung Lim.

In addition we wish to thank the many students, instructors, researchers and practitioners who have contacted us over the years with stimulating comments, positive feedback and provocative questions

We are particularly grateful to Vikram Mehta, Nadia Pantidi, and Mara Balestrini for filming, editing, and compiling a series of on-the-spot “talking heads” videos, where they posed probing questions to the diverse set of attendees at CHI’11, CHI’14, and CHI’18, including a variety of CHI members from across the globe. The questions included asking about the future of interaction design and whether HCI has gone too wild. There are about 75 of these videos, which can be viewed on our website at www.id-book.com. We are also indebted to danah boyd, Harry Brignull, Leah Beuchley, Albrecht Schmidt, Ellen Gottesdiener, and Jon Froehlich for generously contributing in-depth, text-based interviews in the book. We would like to thank Rien Sach, who has been our webmaster for several years, and Deb Yuill who did a thoughtful and thorough job of editing the old reference list.

Danelle Bailey and Jill Reed provided thoughtful critiques and suggestions on all the chapters in the fifth edition, and we thank them.

Finally, we would like to thank our editor and the production team at Wiley who have been very supportive and encouraging throughout the process of developing this fifth edition: Jim Minatel, Pete Gaughan, Gary Schwartz, and Barath Kumar Rajasekaran.

Contents

| | |
|--|-------------|
| What's Inside? | xvii |
| 1 WHAT IS INTERACTION DESIGN? | 1 |
| 1.1 Introduction | 1 |
| 1.2 Good and Poor Design | 2 |
| 1.3 What Is Interaction Design? | 9 |
| 1.4 The User Experience | 13 |
| 1.5 Understanding Users | 15 |
| 1.6 Accessibility and Inclusiveness | 17 |
| 1.7 Usability and User Experience Goals | 19 |
| Interview with Harry Brignull | 34 |
| 2 THE PROCESS OF INTERACTION DESIGN | 37 |
| 2.1 Introduction | 37 |
| 2.2 What Is Involved in Interaction Design? | 38 |
| 2.3 Some Practical Issues | 55 |
| 3 CONCEPTUALIZING INTERACTION | 69 |
| 3.1 Introduction | 69 |
| 3.2 Conceptualizing Interaction | 71 |
| 3.3 Conceptual Models | 74 |
| 3.4 Interface Metaphors | 78 |
| 3.5 Interaction Types | 81 |
| 3.6 Paradigms, Visions, Theories, Models, and Frameworks | 88 |
| Interview with Albrecht Schmidt | 97 |
| 4 COGNITIVE ASPECTS | 101 |
| 4.1 Introduction | 101 |
| 4.2 What Is Cognition? | 102 |
| 4.3 Cognitive Frameworks | 123 |
| 5 SOCIAL INTERACTION | 135 |
| 5.1 Introduction | 135 |
| 5.2 Being Social | 136 |
| 5.3 Face-to-Face Conversations | 139 |

| | | |
|----------|--|------------|
| 5.4 | Remote Conversations | 143 |
| 5.5 | Co-presence | 150 |
| 5.6 | Social Engagement | 158 |
| 6 | EMOTIONAL INTERACTION | 165 |
| 6.1 | Introduction | 165 |
| 6.2 | Emotions and the User Experience | 166 |
| 6.3 | Expressive Interfaces and Emotional Design | 172 |
| 6.4 | Annoying Interfaces | 174 |
| 6.5 | Affective Computing and Emotional AI | 179 |
| 6.6 | Persuasive Technologies and Behavioral Change | 182 |
| 6.7 | Anthropomorphism | 187 |
| 7 | INTERFACES | 193 |
| 7.1 | Introduction | 193 |
| 7.2 | Interface Types | 194 |
| 7.3 | Natural User Interfaces and Beyond | 252 |
| 7.4 | Which Interface? | 253 |
| | Interview with Leah Buechley | 257 |
| 8 | DATA GATHERING | 259 |
| 8.1 | Introduction | 259 |
| 8.2 | Five Key Issues | 260 |
| 8.3 | Data Recording | 266 |
| 8.4 | Interviews | 268 |
| 8.5 | Questionnaires | 278 |
| 8.6 | Observation | 287 |
| 8.7 | Choosing and Combining Techniques | 300 |
| 9 | DATA ANALYSIS, INTERPRETATION, AND PRESENTATION | 307 |
| 9.1 | Introduction | 307 |
| 9.2 | Quantitative and Qualitative | 308 |
| 9.3 | Basic Quantitative Analysis | 311 |
| 9.4 | Basic Qualitative Analysis | 320 |
| 9.5 | Which Kind of Analytic Framework to Use? | 329 |
| 9.6 | Tools to Support Data Analysis | 341 |
| 9.7 | Interpreting and Presenting the Findings | 342 |

| | |
|--|------------|
| 10 DATA AT SCALE | 349 |
| 10.1 Introduction | 349 |
| 10.2 Approaches to Collecting and Analyzing Data | 351 |
| 10.3 Visualizing and Exploring Data | 366 |
| 10.4 Ethical Design Concerns | 375 |
| 11 DISCOVERING REQUIREMENTS | 385 |
| 11.1 Introduction | 385 |
| 11.2 What, How, and Why? | 386 |
| 11.3 What Are Requirements? | 387 |
| 11.4 Data Gathering for Requirements | 395 |
| 11.5 Bringing Requirements to Life: Personas and Scenarios | 403 |
| 11.6 Capturing Interaction with Use Cases | 415 |
| Interview with Ellen Gottesdiener | 418 |
| 12 DESIGN, PROTOTYPING, AND CONSTRUCTION | 421 |
| 12.1 Introduction | 421 |
| 12.2 Prototyping | 422 |
| 12.3 Conceptual Design | 434 |
| 12.4 Concrete Design | 445 |
| 12.5 Generating Prototypes | 447 |
| 12.6 Construction | 457 |
| Interview with Jon Froehlich | 466 |
| 13 INTERACTION DESIGN IN PRACTICE | 471 |
| 13.1 Introduction | 471 |
| 13.2 AgileUX | 473 |
| 13.3 Design Patterns | 484 |
| 13.4 Open Source Resources | 489 |
| 13.5 Tools for Interaction Design | 491 |
| 14 INTRODUCING EVALUATION | 495 |
| 14.1 Introduction | 495 |
| 14.2 The Why, What, Where, and When of Evaluation | 496 |
| 14.3 Types of Evaluation | 500 |
| 14.4 Evaluation Case Studies | 507 |
| 14.5 What Did We Learn from the Case Studies? | 514 |
| 14.6 Other Issues to Consider When Doing Evaluation | 516 |

| | |
|---|------------|
| 15 EVALUATION STUDIES: FROM CONTROLLED TO NATURAL SETTINGS | 523 |
| 15.1 Introduction | 523 |
| 15.2 Usability Testing | 524 |
| 15.3 Conducting Experiments | 533 |
| 15.4 Field Studies | 536 |
| Interview with danah boyd | 546 |
| 16 EVALUATION: INSPECTIONS, ANALYTICS, AND MODELS | 549 |
| 16.1 Introduction | 549 |
| 16.2 Inspections: Heuristic Evaluation and Walk-Throughs | 550 |
| 16.3 Analytics and A/B Testing | 567 |
| 16.4 Predictive Models | 576 |
| References | 581 |
| Index | 619 |

What's Inside?

Welcome to the fifth edition of *Interaction Design: beyond human-computer interaction* and our interactive website at www.id-book.com. Building on the success of the previous editions, we have substantially updated and streamlined the material in this book to provide a comprehensive introduction to the fast-growing and multi-disciplinary field of interaction design. Rather than let the book expand, however, we have again made a conscious effort to keep it at the same size.

Our textbook is aimed at both professionals who want to find out more about interaction design and students from a range of backgrounds studying introductory classes in human-computer interaction, interaction design, information and communications technology, web design, software engineering, digital media, information systems, and information studies. It will appeal to practitioners, designers, and researchers who want to discover what is new in the field or to learn about a specific design approach, method, interface, or topic. It is also written to appeal to a general audience interested in design and technology.

It is called *Interaction Design: beyond human-computer interaction* because interaction design has traditionally been concerned with a broader scope of issues, topics, and methods than was originally the scope of human-computer interaction (HCI)—although nowadays, the two increasingly overlap in scope and coverage of topics. We define interaction design as follows:

Designing interactive products to support the way people communicate and interact in their everyday and working lives.

Interaction design requires an understanding of the capabilities and desires of people and the kinds of technology that are available. Interaction designers use this knowledge to discover requirements and develop and manage them to produce a design. Our textbook provides an introduction to all of these areas. It teaches practical techniques to support development as well as discussing possible technologies and design alternatives.

The number of different types of interface and applications available to today's interaction designers continues to increase steadily, so our textbook, likewise, has been expanded to cover these new technologies. For example, we discuss and provide examples of brain, smart, robotic, wearable, shareable, augmented reality, and multimodal interfaces, as well as more traditional desktop, multimedia, and web-based interfaces. Interaction design in practice is changing fast, so we cover a range of processes, issues, and examples throughout the book.

The book has 16 chapters, and it includes discussion of the different design approaches in common use; how cognitive, social, and affective issues apply to interaction design; and how to gather, analyze, and present data for interaction design. A central theme is that design and evaluation are interwoven, highly iterative processes, with some roots in theory but that rely strongly on good practice to create usable products. The book has a hands-on orientation and explains how to carry out a variety of techniques used to design and evaluate the wide range of new applications coming onto the market. It has a strong pedagogical design and includes many activities (with detailed comments) and more complex activities that can form the basis for student projects. There are also “Dilemmas,” which encourage readers to weigh the pros and cons of controversial issues.

TASTERS

We address topics and questions about the what, why, and how of interaction design. These include the following:

- Why some interfaces are good and others are poor
- Whether people can really multitask
- How technology is transforming the way people communicate with one another
- What are users' needs, and how we can design for them
- How interfaces can be designed to change people's behavior
- How to choose between the many different kinds of interactions that are now available (for example, talking, touching, and wearing)
- What it means to design accessible and inclusive interfaces
- The pros and cons of carrying out studies in the lab versus in the field and in the wild
- When to use qualitative and quantitative methods
- How to construct informed consent forms
- How the type of interview questions posed affects the conclusions that can be drawn from the answers given
- How to move from a set of scenarios and personas to initial low-fidelity prototypes
- How to visualize the results of data analysis effectively
- How to collect, analyze, and interpret data at scale
- Why it is that what people say can be different from what they do
- The ethics of monitoring and recording people's activities
- What are Agile UX and Lean UX and how they relate to interaction design
- How Agile UX can be practically integrated with interaction design throughout different stages of the design process ■

The style of writing throughout the book is intended to be accessible to a range of readers. It is largely conversational in nature and includes anecdotes, cartoons, and case studies. Many of the examples are intended to relate to readers' own experiences. The book and the associated website are also intended to encourage readers to be active when reading and to think about seminal issues. The goal is for readers to understand that much of interaction design needs consideration of the issues, and that they need to learn to weigh the pros and cons and be prepared to make trade-offs. There is rarely a right or wrong answer, although there is a world of difference between a good design and a poor design.

This book is accompanied by a website (www.id-book.com), which provides a variety of resources, including slides for each chapter, comments on chapter activities, and a number of in-depth case studies written by researchers and designers. There are video interviews with a wide range of experts from the field, including professional interaction designers and university professors. Pointers to respected blogs, online tutorials, YouTube videos, and other useful materials are also provided.


Changes from Previous Editions

To reflect the dynamic nature of the field, the fifth edition has been thoroughly updated, and new examples, images, case studies, dilemmas, and so on, have been included to illustrate the changes. Included in this edition is a new chapter called “Data at Scale.” Collecting data has never been easier. However, knowing what to do with it when designing new user experiences is much more difficult. The chapter introduces key methods for collecting data at scale, discusses how to transform data at scale to be meaningful, and reviews a number of methods for visualizing and exploring data at scale while introducing fundamental design principles for making data at scale ethical. This is positioned just after two chapters that introduce data gathering and data analysis that discuss fundamental methods.

In this edition, the chapter on the Process of Interaction Design has been re-located to Chapter 2 in order to better frame the discussion of interaction design. It has been updated with new process models and modified to fit its new location in the book structure. This means that the other chapters have been renumbered to accommodate this and the new chapter.

Chapter 13, “Interaction Design in Practice,” has been updated to reflect recent developments in the use of practical UX methods. Old examples and methods no longer used in the field have been removed to make way for the new material. Some chapters have been completely rewritten, while others have been extensively revised. For example, Chapters 4, 5, and 6 have been substantially updated to reflect new developments in social media and emotional interaction, while also covering the new interaction design issues they raise, such as privacy and addiction. Many examples of new interfaces and technologies have been added to Chapter 7. Chapter 8 and Chapter 9 on data gathering and analysis have also been substantially updated. New case studies and examples have been added to Chapters 14–16 to illustrate how evaluation methods have changed for use with the continuously evolving technology that is being developed for today’s users. The interviews accompanying the chapters have been updated, and two new ones are included with leading figures involved in innovative research, state-of-the-art design, and contemporary practice.

We have decided to continue to provide both a print-based version of the book and an e-book. Both are in full color. The e-book supports note sharing, annotating, contextualized navigating, powerful search features, inserted videos, links, and quizzes.



Chapter 1

WHAT IS INTERACTION DESIGN?

- 1.1 Introduction
- 1.2 Good and Poor Design
- 1.3 What Is Interaction Design?
- 1.4 The User Experience
- 1.5 Understanding Users
- 1.6 Accessibility and Inclusiveness
- 1.7 Usability and User Experience Goals

Objectives

The main goals of this chapter are to accomplish the following:

- Explain the difference between good and poor interaction design.
- Describe what interaction design is and how it relates to human-computer interaction and other fields.
- Explain the relationship between the user experience and usability.
- Introduce what is meant by accessibility and inclusiveness in relation to human-computer interaction.
- Describe what and who is involved in the process of interaction design.
- Outline the different forms of guidance used in interaction design.
- Enable you to evaluate an interactive product and explain what is good and bad about it in terms of the goals and core principles of interaction design.

1.1 Introduction

How many interactive products are there in everyday use? Think for a minute about what you use in a typical day: a smartphone, tablet, computer, laptop, remote control, coffee machine, ticket machine, printer, GPS, smoothie maker, e-reader, smart TV, alarm clock, electric toothbrush, watch, radio, bathroom scales, fitness tracker, game console . . . the list is endless. Now think for a minute about how usable they are. How many are actually easy, effortless, and

enjoyable to use? Some, like the iPad, are a joy to use, where tapping an app and flicking through photos is simple, smooth, and enjoyable. Others, like working out how to buy the cheapest train ticket from a ticket machine that does not recognize your credit card after completing a number of steps and then makes you start again from scratch, can be very frustrating. Why is there a difference?

Many products that require users to interact with them, such as smartphones and fitness trackers, have been designed primarily with the user in mind. They are generally easy and enjoyable to use. Others have not necessarily been designed with the users in mind; rather, they have been engineered primarily as software systems to perform set functions. An example is setting the time on a stove that requires a combination of button presses that are not obvious as to which ones to press together or separately. While they may work effectively, it can be at the expense of how easily they will be learned and therefore used in a real-world context.

Alan Cooper (2018), a well-known user experience (UX) guru, bemoans the fact that much of today's software suffers from the same interaction errors that were around 20 years ago. Why is this still the case, given that interaction design has been in existence for more than 25 years and that there are far more UX designers now in industry than ever before? He points out how many interfaces of new products do not adhere to the interaction design principles validated in the 1990s. For example, he notes that many apps do not follow even the most basic of UX principles, such as offering an “undo” option. He exclaims that it is “inexplicable and unforgivable that these violations continue to resurface in new products today.”

How can we rectify this situation so that the norm is that all new products are designed to provide good user experiences? To achieve this, we need to be able to understand how to reduce the negative aspects (such as frustration and annoyance) of the user experience while enhancing the positive ones (for example, enjoyment and efficacy). This entails developing interactive products that are easy, effective, and pleasurable to use from the users' perspective.

In this chapter, we begin by examining the basics of interaction design. We look at the difference between good and poor design, highlighting how products can differ radically in how usable and enjoyable they are. We then describe what and who is involved in the process of interaction design. The user experience, which is a central concern of interaction design, is then introduced. Finally, we outline how to characterize the user experience in terms of usability goals, user experience goals, and design principles. An in-depth activity is presented at the end of the chapter in which you have the opportunity to put into practice what you have read by evaluating the design of an interactive product.

1.2 Good and Poor Design

A central concern of interaction design is to develop interactive products that are usable. By this we mean products that are generally easy to learn, effective to use, and provide an enjoyable user experience. A good place to start thinking about how to design usable interactive products is to compare examples of well-designed and poorly designed ones. Through identifying the specific weaknesses and strengths of different interactive products, we can begin to

understand what it means for something to be usable or not. Here, we describe two examples of poorly designed products that have persisted over the years—a voice-mail system used in hotels and the ubiquitous remote control—and contrast these with two well-designed examples of the same products that perform the same function.

1.2.1 Voice-Mail System

Imagine the following scenario. You are staying at a hotel for a week while on a business trip. You see a blinking red light on the landline phone beside the bed. You are not sure what this means, so you pick up the handset. You listen to the tone and it goes “beep, beep, beep.” Maybe this means that there is a message for you. To find out how to access the message, you have to read a set of instructions next to the phone. You read and follow the first step:

1. Touch 41.

The system responds: “You have reached the Sunny Hotel voice message center. Please enter the room number for which you would like to leave a message.”

You wait to hear how to listen to a recorded message. But there are no further instructions from the phone. You look down at the instruction sheet again and read:

2. Touch *, your room number, and #.

You do so and the system replies: “You have reached the mailbox for room 106. To leave a message, type in your password.”

You type in the room number again, and the system replies: “Please enter room number again and then your password.”

You don’t know what your password is. You thought it was the same as your room number, but clearly it is not. At this point, you give up and call the front desk for help. The person at the desk explains the correct procedure for listening to messages. This involves typing in, at the appropriate times, the room number and the extension number of the phone (the latter is the password, which is different from the room number). Moreover, it takes six steps to access a message. You give up.

What is problematic with this voice-mail system?

- It is infuriating.
- It is confusing.
- It is inefficient, requiring you to carry out a number of steps for basic tasks.
- It is difficult to use.
- It has no means of letting you know at a glance whether any messages have been left or how many there are. You have to pick up the handset to find out and then go through a series of steps to listen to them.
- It is not obvious what to do: The instructions are provided partially by the system and partially by a card beside the phone.

Now compare it to the phone answering machine shown in Figure 1.1 The illustration shows a small sketch of a phone answering machine. Incoming messages are represented using marbles. The number of marbles that have moved into the pinball-like chute indicates the number of messages. Placing one of these marbles into a dent on the machine causes the recorded message to play. Dropping the same marble into a different dent on the phone dials the caller who left the message.




Figure 1.1 The marble answering machine

Source: Adapted from Crampton Smith (1995)

How does the marble answering machine differ from the voice-mail system?

- It uses familiar physical objects that indicate visually at a glance how many messages have been left.
- It is aesthetically pleasing and enjoyable to use.
- It requires only one-step actions to perform core tasks.
- It is a simple but elegant design.
- It offers less functionality and allows anyone to listen to any of the messages.

The marble answering machine is considered a design classic. It was created by Durrell Bishop while he was a student at the Royal College of Art in London (described by Crampton Smith, 1995). One of his goals was to design a messaging system that represented its basic functionality in terms of the behavior of everyday objects. To do this, he capitalized on people's everyday knowledge of how the physical world works. In particular, he made use of the ubiquitous everyday action of picking up a physical object and putting it down in another place.

This is an example of an interactive product designed with the users in mind. The focus is on providing them with a pleasurable experience but one that also makes efficient the activity of receiving messages. However, it is important to note that although the marble answering machine is an elegant and usable design, it would not be practical in a hotel setting. One of the main reasons is that it is not robust enough to be used in public places; for instance, the marbles could easily get lost or be taken as souvenirs. Also, the need to identify the user before allowing the messages to be played is essential in a hotel setting.

Therefore, when considering the design of an interactive product, it is important to consider where it is going to be used and who is going to use it. The marble answering machine would be more suitable in a home setting—provided that there were no children around who might be tempted to play with the marbles!

Video Durrell Bishop's answering machine: <http://vimeo.com/19930744>.

1.2.2 Remote Control

Every home entertainment system, be it the smart TV, set-top box, stereo system, and so forth, comes with its own remote control. Each one is different in terms of how it looks and works. Many have been designed with a dizzying array of small, multicolored, and double-labeled buttons (one on the button and one above or below it) that often seem arbitrarily positioned in relation to one another. Many viewers, especially when sitting in their living rooms, find it difficult to locate the right ones, even for the simplest of tasks, such as pausing or finding the main menu. It can be especially frustrating for those who need to put on their reading glasses each time to read the buttons. The remote control appears to have been put together very much as an afterthought.

In contrast, much effort and thought went into the design of the classic TiVo remote control with the user in mind (see Figure 1.2). TiVo is a digital video recorder that was originally developed to enable the viewer to record TV shows. The remote control was designed with large buttons that were clearly labeled and logically arranged, making them easy to locate and use in conjunction with the menu interface that appeared on the TV screen. In terms of its physical form, the remote device was designed to fit into the palm of a hand, having a peanut shape. It also has a playful look and feel about it: colorful buttons and cartoon icons are used that are distinctive, making it easy to identify them.



Figure 1.2 The TiVo remote control

Source: <https://business.tivo.com/>

How was it possible to create such a usable and appealing remote device where so many others have failed? The answer is simple: TiVo invested the time and effort to follow a user-centered design process. Specifically, TiVo's director of product design at the time involved potential users in the design process, getting their feedback on everything from the feel of the device in the hand to where best to place the batteries, making them easy to replace but not prone to falling out. He and his design team also resisted the trap of "buttonitis" to which so many other remote controls have fallen victim; that is one where buttons breed like rabbits—a button for every new function. They did this by restricting the number of control buttons embedded in the device to the essential ones. Other functions were then represented as part of the menu options and dialog boxes displayed on the TV screen, which could then be selected via the core set of physical control buttons. The result was a highly usable and pleasing device that has received much praise and numerous design awards.

DILEMMA

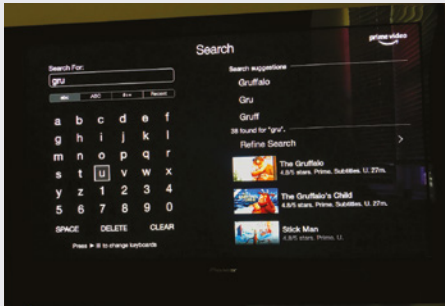
What Is the Best Way to Interact with a Smart TV?

A challenge facing smart TV providers is how to enable users to interact with online content. Viewers can select a whole range of content via their TV screens, but it involves scrolling through lots of menus and screens. In many ways, the TV interface has become more like a computer interface. This raises the question of whether the remote control is the best input device to use for someone who sits on a sofa or chair that is some distance from the wide TV screen. Smart TV developers have addressed this challenge in a number of ways.

An early approach was to provide an on-screen keyboard and numeric keypad that presented a grid of alphanumeric characters (see Figure 1.3a), which were selected by pressing a button repeatedly on a remote control. However, entering the name of a movie or an email address and password using this method can be painstakingly slow; it is also easy to overshoot and select the wrong letter or number when holding a button down on the remote to reach a target character.

More recent remote controls, such as those provided by Apple TV, incorporate a touchpad to enable swiping akin to the control commonly found on laptops. While this form of touch control expedites skipping through a set of letters displayed on a TV screen, it does not make it any easier to type in an email address and password. Each letter, number, or special character still has to be selected. Swiping is also prone to overshooting when aiming for a target letter, number, or character. Instead of providing a grid, the Apple TV interface displays two single lines of letters, numbers, and special characters to swipe across (see Figure 1.3b). While this can make it quicker for someone to reach a character, it is still tedious to select a sequence of characters in this way. For example, if you select a Y and the next letter is an A, you have to swipe all the way back to the beginning of the alphabet.

(a)



(b)

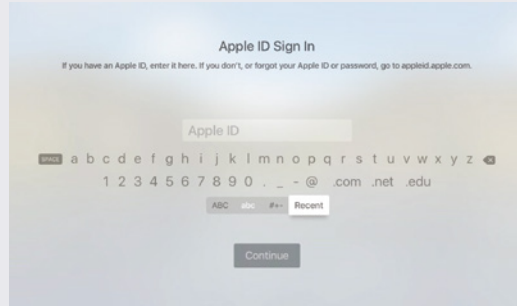


Figure 1.3 Typing on a TV screen (a) by selecting letters and numbers from a square matrix and (b) by swiping along a single line of letters and numbers

Source: (b) <https://support.apple.com/en-us/HT200107>

Might there be a better way to interact with a smart TV while sitting on the sofa? An alternative is to use voice control. Remote controls, like Siri or TiVo, for example, have a speech button that when pressed allows viewers to ask for movies by name or more generally by category, for instance, “What are the best sci-fi movies on Netflix?” Smart speakers, such as Amazon Echo, can also be connected to a smart TV via an HDMI port, and, similarly, the user can ask for something general or more specific, for example, “Alexa, play Big Bang Theory, Season 6, Episode 5, on the TV.” On recognizing the command, it will switch on the TV, switch to the right HDMI channel, open Netflix, and begin streaming the specific episode. Some TV content, however, requires the viewer to say that they are over a certain age by checking a box on the TV display. If the TV could ask the viewer and check that they are over 18, then that would be really smart! Also, if the TV needs the viewer to provide a password to access on-demand content, they won’t want to say it out aloud, character by character, especially in front of others who might also be in the room with them. The use of biometrics, then, may be the answer. ■

1.2.1 What to Design

Designing interactive products requires considering who is going to be using them, how they are going to be used, and where they are going to be used. Another key concern is to understand the kind of activities people are doing when interacting with these products. The appropriateness of different kinds of interfaces and arrangements of input and output devices depends on what kinds of activities are to be supported. For example, if the activity is to enable people to bank online, then an interface that is secure, trustworthy, and easy to navigate is essential. In addition, an interface that allows the user to find out information about new services offered by the bank without it being intrusive would be useful.

The world is becoming suffused with technologies that support increasingly diverse activities. Just think for a minute about what you can currently do using digital technology: send messages, gather information, write essays, control power plants, program, draw, plan, calculate, monitor others, and play games—just to name but a few. Now think about the types of interfaces and interactive devices that are available. They too are equally diverse: multitouch displays, speech-based systems, handheld devices, wearables, and large interactive displays—again, to name but a few. There are also many ways of designing how users can interact with a system, for instance, via the use of menus, commands, forms, icons, gestures, and so on. Furthermore, ever more innovative everyday artifacts are being created using novel materials, such as e-textiles and wearables (see Figure 1.4).



Figure 1.4 Turn signal biking jacket using e-textiles developed by Leah Beuchley

Source: Used courtesy of Leah Buechley

The Internet of Things (IoT) now means that many products and sensors can be connected to each other via the Internet, which enables them to talk to each other. Popular household IoT-enabled products include smart heating and lighting and home security systems where users can change the controls from an app on their phone or check out who is knocking on their door via a doorbell webcam. Other apps that are being developed are meant to make life easier for people, like finding a car parking space in busy areas.

The interfaces for everyday consumer items, such as cameras, microwave ovens, toasters, and washing machines, which used to be physical and the realm of product design, are now predominantly digitally based, requiring interaction design (called consumer electronics). The move toward transforming human-human transactions into solely interface-based ones has also introduced a new kind of customer interaction. Self-checkouts at grocery stores and libraries are now the norm where it is commonplace for customers to check out their own goods or books themselves, and at airports, where passengers check in their own luggage. While more cost-effective and efficient, it is impersonal and puts the onus on the person to interact with the system. Furthermore, accidentally pressing the wrong button or standing in the wrong place at a self-service checkout can result in a frustrating, and sometimes mortifying, experience.

What this all amounts to is a multitude of choices and decisions that interaction designers have to make for an ever-increasing range of products. A key question for interaction design is this: “How do you optimize the users’ interactions with a system, environment, or product so that they support the users’ activities in effective, useful, usable and pleasurable ways?” One could use intuition and hope for the best. Alternatively, one can be more principled in deciding which choices to make by basing them on an understanding of the users. This involves the following:

- Considering what people are good and bad at
- Considering what might help people with the way they currently do things
- Thinking through what might provide quality user experiences
- Listening to what people want and getting them involved in the design
- Using user-centered techniques during the design process

The aim of this book is to cover these aspects with the goal of showing you how to carry out interaction design. In particular, it focuses on how to identify users’ needs and the context of their activities. From this understanding, we move on to consider how to design usable, useful, and pleasurable interactive products.

1.3 What Is Interaction Design?

By interaction design, we mean the following:

Designing interactive products to support the way people communicate and interact in their everyday and working lives

Put another way, it is about creating user experiences that enhance and augment the way people work, communicate, and interact. More generally, Terry Winograd originally described it as “designing spaces for human communication and interaction” (1997, p. 160). John Thackara viewed it as “the why as well as the how of our daily interactions using computers” (2001, p. 50), while Dan Saffer emphasized its artistic aspects: “the art of facilitating interactions between humans through products and services” (2010, p. 4).

A number of terms have been used since to emphasize different aspects of what is being designed, including user interface design (UI), software design, user-centered design, product design, web design, user experience design, and interactive system design. Interaction design is generally used as the overarching term to describe the field, including its methods, theories, and approaches. UX is used more widely in industry to refer to the profession. However, the terms can be used interchangeably. Also, it depends on their ethos and brand.

1.3.1 The Components of Interaction Design

We view interaction design as fundamental to many disciplines, fields, and approaches that are concerned with researching and designing computer-based systems for people. Figure 1.5 presents the core ones along with interdisciplinary fields that comprise one or more of these, such as cognitive ergonomics. It can be confusing to try to work out the differences between them as many overlap. The main differences between interaction design and the other approaches referred to in the figure come largely down to which methods, philosophies, and lenses they use to study, analyze, and design products. Another way they vary is in terms of

the scope and problems they address. For example, information systems is concerned with the application of computing technology in domains such as business, health, and education, whereas ubiquitous computing is concerned with the design, development, and deployment of pervasive computing technologies (for example, IoT) and how they facilitate social interactions and human experiences.




Figure 1.5 Relationship among contributing academic disciplines, design practices, and interdisciplinary fields concerned with interaction design (double-headed arrows mean overlapping)

BOX 1.1

Is Interaction Design Beyond HCI?

We see the main difference between interaction design (ID) and human-computer interaction (HCI) as one of scope. Historically, HCI had a narrow focus on the design and usability of computing systems, while ID was seen as being broader, concerned with the theory, research, and practice of designing user experiences for all manner of technologies, systems, and products. That is one of the reasons why we chose to call our book *Interaction Design: beyond human-computer interaction*, to reflect this wider range. However, nowadays, HCI has greatly expanded in its scope (Churchill et al., 2013), so much so that it overlaps much more with ID (see Figure 1.6). ■



Figure 1.6 HCI out of the box: broadening its reach to cover more areas

1.3.2 Who Is Involved in Interaction Design?

Figure 1.5 also shows that many people are involved in performing interaction design, ranging from social scientists to movie-makers. This is not surprising given that technology has become such a pervasive part of our lives. But it can all seem rather bewildering to the onlooker. How does the mix of players work together?

Designers need to know many different things about users, technologies, and the interactions among them to create effective user experiences. At the least, they need to understand how people act and react to events and how they communicate and interact with each other. To be able to create engaging user experiences, they also need to understand how emotions work, what is meant by aesthetics, desirability, and the role of narrative in human experience. They also need to understand the business side, technical side, manufacturing side, and marketing side. Clearly, it is difficult for one person to be well versed in all of these diverse areas and also know how to apply the different forms of knowledge to the process of interaction design.

Interaction design is ideally carried out by multidisciplinary teams, where the skill sets of engineers, designers, programmers, psychologists, anthropologists, sociologists, marketing people, artists, toy makers, product managers, and others are drawn upon. It is rarely the case,

however, that a design team would have all of these professionals working together. Who to include in a team will depend on a number of factors, including a company's design philosophy, size, purpose, and product line.

One of the benefits of bringing together people with different backgrounds and training is the potential of many more ideas being generated, new methods developed, and more creative and original designs being produced. However, the downside is the costs involved. The more people there are with different backgrounds in a design team, the more difficult it can be to communicate and make progress with the designs being generated. Why? People with different backgrounds have different perspectives and ways of seeing and talking about the world. What one person values as important others may not even see (Kim, 1990). Similarly, a computer scientist's understanding of the term *representation* is often very different from that of a graphic designer or psychologist.

What this means in practice is that confusion, misunderstanding, and communication breakdowns can surface in a team. The various team members may have different ways of talking about design and may use the same terms to mean quite different things. Other problems can arise when a group of people who have not previously worked as a team are thrown together. For example, Aruna Balakrishnan et al. (2011) found that integration across different disciplines and expertise is difficult in many projects, especially when it comes to agreeing on and sharing tasks. The more disparate the team members—in terms of culture, background, and organizational structures—the more complex this is likely to be.

ACTIVITY 1.1

In practice, the makeup of a given design team depends on the kind of interactive product being built. Who do you think should be involved in developing

- A public kiosk providing information about the exhibits available in a science museum?
- An interactive educational website to accompany a TV series?

Comment

Ideally, each team will have a number of different people with different skill sets. For example, the first interactive product would include the following individuals:

- Graphic and interaction designers, museum curators, educational advisers, software engineers, software designers, and ergonomists

The second project would include these types of individuals:

- TV producers, graphic and interaction designers, teachers, video experts, software engineers, and software designers

In addition, as both systems are being developed for use by the general public, representative users, such as school children and parents, should be involved.

In practice, design teams often end up being quite large, especially if they are working on a big project to meet a fixed deadline. For example, it is common to find teams of 15 or more people working on a new product like a health app. This means that a number of people from each area of expertise are likely to be working as part of the project team. ■

1.3.3 Interaction Design Consultancies

Interaction design is now widespread in product and services development. In particular, website consultants and the computing industries have realized its pivotal role in successful interactive products. But it is not just IT companies that are realizing the benefits of having UXers on board. Financial services, retail, governments, and the public sector have realized too the value of interaction design. The presence or absence of good interaction design can make or break a company. Getting noticed in the highly competitive field of web products requires standing out. Being able to demonstrate that your product is easy, effective, and engaging to use is seen as central to this. Marketing departments are also realizing how branding, the number of hits, the customer return rate, and customer satisfaction are greatly affected by the usability of a website.

There are many interaction design consultancies now. These include established companies, such as Cooper, NielsenNorman Group, and IDEO, and more recent ones that specialize in a particular area, such as job board software (for example, Madgex), digital media (think of Cogapp), or mobile design (such as CXpartners). Smaller consultancies, such as Bunnyfoot and Dovetailed, promote diversity, interdisciplinarity, and scientific user research, having psychologists, researchers, interaction designers, usability, and customer experience specialists on board.

Many UX consultancies have impressive websites, providing case studies, tools, and blogs. For example, Holition publishes an annual glossy booklet as part of its UX Series (Javornik et al., 2017) to disseminate the outcomes of their in-house research to the wider community, with a focus on the implications for commercial and cultural aspects. This sharing of UX knowledge enables them to contribute to the discussion about the role of technology in the user experience.

1.4 The User Experience

The user experience refers to how a product behaves and is used by people in the real world. Jakob Nielsen and Don Norman (2014) define it as encompassing “all aspects of the end-user’s interaction with the company, its services, and its products.” As stressed by Jesse Garrett (2010, p. 10), “Every product that is used by someone has a user experience: newspapers, ketchup bottles, reclining armchairs, cardigan sweaters.” More specifically, it is about how people feel about a product and their pleasure and satisfaction when using it, looking at it, holding it, and opening or closing it. It includes their overall impression of how good it is to use, right down to the sensual effect small details have on them, such as how smoothly a switch rotates or the sound of a click and the touch of a button when pressing it. An important aspect is the quality of the experience someone has, be it a quick one, such as taking a photo; a leisurely one, such as playing with an interactive toy; or an integrated one, such as visiting a museum (Law et al., 2009).

It is important to point out that one cannot design a user experience, only design *for* a user experience. In particular, one cannot design a sensual experience, but only create the design features that can evoke it. For example, the outside case of a smartphone can be designed to be smooth, silky, and fit in the palm of a hand; when held, touched, looked at, and interacted with, that can provoke a sensual and satisfying user experience. Conversely, if it is designed to be heavy and awkward to hold, it is much more likely to end up providing a poor user experience—one that is uncomfortable and unpleasant.

Designers sometimes refer to UX as UXD. The addition of the *D* to UX is meant to encourage design thinking that focuses on the quality of the user experience rather than on the set of design methods to use (Allanwood and Beare, 2014). As Don Norman (2004) has stressed for many years, “It is not enough that we build products that function, that are understandable and usable, we also need to build joy and excitement, pleasure and fun, and yes, beauty to people’s lives.”

ACTIVITY 1.2

The iPod Phenomenon

Apple’s classic (and subsequent) generations of portable music players, called iPods, including the iPod Touch, Nano, and Shuffle, released during the early 2000s were a phenomenal success. Why do you think this occurred? Has there been any other product that has matched this quality of experience? With the exception of the iPod Touch, Apple stopped production of them in 2017. Playing music via a smartphone became the norm, superseding the need for a separate device.

Comment

Apple realized early on that successful interaction design involves creating interactive products that have a quality user experience. The sleek appearance of the iPod music player (see Figure 1.7), its simplicity of use, its elegance in style, its distinct family of rainbow colors, a novel interaction style that many people discovered was a sheer pleasure to learn and use, and the catchy naming of its product and content (iTunes, iPod), among many other design features, led to it becoming one of the greatest products of its kind and a must-have fashion item for teenagers, students, and adults alike. While there were many competing players on the market at the time—some with more powerful functionality, others that were cheaper and easier to use, or still others with bigger screens, more memory, and so forth—the quality of the overall user experience paled in comparison to that provided by the iPod.




Figure 1.7 The iPod Nano

Source: David Paul Morris / Getty Images

The nearest overall user experience that has all of the above is not so much for a product but for a physical store. The design of the Apple Store as a completely new customer experience for buying technology has been very successful in how it draws people in and what they do when browsing, discovering, and purchasing goods in the store. The products are laid out in a way to encourage interaction. ■

There are many aspects of the user experience that can be considered and many ways of taking them into account when designing interactive products. Of central importance are the usability, functionality, aesthetics, content, look and feel, and emotional appeal. In addition, Jack Carroll (2004) stresses other wide-reaching aspects, including fun, health, social capital (the social resources that develop and are maintained through social networks, shared values, goals, and norms), and cultural identity, such as age, ethnicity, race, disability, family status, occupation, and education.

Several researchers have attempted to describe the experiential aspect of a user experience. Kasper Hornbæk and Morten Hertzum (2017) note how it is often described in terms of the way that users perceive a product, such as whether a smartwatch is seen as sleek or chunky, and their emotional reaction to it, such as whether people have a positive experience when using it. Marc Hassenzahl's (2010) model of the user experience is the most well-known, where he conceptualizes it in terms of pragmatic and hedonic aspects. By pragmatic, it is meant how simple, practical, and obvious it is for the user to achieve their goals. By hedonic, it is meant how evocative and stimulating the interaction is to them. In addition to a person's perceptions of a product, John McCarthy and Peter Wright (2004) discuss the importance of their expectations and the way they make sense of their experiences when using technology. Their *Technology as Experience* framework accounts for the user experience largely in terms of how it is felt by the user. They recognize that defining experience is incredibly difficult because it is so nebulous and ever-present to us, just as swimming in water is to a fish. Nevertheless, they have tried to capture the essence of human experience by describing it in both holistic and metaphorical terms. These comprise a balance of sensual, cerebral, and emotional threads.

How does one go about producing quality user experiences? There is no secret sauce or magical formula that can be readily applied by interaction designers. However, there are numerous conceptual frameworks, tried and tested design methods, guidelines, and relevant research findings, which are described throughout the book.

1.5 Understanding Users

A main reason for having a better understanding of people in the contexts in which they live, work, and learn is that it can help designers understand how to design interactive products that provide good user experiences or match a user's needs. A collaborative planning tool for a space mission, intended to be used by teams of scientists working in different parts of the world, will have quite different needs from one targeted at customer and sales agents, to be used in a furniture store to draw up kitchen layout plans. Understanding individual

differences can also help designers appreciate that one size does not fit all; what works for one user group may be totally inappropriate for another. For example, children have different expectations than adults about how they want to learn or play. They may find having interactive quizzes and cartoon characters helping them along to be highly motivating, whereas most adults find them annoying. Conversely, adults often like talking-head discussions about topics, but children find them boring. Just as everyday objects like clothes, food, and games are designed differently for children, teenagers, and adults, so too should interactive products be designed for different kinds of users.

Learning more about people and what they do can also reveal incorrect assumptions that designers may have about particular user groups and what they need. For example, it is often assumed that because of deteriorating vision and dexterity, old people want things to be big—be it text or graphical elements appearing on a screen or the physical controls, like dials and switches, used to control devices. This may be true for some elderly people, but studies have shown that many people in their 70s, 80s, and older are perfectly capable of interacting with standard-size information and even small interfaces, for example, smartphones, just as well as those in their teens and 20s, even though, initially, some might think they will find it difficult (Siek et al., 2005). It is increasingly the case that as people get older, they do not like to consider themselves as lacking in cognitive and manual skills. Being aware of people’s sensitivities, such as aging, is as important as knowing how to design for their capabilities (Johnson and Finn, 2017). In particular, while many older adults now feel comfortable with and use a range of technologies (for instance, email, online shopping, online games, or social media), they may resist adopting new technologies. This is not because they don’t perceive them as being useful to their lives but because they don’t want to waste their time getting caught up by the distractions that digital life brings (Knowles and Hanson, 2018), for example, not wanting to be “glued to one’s mobile phone” like younger generations.

Being aware of cultural differences is also an important concern for interaction design, particularly for products intended for a diverse range of user groups from different countries. An example of a cultural difference is the dates and times used in different countries. In the United States, for example, the date is written as month, day, year (05/21/20), whereas in other countries, it is written in the sequence of day, month, year (21/05/20). This can cause problems for designers when deciding on the format of online forms, especially if intended for global use. It is also a concern for products that have time as a function, such as operating systems, digital clocks, or car dashboards. To which cultural group do they give preference? How do they alert users to the format that is set as default? This raises the question of how easily an interface designed for one user group can be used and accepted by another. Why is it that certain products, like a fitness tracker, are universally accepted by people from all parts of the world, whereas websites are designed differently and reacted to differently by people from different cultures?

To understand more about users, we have included three chapters (Chapters 4–6) that explain in detail how people act and interact with one another, with information, and with various technologies, together with describing their abilities, emotions, needs, desires, and what causes them to get annoyed, frustrated, lose patience, and get bored. We draw upon relevant psychological theory and social science research. Such knowledge enables designers to determine which solutions to choose from the many design alternatives available and how to develop and test these further.

1.6 Accessibility and Inclusiveness

Accessibility refers to the extent to which an interactive product is accessible by as many people as possible. Companies like Google and Apple provide tools for their developers to promote this. The focus is on people with disabilities. For example, Android OS provides a range of tools for those with disabilities, such as hearing aid compatibility to a built-in screen reader, while Apple VoiceOver lets the user know what's happening on its devices, so they can easily navigate and even know who is in a selfie just taken, by listening to the phone. Inclusiveness means being fair, open, and equal to everyone. Inclusive design is an overarching approach where designers strive to make their products and services accommodate the widest possible number of people. An example is ensuring that smartphones are being designed for all and made available to everyone—regardless of their disability, education, age, or income.

Whether or not a person is considered to be disabled changes over time with age, or as recovery from an accident progresses throughout their life. In addition, the severity and impact of an impairment can vary over the course of a day or in different environmental conditions. Disability can result because technologies are often designed in such a way as to necessitate a certain type of interaction that is impossible for someone with an impairment. Disability in this context is viewed as the result of poor interaction design between a user and the technology, not the impairment alone. Accessibility, on the other hand, opens up experiences so that they are accessible to all. Technologies that are now mainstream once started out as solutions to accessibility challenges. For example, SMS was designed for hearing-impaired people before it became a mainstream technology. Furthermore, designing for accessibility inherently results in inclusive design for all.

Accessibility can be achieved in two ways: first, through the inclusive design of technology, and second, through the design of assistive technology. When designing for accessibility, it is essential to understand the types of impairments that can lead to disability as they come in many forms. They are often classified by the type of impairment, for example:

- Sensory impairment (such as loss of vision or hearing)
- Physical impairment (having loss of functions to one or more parts of the body, for example, after a stroke or spinal cord injury)
- Cognitive (for instance, learning impairment or loss of memory/cognitive function due to old age or a condition such as Alzheimer's disease)

Within each type is a complex mix of people and capabilities. For example, a person might have only peripheral vision, be color blind, or have no light perception (and be registered blind). All are forms of visual impairment, and all require different design approaches. Color blindness can be overcome by an inclusive design approach. Designers can choose colors that will appear as separate colors to everyone. However, peripheral vision loss or complete blindness will often need an assistive technology to be designed.

Impairment can also be categorized as follows:

- Permanent (for example, long-term wheelchair user)
- Temporary (such as after an accident or illness)
- Situational (for instance, a noisy environment means a person can't hear)

The number of people living with permanent disability increases with age. Fewer than 20 percent of people are born with a disability, whereas 80 percent of people will have a disability once they reach 85. As people age, their functional abilities diminish. For example, people older than 50 often find it difficult to hear conversations in rooms with hard surfaces and lots of background noise. This is a disability that will come to most of us at some point.

People with permanent disabilities often use assistive technology in their everyday life, which they consider to be life-essential and an extension of their self (Holloway and Dawes, 2016). Examples include wheelchairs (people now refer to “wearing their wheels,” rather than “using a wheelchair”) and augmented and alternative communication aids. Much current HCI research into disability explores how new technologies, such as IoT, wearables, and virtual reality, can be used to improve upon existing assistive technologies.

Aimee Mullens is an athlete, actor, and fashion model who has shown how prosthetics can be designed to move beyond being purely functional (and often ugly) to being desirable and highly fashionable. She became a bilateral amputee when her legs were amputated below the knee as a one-year-old. She has done much to blur the boundary between disabled and nondisabled people, and she uses fashion as a tool to achieve this. Several prosthetic companies now incorporate fashion design into their products, including striking leg covers that are affordable by all (see Figure 1.8).



Figure 1.8 Fashionable leg cover designed by Alleles Design Studio

Source: <https://alleles.ca/>. Used courtesy of Alison Andersen

1.7 Usability and User Experience Goals

Part of the process of understanding users is to be clear about the primary objective of developing an interactive product for them. Is it to design an efficient system that will allow them to be highly productive in their work? Is it to design a learning tool that will be challenging and motivating? Or, is it something else? To help identify the objectives, we suggest classifying them in terms of usability and user experience goals. Traditionally, usability goals are concerned with meeting specific usability criteria, such as efficiency, whereas user experience goals are concerned with explicating the nature of the user experience, for instance, to be aesthetically pleasing. It is important to note, however, that the distinction between the two types of goals is not clear-cut since usability is often fundamental to the quality of the user experience and, conversely, aspects of the user experience, such as how it feels and looks, are inextricably linked with how usable the product is. We distinguish between them here to help clarify their roles but stress the importance of considering them together when designing for a user experience. Also, historically HCI was concerned primarily with usability, but it has since become concerned with understanding, designing for, and evaluating a wider range of user experience aspects.

1.7.1 Usability Goals

Usability refers to ensuring that interactive products are easy to learn, effective to use, and enjoyable from the user's perspective. It involves optimizing the interactions people have with interactive products to enable them to carry out their activities at work, at school, and in their everyday lives. More specifically, usability is broken down into the following six goals:

- Effective to use (effectiveness)
- Efficient to use (efficiency)
- Safe to use (safety)
- Having good utility (utility)
- Easy to learn (learnability)
- Easy to remember how to use (memorability)

Usability goals are typically operationalized as questions. The purpose is to provide the interaction designer with a concrete means of assessing various aspects of an interactive product and the user experience. Through answering the questions, designers can be alerted very early on in the design process to potential design problems and conflicts that they might not have considered. However, simply asking “Is the system easy to learn?” is not going to be very helpful. Asking about the usability of a product in a more detailed way—for example, “How long will it take a user to figure out how to use the most basic functions for a new smartwatch; how much can they capitalize on from their prior experience; and how long would it take the user to learn the whole set of functions?”—will elicit far more information.

The following are descriptions of the usability goals and a question for each one:

- (i) *Effectiveness* is a general goal, and it refers to how good a product is at doing what it is supposed to do.

Question: Is the product capable of allowing people to learn, carry out their work efficiently, access the information that they need, or buy the goods that they want?

- (ii) *Efficiency* refers to the way a product supports users in carrying out their tasks. The marble answering machine described earlier in this chapter was considered efficient in that it let the user carry out common tasks, for example, listening to messages, through a minimal number of steps. In contrast, the voice-mail system was considered inefficient because it required the user to carry out many steps and learn an arbitrary set of sequences for the same common task. This implies that an efficient way of supporting common tasks is to let the user use single button or key presses. An example of where this kind of efficiency mechanism has been employed effectively is in online shopping. Once users have entered all of the necessary personal details in an online form to make a purchase, they can let the website save all of their personal details. Then, if they want to make another purchase at that site, they don't have to re-enter all of their personal details. A highly successful mechanism patented by Amazon.com is the one-click option, which requires users to click only a single button when they want to make another purchase.

Question: Once users have learned how to use a product to carry out their tasks, can they sustain a high level of productivity?

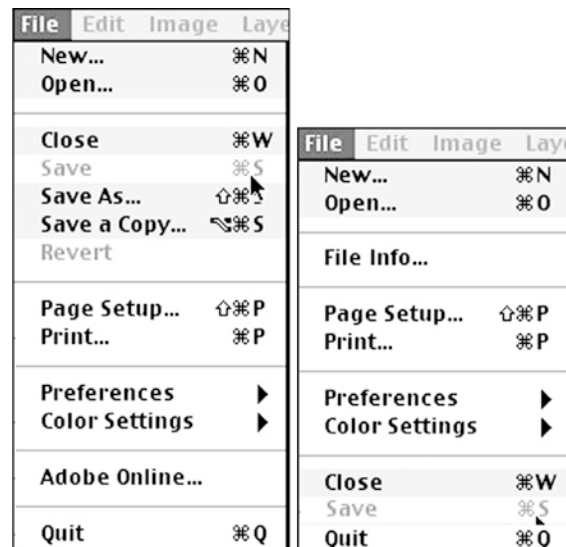
- (iii) *Safety* involves protecting the user from dangerous conditions and undesirable situations. In relation to the first ergonomic aspect, it refers to the external conditions where people work. For example, where there are hazardous conditions—such as X-ray machines or toxic chemicals—operators should be able to interact with and control computer-based systems remotely. The second aspect refers to helping any kind of user in any kind of situation to avoid the dangers of carrying out unwanted actions accidentally. It also refers to the perceived fears that users might have of the consequences of making errors and how this affects their behavior. Making interactive products safer in this sense involves (1) preventing the user from making serious errors by reducing the risk of wrong keys/buttons being mistakenly activated (an example is not placing the quit or delete-file command right next to the save command on a menu) and (2) providing users with various means of recovery should they make errors, such as an undo function. Safe interactive systems should engender confidence and allow the user the opportunity to explore the interface to try new operations (see Figure 1.9a). Another safety mechanism is confirming dialog boxes that give users another chance to consider their intentions (a well-known example is the appearance of a dialog box after issuing the command to delete everything in the trash, saying: “Are you sure you want to remove the items in the Trash permanently?”) (see Figure 1.9b).

Question: What is the range of errors that are possible using the product, and what measures are there to permit users to recover easily from them?

- (iv) *Utility* refers to the extent to which the product provides the right kind of functionality so that users can do what they need or want to do. An example of a product with high utility is an accounting software package that provides a powerful computational tool that accountants can use to work out tax returns. An example of a product with low utility is a software drawing tool that does not allow users to draw freehand but forces them to use a mouse to create their drawings, using only polygon shapes.

Question: Does the product provide an appropriate set of functions that will enable users to carry out all of their tasks in the way they want to do them?

- (v) *Learnability* refers to how easy a system is to learn to use. It is well known that people don't like spending a long time learning how to use a system. They want to get started right away and become competent at carrying out tasks without too much effort. This is



(a)



(b)

Figure 1.9 (a) A safe and unsafe menu. Which is which and why? (b) A warning dialog box for Mac OS X

especially true for interactive products intended for everyday use (for example social media, email, or a GPS) and those used only infrequently (for instance, online tax forms). To a certain extent, people are prepared to spend a longer time learning more complex systems that provide a wider range of functionality, such as web authoring tools. In these situations, pop-up tutorials can help by providing contextualized step-by-step material with hands-on exercises. A key concern is determining how much time users are prepared to spend learning a product. It seems like a waste if a product provides a range of functionality that the majority of users are unable or unprepared to spend the time learning how to use.

Question: Is it possible for the user to work out how to use the product by exploring the interface and trying certain actions? How hard will it be to learn the whole set of functions in this way?

- (vi) *Memorability* refers to how easy a product is to remember how to use, once learned. This is especially important for interactive products that are used infrequently. If users haven't used an operation for a few months or longer, they should be able to remember or at least rapidly be reminded how to use it. Users shouldn't have to keep relearning how to carry

out tasks. Unfortunately, this tends to happen when the operations required to be learned are obscure, illogical, or poorly sequenced. Users need to be helped to remember how to do tasks. There are many ways of designing the interaction to support this. For example, users can be helped to remember the sequence of operations at different stages of a task through contextualized icons, meaningful command names, and menu options. Also, structuring options and icons so that they are placed in relevant categories of options, for example, placing all of the drawing tools in the same place on the screen, can help the user remember where to look to find a particular tool at a given stage of a task.

Question: What types of interface support have been provided to help users remember how to carry out tasks, especially for products and operations they use infrequently?

In addition to couching usability goals in terms of specific questions, they are turned into usability criteria. These are specific objectives that enable the usability of a product to be assessed in terms of how it can improve (or not improve) a user's performance. Examples of commonly used usability criteria are time to complete a task (efficiency), time to learn a task (learnability), and the number of errors made when carrying out a given task over time (memorability). These can provide quantitative indicators of the extent to which productivity has increased, or how work, training, or learning have been improved. They are also useful for measuring the extent to which personal, public, and home-based products support leisure and information gathering activities. However, they do not address the overall quality of the user experience, which is where user experience goals come into play.

1.7.2 User Experience Goals

A diversity of user experience goals has been articulated in interaction design, which covers a range of emotions and felt experiences. These include desirable and undesirable ones, as shown in Table 1.1.

| | | |
|----------------------------|-------------------------|------------------------|
| Desirable aspects | | |
| Satisfying | Helpful | Fun |
| Enjoyable | Motivating | Provocative |
| Engaging | Challenging | Surprising |
| Pleasurable | Enhancing sociability | Rewarding |
| Exciting | Supporting creativity | Emotionally fulfilling |
| Entertaining | Cognitively stimulating | Experiencing flow |
| Undesirable aspects | | |
| Boring | Unpleasant | |
| Frustrating | Patronizing | |
| Making one feel guilty | Making one feel stupid | |
| Annoying | Cutesy | |
| Childish | Gimmicky | |

Table 1.1 Desirable and undesirable aspects of the user experience

Many of these are subjective qualities and are concerned with how a system feels to a user. They differ from the more objective usability goals in that they are concerned with how users experience an interactive product from their perspective, rather than assessing how useful or productive a system is from its own perspective. Whereas the terms used to describe usability goals comprise a small distinct set, many more terms are used to describe the multifaceted nature of the user experience. They also overlap with what they are referring to. In so doing, they offer subtly different options for expressing the way an experience varies for the same activity over time, technology, and place. For example, we may describe listening to music in the shower as highly pleasurable, but consider it more apt to describe listening to music in the car as enjoyable. Similarly, listening to music on a high-end powerful music system may invoke exciting and emotionally fulfilling feelings, while listening to it on a smartphone that has a shuffle mode may be serendipitously enjoyable, especially not knowing what tune is next. The process of selecting terms that best convey a user's feelings, state of being, emotions, sensations, and so forth when using or interacting with a product at a given time and place can help designers understand the multifaceted and changing nature of the user experience.

The concepts can be further defined in terms of elements that contribute to making a user experience pleasurable, fun, exciting, and so on. They include attention, pace, play, interactivity, conscious and unconscious control, style of narrative, and flow. The concept of flow (Csikszentmihalyi, 1997) is popular in interaction design for informing the design of user experiences for websites, video games, and other interactive products. It refers to a state of intense emotional involvement that comes from being completely involved in an activity, like playing music, and where time flies. Instead of designing web interfaces to cater to visitors who know what they want, they can be designed to induce a state of flow, leading the visitor to some unexpected place, where they become completely absorbed. In an interview with *Wired* magazine, Mihaly Csikszentmihalyi (1996) uses the analogy of a gourmet meal to describe how a user experience can be designed to be engrossing, “starting off with the appetizers, moving on to the salads and entrées, and building toward dessert and not knowing what will follow.”

The quality of the user experience may also be affected by single actions performed at an interface. For example, people can get much pleasure from turning a knob that has the perfect level of gliding resistance; they may enjoy flicking their finger from the bottom of a smartphone screen to reveal a new menu, with the effect that it appears by magic, or enjoy the sound of trash being emptied from the trashcan on a screen. These one-off actions can be performed infrequently or several times a day—which the user never tires of doing. Dan Saffer (2014) has described these as *micro-interactions* and argues that designing these moments of interaction at the interface—despite being small—can have a big impact on the user experience.

ACTIVITY 1.3

There are more desirable than undesirable aspects of the user experience listed in Table 1.1. Why do you think this is so? Should you consider all of these when designing a product?

(Continued)

Comment

The two lists we have come up with are not meant to be exhaustive. There are likely to be more—both desirable and undesirable—as new products surface. The reason for there being more of the former is that a primary goal of interaction design is to create positive experiences. There are many ways of achieving this.

Not all usability and user experience goals will be relevant to the design and evaluation of an interactive product being developed. Some combinations will also be incompatible. For example, it may not be possible or desirable to design a process control system that is both safe and fun. Recognizing and understanding the nature of the relationship between usability and user experience goals is central to interaction design. It enables designers to become aware of the consequences of pursuing different combinations when designing products and highlighting potential trade-offs and conflicts. As suggested by Jack Carroll (2004), articulating the interactions of the various components of the user's experience can lead to a deeper and more significant interpretation of the role of each component. ■

BOX 1.3**Beyond Usability: Designing to Persuade**

Eric Schaffer (2009) argues that we should be focusing more on the user experience and less on usability. He points out how many websites are designed to persuade or influence rather than enable users to perform their tasks in an efficient manner. For example, many online shopping sites are in the business of selling services and products, where a core strategy is to entice people to buy what they might not have thought they needed. Online shopping experiences are increasingly about persuading people to buy rather than being designed to make shopping easy. This involves designing for persuasion, emotion, and trust, which may or may not be compatible with usability goals.

This entails determining what customers will do, whether it is to buy a product or renew a membership, and it involves encouraging, suggesting, or reminding the user of things that they might like or need. Many online travel sites try to lure visitors to purchase additional items (such as hotels, insurance, car rental, car parking, or day trips) besides the flight they originally wanted to book, and they will add a list full of tempting graphics to the visitor's booking form, which then has to be scrolled through before being able to complete the transaction. These opportunities need to be designed to be eye-catching and enjoyable, in the same way that an array of products are attractively laid out in the aisles of a grocery store that one is required to walk past before reaching one's desired product.

Some online sites, however, have gone too far, for example, adding items to the customer's shopping basket (for example, insurance, special delivery, and care and handling) that the shopper has to deselect if not desired or start all over again. This sneaky add-on approach can often result in a negative experience. More generally, this deceptive approach

to UX has been described by Harry Brignull as *dark patterns* (see <http://darkpatterns.org/>). Shoppers often become annoyed if they notice decisions that add cost to their purchase have been made on their behalf without even being asked. For example, on clicking the unsubscribe button on the website of a car rental company, as indicated in Figure 1.10, the user is taken to another page where they have to uncheck additional boxes and then Update. They are then taken to yet another page where they are asked for their reason. The next screen says “Your email preferences have been updated. Do you need to hire a vehicle?” without letting the user know whether they have been unsubscribed from their mailing list.

The figure consists of two screenshots of an email preferences form. The top screenshot shows a form titled "Email preferences" with a text input field containing the email address "y.rogers@ucl.ac.uk". Below the input field, there is a horizontal line and the text "Uncheck the emails you do not want to receive". Underneath, there are three checkboxes, all of which are checked: "Newsletters UK", "NiftyCars Partners offers", and "About your rental". A yellow "Update" button is positioned below these options. At the bottom left of the form, there is a small asterisk and the text "* required fields". The bottom screenshot also shows a form titled "Email preferences". It features the text "We'd love to get some feedback on why you're unsubscribing." followed by four radio button options: "Emails were too frequent", "Emails were not relevant", "I am no longer interested in this content", and "I never signed up for newsletters from NiftyCars". A yellow "Update" button is located below these options.

Figure 1.10 Dark pattern for a car rental company

(Continued)

The key is to nudge people in subtle and pleasant ways with which they can trust and feel comfortable. Natasha Loma (2018) points out how dark pattern design is “deception and dishonesty by design.” She describes in a TechCrunch article the many kinds of dark patterns that are now used to deceive users. A well-known example that most of us have experienced is unsubscribing from a marketing mailing list. Many sites go to great lengths to make it difficult for you to leave; you think you have unsubscribed, but then you discover that you need to type in your email address and click several more buttons to reaffirm that you really want to quit. Then, just when you think you are safe, they post a survey asking you to answer a few questions about why you want to leave. Like Harry Brignull, she argues that companies should adopt fair and ethical design where users have to opt in to any actions that benefit the company at the expense of the users’ interests. ■

1.7.3 Design Principles

Design principles are used by interaction designers to aid their thinking when designing for the user experience. These are generalizable abstractions intended to orient designers toward thinking about different aspects of their designs. A well-known example is feedback: Products should be designed to provide adequate feedback to the users that informs them about what has already been done so that they know what to do next in the interface. Another one that is important is findability (Morville, 2005). This refers to the degree to which a particular object is easy to discover or locate—be it navigating a website, moving through a building, or finding the delete image option on a digital camera. Related to this is the principle of navigability: Is it obvious what to do and where to go in an interface; are the menus structured in a way that allows the user to move smoothly through them to reach the option they want?

Design principles are derived from a mix of theory-based knowledge, experience, and common sense. They tend to be written in a prescriptive manner, suggesting to designers what to provide and what to avoid at the interface—if you like, the dos and don’ts of interaction design. More specifically, they are intended to help designers explain and improve their designs (Thimbleby, 1990). However, they are not intended to specify how to design an actual interface, for instance, telling the designer how to design a particular icon or how to structure a web portal, but to act more like triggers for designers, ensuring that they provide certain features in an interface.

A number of design principles have been promoted. The best known are concerned with how to determine what users should see and do when carrying out their tasks using an interactive product. Here we briefly describe the most common ones: visibility, feedback, constraints, consistency, and affordance.

Visibility

The importance of visibility is exemplified by our contrasting examples at the beginning of the chapter. The voice-mail system made the presence and number of waiting messages invisible, while the answering machine made both aspects highly visible. The more visible functions are, the more likely it is that users will be able to know what to do next. Don Norman (1988) describes the controls of a car to emphasize this point. The controls for different operations are clearly visible, such as indicators, headlights, horn, and hazard warning lights, indicating what

can be done. The relationship between the way the controls have been positioned in the car and what they do makes it easy for the driver to find the appropriate control for the task at hand.

In contrast, when functions are out of sight, it makes them more difficult to find and to know how to use. For example, devices and environments that have become automated through the use of sensor technology (usually for hygiene and energy-saving reasons)—like faucets, elevators, and lights—can sometimes be more difficult for people to know how to control, especially how to activate or deactivate them. This can result in people getting caught short and frustrated. Figure 1.11 shows a sign that explains how to use the automatically controlled faucet for what is normally an everyday and well-learned activity. It also states that the faucets cannot be operated if wearing black clothing. It does not explain, however, what to do if you are wearing black clothing! Increasingly, highly visible controlling devices, like knobs, buttons, and switches, which are intuitive to use, have been replaced by invisible and ambiguous activating zones where people have to guess where to move their hands, bodies, or feet—on, into, or in front of—to make them work.



Figure 1.11 A sign in the restrooms at the Cincinnati airport

Source: <http://www.baddesigns.com>

Feedback

Related to the concept of visibility is feedback. This is best illustrated by an analogy to what everyday life would be like without it. Imagine trying to play a guitar, slice bread using a knife, or write using a pen if none of the actions produced any effect for several seconds.

There would be an unbearable delay before the music was produced, the bread was cut, or the words appeared on the paper, making it almost impossible for the person to continue with the next strum, cut, or stroke.

Feedback involves sending back information about what action has been done and what has been accomplished, allowing the person to continue with the activity. Various kinds of feedback are available for interaction design—audio, tactile, verbal, visual, and combinations of these. Deciding which combinations are appropriate for different types of activities and interactivities is central. Using feedback in the right way can also provide the necessary visibility for user interaction.

Constraints

The design concept of *constraining* refers to determining ways of restricting the kinds of user interaction that can take place at a given moment. There are various ways that this can be achieved. A common design practice in graphical user interfaces is to deactivate certain menu options by shading them gray, thereby restricting the user only to actions permissible at that stage of the activity (see Figure 1.12). One of the advantages of this form of constraining is that it prevents the user from selecting incorrect options and thereby reduces the chance of making a mistake.

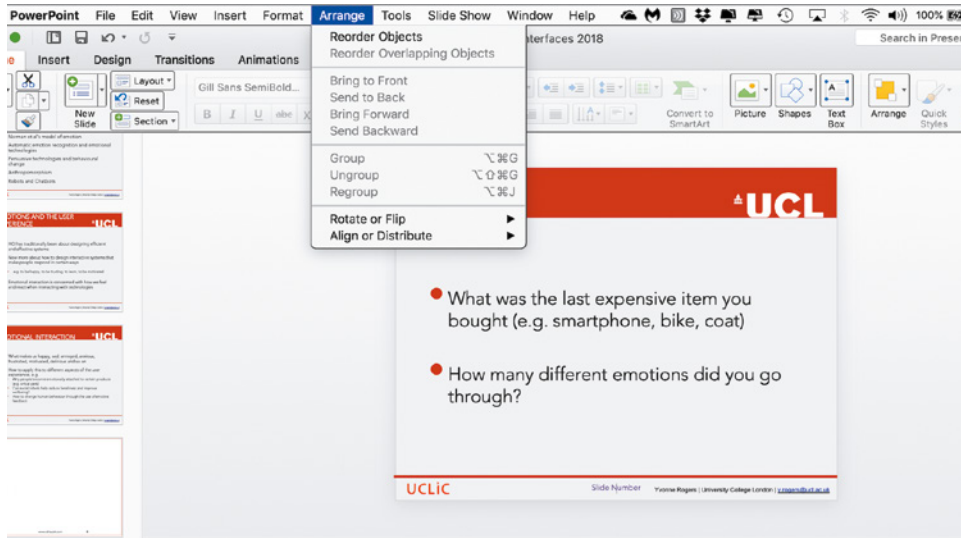


Figure 1.12 A menu showing restricted availability of options as an example of logical constraining. Gray text indicates deactivated options.

Source: <https://www.ucl.ac.uk>

The use of different kinds of graphical representations can also constrain a person's interpretation of a problem or information space. For example, flow chart diagrams show which objects are related to which, thereby constraining the way that the information can be perceived. The physical design of a device can also constrain how it is used; for example, the

external slots in a computer have been designed to allow a cable or card to be inserted in a certain way only. Sometimes, however, the physical constraint is ambiguous, as shown in Figure 1.13. The figure shows part of the back of a computer. There are two sets of connectors; the two on the right are for a mouse and a keyboard. They look identical and are physically constrained in the same way. How do you know which is which? Do the labels help?



Figure 1.13 Ambiguous constraints on the back of a computer

Source: <http://www.baddesigns.com>

Consistency

This refers to designing interfaces to have similar operations and use similar elements for achieving similar tasks. In particular, a consistent interface is one that follows rules, such as using the same operation to select all objects. For example, a consistent operation is using the same input action to highlight any graphical object on the interface, such as always clicking the left mouse button. Inconsistent interfaces, on the other hand, allow exceptions to a rule. An example is where certain graphical objects (for example, email messages presented in a table) can be highlighted only by using the right mouse button, while all other operations are highlighted using the left mouse button. The problem with this kind of inconsistency is that it is quite arbitrary, making it difficult for users to remember and making its use more prone to mistakes.

One of the benefits of consistent interfaces, therefore, is that they are easier to learn and use. Users have to learn only a single mode of operation that is applicable to all objects. This principle works well for simple interfaces with limited operations, such as a portable radio with a small number of operations mapped onto separate buttons. Here, all the user has to do is to learn what each button represents and select accordingly. However, it can be more problematic to apply the concept of consistency to more complex interfaces, especially when many different operations need to be designed. For example, consider how to design an interface for an application that offers hundreds of operations, such as a word-processing application. There is simply not enough space for a thousand buttons, each of which maps to an individual operation. Even if there were, it would be extremely difficult and time-consuming for the user to search through all of them to find the desired operation. A much more effective design solution is to create categories of commands that can be mapped into subsets of operations that can be displayed at the interface, for instance, via menus.

Affordance

This is a term used to refer to an attribute of an object that allows people to know how to use it. For example, a mouse button invites pushing (in so doing, activating clicking) by the way it is physically constrained in its plastic shell. At a simple level, to afford means “to give a clue” (Norman, 1988). When the affordances of a physical object are perceptually obvious, it is easy to know how to interact with it. For example, a door handle affords pulling, a cup handle affords grasping, and a mouse button affords pushing. The term has since been much popularized in interaction design, being used to describe how interfaces should make it obvious as to what can be done when using them. For example, graphical elements like buttons, icons, links, and scrollbars are discussed with respect to how to make it appear obvious how they should be used: icons should be designed to afford clicking, scrollbars to afford moving up and down, and buttons to afford pushing.

Don Norman (1999) suggests that there are two kinds of affordance: perceived and real. Physical objects are said to have real affordances, like grasping, that are perceptually obvious and do not have to be learned. In contrast, user interfaces that are screen-based are virtual and do not have these kinds of real affordances. Using this distinction, he argues that it does not make sense to try to design for real affordances at the interface, except when designing physical devices, like control consoles, where affordances like pulling and pressing are helpful in guiding the user to know what to do. Alternatively, screen-based interfaces are better conceptualized as perceived affordances, which are essentially learned conventions. However, watching a one-year-old swiping smartphone screens, zooming in and out on images with their finger and thumb, and touching menu options suggests that kind of learning comes naturally.

Applying Design Principles in Practice

One of the challenges of applying more than one of the design principles in interaction design is that trade-offs can arise among them. For example, the more you try to constrain an interface, the less visible information becomes. The same can also happen when trying to apply a single design principle. For example, the more an interface is designed to afford through trying to resemble the way physical objects look, the more it can become cluttered and difficult to use. It can also be the case that the more an interface is designed to be aesthetic, the less usable it becomes. Consistency can be a problematic design principle; trying to design an interface to be consistent with something can make it inconsistent with something else. Furthermore, sometimes inconsistent interfaces are actually easier to use than consistent interfaces. This is illustrated by Jonathan Grudin’s classic (1989) use of the analogy of where knives are stored in a house. Knives come in a variety of forms, including butter knives, steak knives, table knives, and fish knives. An easy place to put them all and subsequently locate them is in the top drawer by the sink. This makes it easy for everyone to find them and follows a simple consistent rule. But what about the knives that don’t fit or are too sharp to put in the drawer, like carving knives and bread knives? They are placed in a wooden block. And what about the best knives kept only for special occasions? They are placed in the cabinet in another room for safekeeping. And what about other knives like putty knives and paint-scraping knives used in home improvement projects (kept in the garage) and jack-knives (kept in one’s pockets or backpack)? Very quickly, the consistency rule begins to break down.

Jonathan Grudin notes how, in extending the number of places where knives are kept, inconsistency is introduced, which in turn increases the time needed to learn where they are all stored. However, the placement of the knives in different places often makes it easier to find them because they are at hand for the context in which they are used and are also next to the other objects used for a specific task; for instance, all of the home improvement project tools are stored together in a box in the garage. The same is true when designing interfaces: introducing inconsistency can make it more difficult to learn an interface, but in the long run it can make it easier to use.

ACTIVITY 1.4

One of the main design principles for website design is simplicity. Jakob Nielsen (1999) proposed that designers go through all of their design elements and remove them one by one. If a design works just as well without an element, then remove it. Do you think this is a good design principle? If you have your own website, try doing this and seeing what happens. At what point does the interaction break down?

Comment

Simplicity is certainly an important design principle. Many designers try to cram too much into a screenful of space, making it unwieldy for people to find the element in which they are interested. Removing design elements to see what can be discarded without affecting the overall function of the website can be a salutary lesson. Unnecessary icons, buttons, boxes, lines, graphics, shading, and text can be stripped, leaving a cleaner, crisper, and easier-to-navigate website. However, graphics, shading, coloring, and formatting can make a site aesthetically pleasing and enjoyable to use. Plain vanilla sites consisting solely of lists of text and a few links may not be as appealing and may put certain visitors off, never to return. Good interaction design involves getting the right balance between aesthetic appeal and the optimal amount and kind of information per page. ■

In-Depth Activity

This activity is intended for you to put into practice what you have studied in this chapter. Specifically, the objective is to enable you to define usability and user experience goals and to transform these and other design principles into specific questions to help evaluate an interactive product.

Find an everyday handheld device, for example, a remote control, digital camera, or smartphone and examine how it has been designed, paying particular attention to how the user is meant to interact with it.

(Continued)

- (a) From your first impressions, write down what is good and bad about the way the device works.
- (b) Give a description of the user experience resulting from interacting with it.
- (c) Outline some of the core micro-interactions that are supported by it. Are they pleasurable, easy, and obvious?
- (d) Based on your reading of this chapter and any other material you have come across about interaction design, compile a set of usability and user experience goals that you think will be most relevant in evaluating the device. Decide which are the most important ones and explain why.
- (e) Translate each of your sets of usability and user experience goals into two or three specific questions. Then use them to assess how well your device fares.
- (f) Repeat steps (c) and (d), but this time use the design principles outlined in the chapter.
- (g) Finally, discuss possible improvements to the interface based on the answers obtained in steps (d) and (e).

Summary

In this chapter, we have looked at what interaction design is and its importance when developing apps, products, services, and systems. To begin, a number of good and bad designs were presented to illustrate how interaction design can make a difference. We described who and what is involved in interaction design and the need to understand accessibility and inclusiveness. We explained in detail what usability and user experience are, how they have been characterized, and how to operationalize them to assess the quality of a user experience resulting from interacting with an interactive product. The increasing emphasis on designing for the user experience and not just products that are usable was stressed. A number of core design principles were also introduced that provide guidance for helping to inform the interaction design process.

Key Points

- Interaction design is concerned with designing interactive products to support the way people communicate and interact in their everyday and working lives.
- Interaction design is multidisciplinary, involving many inputs from wide-ranging disciplines and fields.
- The notion of the user experience is central to interaction design.
- Optimizing the interaction between users and interactive products requires consideration of a number of interdependent factors, including context of use, types of activity, UX goals, accessibility, cultural differences, and user groups.
- Identifying and specifying relevant usability and user experience goals can help lead to the design of good interactive products.
- Design principles, such as feedback and simplicity, are useful heuristics for informing, analyzing, and evaluating aspects of an interactive product.

Further Reading

Here we recommend a few seminal readings on interaction design and the user experience (in alphabetical order).


COOPER, A., REIMANN, R., CRONIN, D. AND NOESSEL, C. (2014) *About Face: The Essentials of Interaction Design* (4th ed.). John Wiley & Sons Inc. This fourth edition of *About Face* provides an updated overview of what is involved in interaction design, and it is written in a personable style that appeals to practitioners and students alike.

GARRETT, J. J. (2010) *The Elements of User Experience: User-Centered Design for the Web and Beyond* (2nd ed.). New Riders Press. This is the second edition of the popular coffee-table introductory book to interaction design. It focuses on how to ask the right questions when designing for a user experience. It emphasizes the importance of understanding how products work on the outside, that is, when a person comes into contact with those products and tries to work with them. It also considers a business perspective.

LIDWELL, W., HOLDEN, K. AND BUTLER, J. (2010) *Revised and Updated: 125 Ways to Enhance Usability, Influence Perception, Increase Appeal, Make Better Design Decisions and Teach Through Design*. Rockport Publishers, Inc. This book presents classic design principles such as consistency, accessibility, and visibility in addition to some lesser-known ones, such as constancy, chunking, and symmetry. They are alphabetically ordered (for easy reference) with a diversity of examples to illustrate how they work and can be used.

NORMAN, D.A. (2013) *The Design of Everyday Things: Revised and Expanded Edition*. MIT Press. This book was first published in 1988 and became an international best seller, introducing the world of technology to the importance of design and psychology. It covers the design of everyday things, such as refrigerators and thermostats, providing much food for thought in relation to how to design interfaces. This latest edition is comprehensively revised showing how principles from psychology apply to a diversity of old and new technologies. The book is highly accessible with many illustrative examples.

SAFFER, D. (2014) *Microinteractions: Designing with Details*. O'Reilly. This highly accessible book provides many examples of the small things in interaction design that make a big difference between a pleasant experience and a nightmare one. Dan Saffer describes how to design them to be efficient, understandable, and enjoyable user actions. He goes into detail about their structure and the different kinds, including many examples with lots of illustrations. The book is a joy to dip into and enables you to understand right away why and how it is important to get the micro-interactions right.



INTERVIEW with Harry Brignull

Harry Brignull is a user experience consultant based in the United Kingdom. He has a PhD in cognitive science, and his work involves building better experiences by blending user research and interaction design. In his work, Harry has consulted for companies including Spotify, Smart Pension, The Telegraph, British Airways, Vodafone, and many others. In his spare time, Harry also runs a blog on interaction design that has attracted a lot of eyeballs. It is called 90percentofeverything.com, and it is well worth checking out.

What are the characteristics of a good interaction designer?

I think of interaction design, user experience design, service design, and user research as a combined group of disciplines that are tricky to tease apart. Every company has slightly different terminology, processes, and approaches. I'll let you into a secret, though. They're all making it up as they go along. When you see any organization portraying its design and research publicly, they're showing you a fictionalized view of it for recruitment and marketing purposes. The reality of the work is usually very different. Research and design is naturally messy. There's a lot of waste, false assumptions, and blind alleys you have to go down before you can define and understand a problem well enough to solve it. If an employer doesn't understand this and they don't give you the space and time you need, then you won't be able to do a good job, regardless of your skills and training.

A good interaction designer has skills that work like expanding foam. You expand to fill the skill gaps in your team. If you don't have a writer present, you need to be able to step up and do it yourself, at least to the level of a credible draft. If you don't have a researcher, you'll need to step up and do it yourself. The same goes for developing code-based prototypes, planning the user journeys, and so on. You'll soon learn to become used to working outside of your comfort zone and relish the new challenges that each project brings.

How has interaction design changed in the past few years?

In-housing of design teams is a big trend at the moment. When I started my consultancy career in the mid-2000s, the main route to getting a career in industry was to get a role at an agency, like a UX consultancy, a research agency, or a full-service agency. Big organizations didn't even know where to start with hiring and building their own teams, so they paid enormous sums to agencies to design and build their products. This turned out to be a pretty ineffective model—when the agencies finish a project, they take all the acquired expertise away with them to their next clients.

These days, digital organizations have wised up, and they've started building their own in-house teams. This means that a big theme in design these days is organizational change. You can't do good design in an organization that isn't set up for it. In fact, in old, large organizations, the political

structure often seems to be set up to sabotage good design and development practices. It sounds crazy, but it's very common to walk into an organization to find a project manager brandishing a waterfall Gantt chart while ranting obsessively about Agile (which is a contradiction in terms) or to find a product owner saying in one breath they value user research yet in the next breath getting angry with researchers for bringing them bad news. As well as "legacy technology," organizations naturally end up with "legacy thinking." It's really tricky to change it. Design used to be just a department. Nowadays it's understood that good design requires the entire organization to work together in a cohesive way.

What projects are you working on now?

I'm currently head of UX at a FinTech startup called Smart Pension in London. Pensions pose a really fascinating user-centered design challenge. Consumers hate thinking about pensions, but they desperately need them. In a recent research session, one of the participants said something that really stuck with me: "Planning your pension is like planning for your own funeral." Humans are pretty terrible at long-term planning over multiple decades. Nobody likes to think about their own mortality. But this is exactly what you need to do if you want to have a happy retirement.

The pension industry is full of jargon and off-putting technical complexity. Even fundamental financial concepts like *risk* aren't well understood by many consumers. In some recent research, one of our participants got really tongue-tied trying to understand the idea that since they were young, it would be "high risk" (in the loose nontechnical definition of the word) to put their money into a "low-risk" fund

(in the technical definition of the word) since they'd probably end up with lower returns when they got older. Investment is confusing unless you've had training. Then, there's the problem that "a little knowledge can hurt." Some consumers who think they know what they're doing can end up suffering when they think they can beat the market by moving their money around between funds every week.

Self-service online pension (retirement plans) platforms don't do anything to help people make the right decisions because that would count as advice, which they're not able to give because of the way it's regulated. Giving an average person a self-service platform and telling them to go sort out their pension is like giving them a Unix terminal and telling them to sort out their own web server. A few PDF fact sheets just aren't going to help. If consumers want advice, they have to go to a financial advisor, which can be expensive and doesn't make financial sense unless you have a lot of money in the first place. There's a gap in the market, and we're working these sorts of challenges in my team at Smart Pension.

What would you say are the biggest challenges facing you and other consultants doing interaction design these days?

A career in interaction design is one of continual education and training. The biggest challenge is to keep this going. Even if you feel that you're at the peak of your skills, the technology landscape will be shifting under your feet, and you need to keep an eye on what's coming next so you don't get left behind. In fact, things move so quickly in interaction design that by the time you read this interview, it will already be dated.

If you ever find yourself in a "comfortable" role doing the same thing every

(Continued)

day, then beware—you're doing yourself a disservice. Get out there, stretch yourself, and make sure you spend some time every week outside your comfort zone.

If you're asked to evaluate a prototype service or product and you discover it is really bad, how do you break the news?

It depends what your goal is. If you want to just deliver the bad news and leave, then by all means be totally brutal and don't pull any punches. But if you want to build a relationship with the client, you're going

to need to help them work out how to move forward.

Remember, when you deliver bad news to a client, you're basically explaining to them that they're in a dark place and it's their fault. It can be quite embarrassing and depressing. It can drive stakeholders apart when really you need to bring them together and give them a shared vision to work toward. Discovering bad design is an opportunity for improvement. Always pair the bad news with a recommendation of what to do next.

NOTE

We use the term *interactive products* generically to refer to all classes of interactive systems, technologies, environments, tools, applications, services, and devices.

Chapter 2

THE PROCESS OF INTERACTION DESIGN

2.1 Introduction

2.2 What Is Involved in Interaction Design?

2.3 Some Practical Issues

Objectives

The main goals of this chapter are to accomplish the following:

- Reflect on what interaction design involves.
- Explain some of the advantages of involving users in development.
- Explain the main principles of a user-centered approach.
- Introduce the four basic activities of interaction design and how they are related in a simple lifecycle model.
- Ask some important questions about the interaction design process and provide the answers.
- Consider how interaction design activities can be integrated into other development lifecycles.

2.1 Introduction

Imagine that you have been asked to design a cloud-based service to enable people to share and curate their photos, movies, music, chats, documents, and so on, in an efficient, safe, and enjoyable way. What would you do? How would you start? Would you begin by sketching how the interface might look, work out how the system architecture should be structured, or just start coding? Or, would you start by asking users about their current experiences with sharing files and examine the existing tools, for example, Dropbox and Google Drive, and based on this begin thinking about how you were going to design the new service? What would you do next? This chapter discusses the process of interaction design, that is, how to design an interactive product.

There are many fields of design, such as graphic design, architectural design, industrial design, and software design. Although each discipline has its own approach to design, there

are commonalities. The Design Council of the United Kingdom captures these in the double-diamond of design, as shown in Figure 2.1. This approach has four phases which are iterated:

- *Discover*: Designers try to gather insights about the problem.
- *Define*: Designers develop a clear brief that frames the design challenge.
- *Develop*: Solutions or concepts are created, prototyped, tested, and iterated.
- *Deliver*: The resulting project is finalized, produced, and launched.

Interaction design also follows these phases, and it is underpinned by the philosophy of user-centered design, that is, involving users throughout development. Traditionally, interaction designers begin by doing user research and then sketching their ideas. But who are the users to be researched, and how can they be involved in development? Will they know what they want or need if we just ask them? From where do interaction designers get their ideas, and how do they generate designs?

In this chapter, we raise and answer these kinds of questions, discuss user-centered design, and explore the four basic activities of the interaction design process. We also introduce a lifecycle model of interaction design that captures these activities and the relationships among them.

2.2 What Is Involved in Interaction Design?

Interaction design has specific activities focused on discovering requirements for the product, designing something to fulfill those requirements, and producing prototypes that are then evaluated. In addition, interaction design focuses attention on users and their goals.

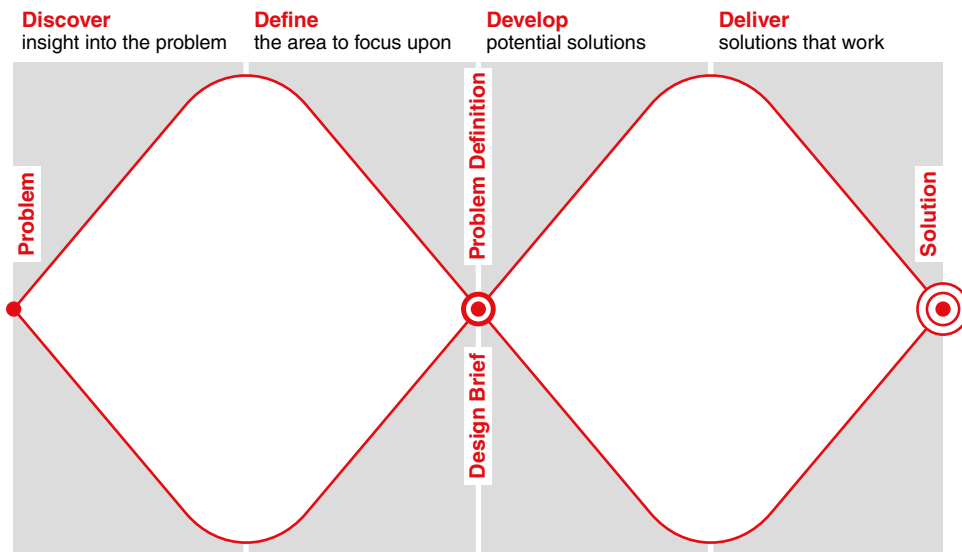


Figure 2.1 The double diamond of design

Source: Adapted from <https://www.designcouncil.org.uk/news-opinion/design-process-what-double-diamond>

For example, the artifact's use and target domain are investigated by taking a user-centered approach to development, users' opinions and reactions to early designs are sought, and users are involved appropriately in the development process itself. This means that users' concerns direct the development rather than just technical concerns.

Design is also about trade-offs—about balancing conflicting requirements. One common form of trade-off when developing a system to offer advice, for example, is deciding how much choice will be given to the user and how much direction the system should offer. Often, the division will depend on the purpose of the system, for example, whether it is for playing music tracks or for controlling traffic flow. Getting the balance right requires experience, but it also requires the development and evaluation of alternative solutions.

Generating alternatives is a key principle in most design disciplines and one that is also central to interaction design. Linus Pauling, twice a Nobel Prize winner, once said, “The best way to get a good idea is to get lots of ideas.” Generating lots of ideas is not necessarily hard, but choosing which of them to pursue is more difficult. For example, Tom Kelley (2016) describes seven secrets for successful brainstorming, including sharpening the focus (having a well-honed problem statement), having playful rules (to encourage ideas), and getting physical (using visual props).

Involving users and others in the design process means that the designs and potential solutions will need to be communicated to people other than the original designer. This requires the design to be captured and expressed in a form that allows review, revision, and improvement. There are many ways of doing this, one of the simplest being to produce a series of sketches. Other common approaches are to write a description in natural language, to draw a series of diagrams, and to build a prototype, that is, a limited version of the final product. A combination of these techniques is likely to be the most effective. When users are involved, capturing and expressing a design in a suitable format is especially important since they are unlikely to understand jargon or specialist notations. In fact, a form with which users can interact is most effective, so building prototypes is an extremely powerful approach.

ACTIVITY 2.1

This activity asks you to apply the double diamond of design to produce an innovative interactive product for your own use. By focusing on a product for yourself, the activity deliberately de-emphasizes issues concerned with involving other users, and instead it emphasizes the overall process.

Imagine that you want to design a product that helps you organize a trip. This might be for a business or vacation trip, to visit relatives halfway around the world, or for a bike ride on the weekend—whatever kind of trip you like. In addition to planning the route or booking tickets, the product may help to check visa requirements, arrange guided tours, investigate the facilities at a location, and so on.

1. Using the first three phases of the double diamond of design, produce an initial design using a sketch or two, showing its main functionality and its general look and feel. This activity omits the fourth phase, as you are not expected to deliver a working solution.
2. Now reflect on how your activities fell into these phases. What did you do first? What was your instinct to do first? Did you have any particular artifacts or experiences upon which to base your design?

(Continued)

Comment

1. The first phase focuses on discovering insights about the problem, but is there a problem? If so, what is it? Although most of us manage to book trips and travel to destinations with the right visas and in comfort, upon reflection the process and the outcome can be improved. For example, dietary requirements are not always fulfilled, and the accommodation is not always in the best location. There is a lot of information available to support organizing travel, and there are many agents, websites, travel books, and tourist boards that can help. The problem is that it can be overwhelming.

The second phase is about defining the area on which to focus. There are many reasons for travelling—both individual and family—but in my experience organizing business trips to meetings worldwide is stressful, and minimizing the complexity involved in these would be worthwhile. The experience would be improved if the product offers advice from the many possible sources of information and tailors that advice to individual preferences.

The third phase focuses on developing solutions, which in this case is a sketch of the design itself. Figure 2.2 shows an initial design. This has two versions of the product—one as an app to run on a mobile device and one to run on a larger screen. The assumptions underlying the choice to build two versions are based on my experience; I would normally plan the details of the trip at my desk, while requiring updates and local information while traveling. The mobile app has a simple interaction style that is easy to use on the go, while the larger-screen version is more sophisticated and shows a lot of information and the various choices available.

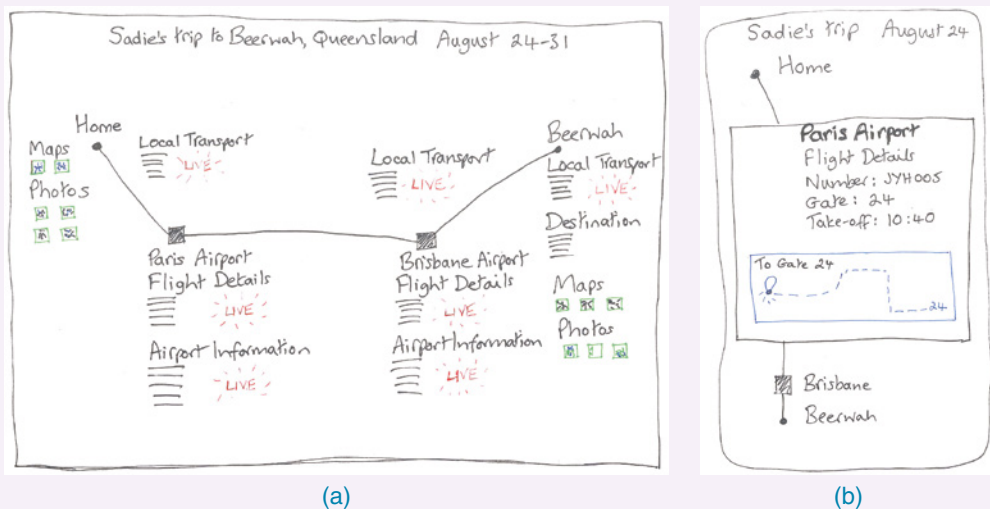


Figure 2.2 Initial sketches of the trip organizer showing (a) a large screen covering the entire journey from home to Beerwah in Australia and (b) the smartphone screen available for the leg of the journey at Paris (Charles de Gaulle) airport

2. Initially, it wasn't clear that there was a problem to address, but on reflection the complexity of the available information and the benefit of tailoring choices became clearer. The second phase guided me toward thinking about the area on which to focus. Worldwide business trips are the most difficult, and reducing the complexity of information sources through customization would definitely help. It would be good if the product learned about my preferences, for example, recommending flights from my favorite airline and finding places to have a vegan meal.

Developing solutions (the third phase) led me to consider how to interact with the product—seeing detail on a large screen would be useful, but a summary that can be shown on a mobile device is also needed. The type of support also depends on where the meeting is being held. Planning a trip abroad requires both a high-level view to check visas, vaccinations, and travel advice, as well as a detailed view about the proximity of accommodation to the meeting venue and specific flight times. Planning a local trip is much less complicated.

The exact steps taken to create a product will vary from designer to designer, from product to product, and from organization to organization (see Box 2.1). Capturing concrete ideas, through sketches or written descriptions, helps to focus the mind on what is being designed, the context of the design, and what user experience is to be expected. The sketches can capture only some elements of the design, however, and other formats are needed to capture everything intended. Throughout this activity, you have been making choices between alternatives, exploring requirements in detail, and refining your ideas about what the product will do. ■

2.2.1 Understanding the Problem Space

Deciding what to design is key, and exploring the problem space is one way in which to decide. This is the first phase in the double diamond, but it can be overlooked by those new to interaction design, as you may have discovered in Activity 2.1. In the process of creating an interactive product, it can be tempting to begin at the nuts and bolts level of design. By this we mean working out how to design the physical interface and what technologies and interaction styles to use, for example, whether to use multitouch, voice, graphical user interface, heads-up display, augmented reality, gesture-based, and so forth. The problem with starting here is that potential users and their context can be misunderstood, and usability and user experience goals can be overlooked, both of which were discussed in Chapter 1, “What Is Interaction Design?”

For example, consider the augmented reality displays and holographic navigation systems that are available in some cars nowadays (see Figure 2.3). They are the result of decades of research into human factors of information displays (for instance, Campbell et al., 2016), the driving experience itself (Perterer et al., 2013; Lee et al., 2005), and the suitability of different technologies (for example, Jose et al., 2016), as well as improvements in technology. Understanding the problem space has been critical in arriving at workable solutions that are safe and trusted. Having said that, some people may not be comfortable using a holographic navigation system and choose not to have one installed.



(a)



(b)

Figure 2.3 (a) Example of the holographic navigation display from WayRay which overlays GPS navigation instructions onto the road ahead and gathers and shares driver statistics (b) an augmented reality navigation system available in some cars today

Sources: (a) Used courtesy of WayRay, (b) Used courtesy of Muhammad Saad

While it is certainly necessary at some point to choose which technology to employ and decide how to design the physical aspects, it is better to make these decisions after articulating the nature of the problem space. By this we mean understanding what is currently the user experience or the product, why a change is needed, and how this change will improve the user experience. In the previous example, this involves finding out what is problematic with existing support for navigating while driving. An example is ensuring that drivers can continue to drive safely without being distracted when looking at a small GPS display mounted on the dashboard to figure out on which road it is asking them to “turn left.” Even when designing for a new user experience, it still requires understanding the context for which it will be used and the possible current user expectations.

The process of articulating the problem space is typically done as a team effort. Invariably, team members will have differing perspectives on it. For example, a project manager is likely to be concerned about a proposed solution in terms of budgets, timelines, and staffing costs, whereas a software engineer will be thinking about breaking it down into specific technical concepts. The implications of pursuing each perspective need to be considered in relation to one another. Although time-consuming and sometimes resulting in disagreements among the design team, the benefits of this process can far outweigh the associated costs: there will be much less chance of incorrect assumptions and unsupported

claims creeping into a design solution that later turn out to be unusable or unwanted. Spending time enumerating and reflecting upon ideas during the early stages of the design process enables more options and possibilities to be considered. Furthermore, designers are increasingly expected to justify their choice of problems and to be able to present clearly and convincingly their rationale in business as well as design language. Being able to think and analyze, present, and argue is valued as much as the ability to create a product (Kolko, 2011).

BOX 2.1

Four Approaches to Interaction Design

Dan Saffer (2010) suggests four main approaches to interaction design, each of which is based on a distinct underlying philosophy: User-centered design, Activity-centered design, Systems design, and Genius design.

Dan Saffer acknowledges that the purest form of any of these approaches is unlikely to be realized, and he takes an extreme view of each in order to distinguish among them. In user-centered design, the user knows best and is the guide to the designer; the designer's role is to translate the users' needs and goals into a design solution.

Activity-centered design focuses on the behavior surrounding particular tasks. Users still play a significant role, but it is their behavior rather than their goals and needs that is important. Systems design is a structured, rigorous, and holistic design approach that focuses on context and is particularly appropriate for complex problems. In systems design, it is the system (that is, the people, computers, objects, devices, and so on) that the center of attention, while the users' role is to set the goals of the system.

Finally, genius design is different from the other three approaches because it relies largely on the experience and creative flair of a designer. Jim Leftwich, an experienced interaction designer interviewed by Dan Saffer (2010, pp. 44–45), prefers the term *rapid expert design*. In this approach, the users' role is to validate ideas generated by the designer, and users are not involved during the design process itself. Dan Saffer points out that this is not necessarily by choice, but it may be because of limited or no resources for user involvement.

Different design problems lend themselves more easily to different approaches, and different designers will tend to gravitate toward using the approach that suits them best. Although an individual designer may prefer a particular approach, it is important that the approach for any one design problem is chosen with that design problem in mind. ■

2.2.2 The Importance of Involving Users

Chapter 1 stressed the importance of understanding users, and the previous description emphasizes the need to involve users in interaction design. Involving users in development is important because it's the best way to ensure that the end product is usable and that it indeed will be used. In the past, it was common for developers to talk only to managers, experts, or

proxy users, or even to use their own judgment without reference to anyone else. While others involved in designing the product can provide useful information, they will not have the same perspective as the target user who performs the activity every day or who will use the intended product on a regular basis.

In commercial projects, a role called the *product owner* is common. The product owner's job is to filter user and customer input to the development cycle and to prioritize requirements or features. This person is usually someone with business and technical knowledge, but not interaction design knowledge, and they are rarely (if ever) a direct user of the product. Although the product owner may be called upon to assess designs, they are a proxy user at best, and their involvement does not avoid the need for user involvement.

The best way to ensure that developers gain a good understanding of users' goals, leading to a more appropriate, more usable product, is to involve target users throughout development. However, two other aspects unrelated to functionality are equally as important if the product is to be usable and used: expectation management and ownership.

Expectation management is the process of making sure that the users' expectations of the new product are realistic. Its purpose is to ensure that there are no surprises for users when the product arrives. If users feel they have been cheated by promises that have not been fulfilled, then this will cause resistance and even rejection. Marketing of the new arrival must be careful not to misrepresent the product, although it may be particularly difficult to achieve with a large and complex system (Nevo and Wade, 2007). How many times have you seen an advertisement for something that you thought would be really good to have, but when you actually see one, you discover that the marketing hype was a little exaggerated? We expect that you felt quite disappointed and let down. This is the kind of feeling that expectation management tries to avoid.

Involving users throughout development helps with expectation management because they can see the product's capabilities from an early stage. They will also understand better how it will affect their jobs and lives and why the features are designed that way. Adequate and timely training is another technique for managing expectations. If users have the chance to work with the product before it is released through training or hands-on demonstrations of a prerelease version, then they will understand better what to expect when the final product is available.

A second reason for user involvement is ownership. Users who are involved and feel that they have contributed to a product's development are more likely to feel a sense of ownership toward it and support its use (Bano et al., 2017).

How to involve users, in what roles, and for how long, needs careful planning, as discussed in the next Dilemma box.

DILEMMA

Too Much of a Good Thing?

Involving users in development is a good thing, but what evidence is there that user involvement is productive? How much should users be involved and in what role(s)? Is it appropriate for users to lead a technical development project, or is it more beneficial for them to focus on evaluating prototypes?

Uli Abelein et al. (2013) performed a detailed review of the literature in this area and concluded that, overall, the evidence indicates that user involvement has a positive effect on user satisfaction and system use. However, they also found that even though the data clearly indicates this positive effect, some links have a large variation, suggesting that there is still no clear way to measure the effects consistently. In addition, they found that most studies with negative correlations involving users and system success were published more than 10 years previously.

Ramanath Subrayaman et al. (2010) investigated the impact of user participation on levels of satisfaction with the product by both developers and users. They found that for new products, developer satisfaction increased as user participation increased. On the other hand, user satisfaction was higher if their participation was low, and satisfaction dropped as their participation increased. They also identified that high levels of user involvement can generate conflicts and increased reworking. For maintenance projects, both developers and users were most satisfied with a moderate level of participation (approximately 20 percent of overall project development time). Based just on user satisfaction as an indicator of project success, then, it seems that low user participation is most beneficial.

The kind of product being developed, the kind of user involvement possible, the activities in which they are involved, and the application domain all have an impact on the effectiveness of user input (Bano and Zowghi, 2015). Peter Richard et al. (2014) investigated the effect of user involvement in transport design projects. They found that involving users at later stages of development mainly resulted in suggestions for service improvement, whereas users involved at earlier stages of innovation suggested more creative ideas.

Recent moves toward an agile way of working (see Chapter 13, “Interaction Design in Practice”) has emphasized the need for feedback from customers and users, but this also has its challenges. Kurt Schmitz et al. (2018) suggests that in tailoring their methods, teams consider the distinction between frequent participation in activities and effective engagement.

User involvement is undoubtedly beneficial, but the levels and types of involvement require careful consideration and balance. ■

2.2.3 Degrees of User Involvement

Different degrees of user involvement are possible, ranging from fully engaged throughout all iterations of the development process to targeted participation in specific activities and from small groups of individual users in face-to-face contexts to hundreds of thousands of potential users and stakeholders online. Where available, individual users may be co-opted onto the design team so that they are major contributors to the development. This has pros and cons. On the downside, full-time involvement may mean that they become out of touch with their user community, while part-time involvement might result in a high workload for them. On the positive side, having a user engaged full or part-time does mean that the input is available continually throughout development. On the other hand, users may take part in specific activities to inform the development or to evaluate designs once they are available. This is a valuable form of involvement, but the users’ input is limited to that particular activity. Where the circumstances around a project limit user involvement in this way, there are techniques to keep users’ concerns uppermost in developers’ minds, such as through personas (see Chapter 11, “Discovering Requirements”).

Initially, user involvement took the form of small groups or individuals taking part in face-to-face information-gathering design, or evaluation sessions, but increasing online connectivity has led to a situation in which many thousands of potential users can contribute to product development. There is still a place for face-to-face user involvement and *in situ* studies, but the range of possibilities for user involvement is now much wider. One example of this is online feedback exchange (OFE) systems, which are increasingly used to test design concepts with millions of target users before going to market (Foong et al., 2017).

In fact, design is becoming increasingly participative through crowdsourcing design ideas and examples, for instance (Yu et al., 2016). Where crowdsourcing is used, a range of different people are encouraged to contribute, and this can include any and all of the stakeholders. This wide participation helps to bring different perspectives to the process, which enhances the design itself, produces more user satisfaction with the final product, and engenders a sense of ownership. Another example of involving users at scale is citizen engagement, the goal of which is to engage a population—civic or otherwise—with the aim of promoting empowerment through technology. The underlying aim is to involve members of the public in helping them make a change in their lives where technology is often viewed as an integral part of the process.

Participatory design, also sometimes referred to as *cooperative design* or *co-design*, is an overarching design philosophy that places those for whom systems, technologies, and services are being designed, as central actors in creation activities. The idea is that instead of being passive receivers of new technological or industrial artifacts, end users and stakeholders are active participants in the design process. Chapter 12, “Design, Prototyping, and Construction,” provides more information on participatory design.

The individual circumstances of the project affect what is realistic and appropriate. If the end-user groups are identifiable, for example, the product is for a particular company, then it is easier to involve them. If, however, the product is intended for the open market, it is unlikely that users will be available to join the design team. In this case, targeted activities and online feedback systems may be employed. Box 2.2 outlines an alternative way to obtain user input from an existing product, and Box 2.5 discusses A/B testing, which draws on user feedback to choose between alternative designs.

BOX 2.2

User Involvement After Product Release

Once a product has been released, a different kind of user involvement is possible—one that captures data and user feedback based on day-to-day use of the product. The prevalence of customer reviews has grown considerably in recent years, and they significantly affect the popularity and success of a product (Harman et al., 2012). These reviews provide useful and far-ranging user feedback. For example, Hammad Khalid et al. (2015) studied reviews of mobile apps to see what reviewers complained about. They identified 12 complaint types, including privacy and ethics, interface, and feature removal. Customer reviews can provide useful insight to help improve products, but detailed analysis of feedback gathered this way is time-consuming.

Error reporting systems (ERSs, also called *online crashing analysis*) automatically collect information from users that is used to improve applications in the longer term. This is done with users' permission, but with a minimal reporting burden. Figure 2.4 shows two dialog boxes for the Windows error reporting system that is built into Microsoft operating systems. This kind of reporting can have a significant effect on the quality of applications. For example, 29 percent of the errors fixed by the Windows XP (Service Pack 1) team were based on information collected through their ERS (Kinshumann et al., 2011). While Windows XP is no longer being supported, this statistic illustrates the impact ERSs can have. The system uses a sophisticated approach to error reporting based on five strategies: automatic aggregation of error reports; progressive data collection so that the data collected (such as abbreviated or full stack and memory dumps) varies depending on the level of data needed to diagnose the error; minimal user interaction; preserving user privacy; and providing solutions directly to users where possible. By using these strategies, plus statistical analysis, effort can be focused on the bugs that have the highest impact on the most users. ■

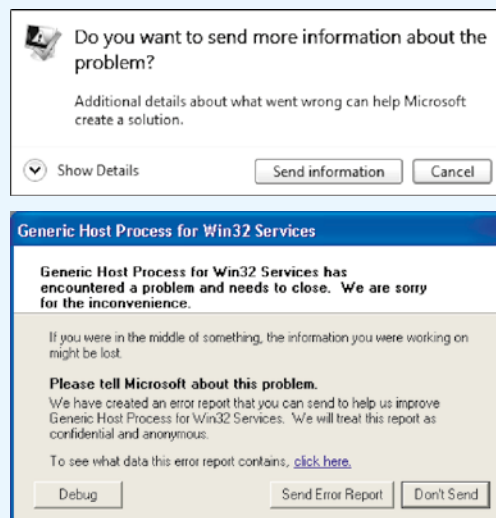


Figure 2.4 Two typical dialog boxes from the Windows error reporting system

2.2.4 What Is a User-Centered Approach?

Throughout this book, we emphasize the need for a user-centered approach to development. By this we mean that the real users and their goals, not just technology, are the driving force behind product development. As a consequence, a well-designed system will make the most of human skill and judgment, will be directly relevant to the activity in hand, and will support rather than constrain the user. This is less of a technique and more of a philosophy.

When the field of HCI was being established, John Gould and Clayton Lewis (1985) laid down three principles that they believed would lead to a “useful and easy to use computer system.” These principles are as follows:

1. *Early focus on users and tasks.* This means first understanding who the users will be by directly studying their cognitive, behavioral, anthropomorphic, and attitudinal characteristics. This requires observing users doing their normal tasks, studying the nature of those tasks, and then involving users in the design process.
2. *Empirical measurement.* Early in development, the reactions and performance of intended users to printed scenarios, manuals, and so forth, are observed and measured. Later, users interact with simulations and prototypes, and their performance and reactions are observed, recorded, and analyzed.
3. *Iterative design.* When problems are found in user testing, they are fixed, and then more tests and observations are carried out to see the effects of the fixes. This means that design and development are iterative, with cycles of design-test-measure-redesign being repeated as often as necessary.

These three principles are now generally accepted as the basis for a user-centered approach. When this paper was written, however, they were not accepted by most developers. We discuss these principles in more detail in the following sections.

Early Focus on Users and Tasks

This principle can be expanded and clarified through the following five further principles:

1. *Users’ tasks and goals are the driving force behind the development.*
While technology will inform design options and choices, it is not the driving force. Instead of saying “Where can we deploy this new technology?” say “What technologies are available to provide better support for users’ goals?”
2. *Users’ behavior and context of use are studied, and the system is designed to support them.*
This is not just about capturing users’ tasks and goals. How people perform their tasks is also significant. Understanding behavior highlights priorities, preferences, and implicit intentions.
3. *Users’ characteristics are captured and designed for.*
When things go wrong with technology, people often think it is their fault. People are prone to making errors and have certain limitations, both cognitive and physical. Products designed to support people should take these limitations into account and try to prevent mistakes from being made. Cognitive aspects, such as attention, memory, and perception issues are introduced in Chapter 4, “Cognitive Aspects.” Physical aspects include height, mobility, and strength. Some characteristics are general, such as color blindness, which affects about 4.5 percent of the population, but some characteristics are associated with a particular job or task. In addition to general characteristics, those traits specific to the intended user group also need to be captured.
4. *Users are consulted throughout development from earliest phases to the latest.*
As discussed earlier, there are different levels of user involvement, and there are different ways in which to consult users.
5. *All design decisions are taken within the context of the users, their activities, and their environment.*
This does not necessarily mean that users are actively involved in design decisions, but that is one option.

ACTIVITY 2.2

Assume you are involved in developing a novel online experience for buying garden plants. Although many websites exist for buying plants online, you want to produce a distinct experience to increase the organization's market share. Suggest ways of applying the previous principles in this task.

Comment

To address the first three principles, you would need to find out about the tasks and goals, behavior, and characteristics of potential customers of the new experience, together with any different contexts of use. Studying current users of existing online plant shops will provide some information, and it will also identify some challenges to be addressed in the new experience. However, as you want to increase the organization's market share, consulting existing users alone would not be enough. Alternative avenues of investigation include physical shopping situations—for example, shopping at the market, in the local corner shop, and so on, and local gardening clubs, radio programs, or podcasts. These alternatives will help you find the advantages and disadvantages of buying plants in different settings, and you will observe different behaviors. By looking at these options, a new set of potential users and contexts can be identified.

For the fourth principle, the set of new users will emerge as investigations progress, but people who are representative of the user group may be accessible from the beginning. Workshops or evaluation sessions could be run with them, possibly in one of the alternative shopping environments such as the market. The last principle could be supported through the creation of a design room that houses all of the data collected, and it is a place where the development team can go to find out more about the users and the product goals. ■

Empirical Measurement

Where possible, specific usability and user experience goals should be identified, clearly documented, and agreed upon at the beginning of the project. They can help designers choose between alternative designs and check on progress as the product is developed. Identifying specific goals up front means that the product can be empirically evaluated at regular stages throughout development.

Iterative Design

Iteration allows designs to be refined based on feedback. As users and designers engage with the domain and start to discuss requirements, needs, hopes, and aspirations, then different insights into what is needed, what will help, and what is feasible will emerge. This leads to a need for iteration—for the activities to inform each other and to be repeated. No matter how good the designers are and however clear the users may think their vision is of the required artifact, ideas will need to be revised in light of feedback, likely several times. This is particularly true when trying to innovate. Innovation rarely emerges whole and ready to go. It takes time, evolution, trial and error, and a great deal of patience. Iteration is inevitable because designers never get the solution right the first time (Gould and Lewis, 1985).

2.2.5 Four Basic Activities of Interaction Design

The four basic activities for interaction design are as follows:

1. Discovering requirements for the interactive product.
2. Designing alternatives that meet those requirements.
3. Prototyping the alternative designs so that they can be communicated and assessed.
4. Evaluating the product and the user experience it offers throughout the process.

Discovering Requirements

This activity covers the left side of the double diamond of design, and it is focused on discovering something new about the world and defining what will be developed. In the case of interaction design, this includes understanding the target users and the support an interactive product could usefully provide. This understanding is gleaned through data gathering and analysis, which are discussed in Chapters 8–10. It forms the basis of the product's requirements and underpins subsequent design and development. The requirements activity is discussed further in Chapter 11.

Designing Alternatives

This is the core activity of designing and is part of the Develop phase of the double diamond: proposing ideas for meeting the requirements. For interaction design, this activity can be viewed as two subactivities: conceptual design and concrete design. Conceptual design involves producing the conceptual model for the product, and a conceptual model describes an abstraction outlining what people can do with a product and what concepts are needed to understand how to interact with it. Concrete design considers the detail of the product including the colors, sounds, and images to use, menu design, and icon design. Alternatives are considered at every point. Conceptual design is discussed in Chapter 3, and more design issues for specific interface types are in Chapter 7; more detail about how to design an interactive product is in Chapter 12.

Prototyping

Prototyping is also part of the Develop phase of the double diamond. Interaction design involves designing the behavior of interactive products as well as their look and feel. The most effective way for users to evaluate such designs is to interact with them, and this can be achieved through prototyping. This does not necessarily mean that a piece of software is required. There are different prototyping techniques, not all of which require a working piece of software. For example, paper-based prototypes are quick and cheap to build and are effective for identifying problems in the early stages of design, and through role-playing users can get a real sense of what it will be like to interact with the product. Prototyping is covered in Chapter 12.

Evaluating

Evaluating is also part of the Develop phase of the double diamond. It is the process of determining the usability and acceptability of the product or design measured in terms of a variety of usability and user-experience criteria. Evaluation does not replace activities concerned with quality assurance and testing to make sure that the final product is fit for its intended purpose, but it complements and enhances them. Chapters 14–16 cover evaluation.

The activities to discover requirements, design alternatives, build prototypes, and evaluate them are intertwined: alternatives are evaluated through the prototypes, and the results are fed back into further design or to identify alternative requirements.

2.2.6 A Simple Lifecycle Model for Interaction Design

Understanding what activities are involved in interaction design is the first step to being able to do it, but it is also important to consider how the activities are related to one another. The term *lifecycle model* (or *process model*) is used to represent a model that captures a set of activities and how they are related. Existing models have varying levels of sophistication and complexity and are often not prescriptive. For projects involving only a few experienced developers, a simple process is adequate. However, for larger systems involving tens or hundreds of developers with hundreds or thousands of users, a simple process just isn't enough to provide the management structure and discipline necessary to engineer a usable product.



Source: Fran / Cartoon Stock

Many lifecycle models have been proposed in fields related to interaction design. For example, software engineering lifecycle models include the waterfall, spiral, and V models (for more information about these models, see Pressman and Maxim [2014]). HCI has been less associated with lifecycle models, but two well-known ones are the Star (Hartson and Hix, 1989) and an international standard model ISO 9241-210. Rather than explaining the details of these models, we focus on the classic lifecycle model shown in Figure 2.5. This model shows how the four activities of interaction design are related, and it incorporates the three principles of user-centered design discussed earlier.

Many projects start by discovering requirements from which alternative designs are generated. Prototype versions of the designs are developed and then evaluated. During prototyping or based on feedback from evaluations, the team may need to refine the requirements or to redesign. One or more alternative designs may follow this iterative cycle in parallel. Implicit in this cycle is that the final product will emerge in an evolutionary fashion from an initial idea through to the finished product or from limited functionality to sophisticated functionality. Exactly how this evolution happens varies from project to project. However many times through the cycle the product goes, development ends with an evaluation activity that ensures that the final product meets the prescribed user experience and usability criteria. This evolutionary production is part of the Delivery phase of the double diamond.

In recent years, a wide range of lifecycle models has emerged, all of which encompass these activities but with different emphases on activities, relationships, and outputs. For example, Google Design Sprints (Box 2.3) emphasize problem investigation, solution development, and testing with customers all in one week. This does not result in a robust final product, but it does make sure that the solution idea is acceptable to customers. The in-the-wild approach (Box 2.4) emphasizes the development of novel technologies that are not necessarily designed for specific user needs but to augment people, places, and settings. Further models are discussed in Chapter 13.

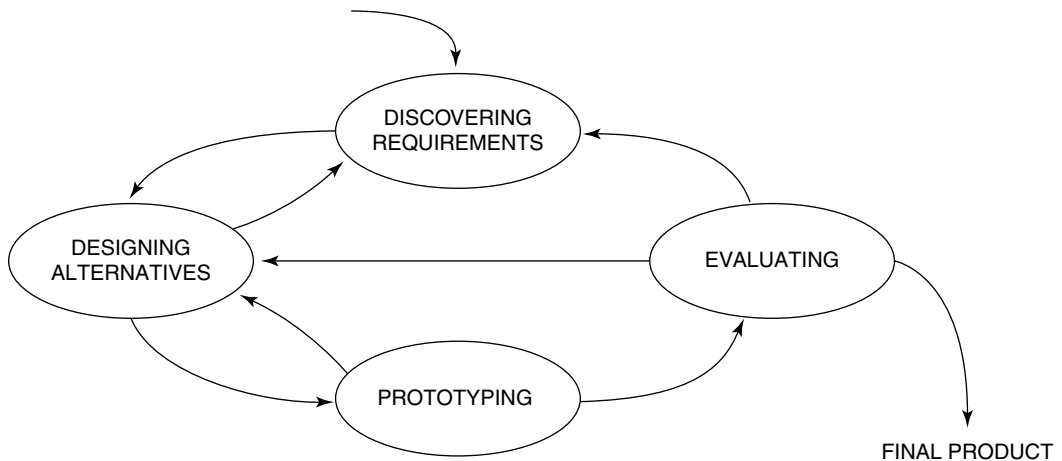


Figure 2.5 A simple interaction design lifecycle model

BOX 2.3

Google Design Sprints (Adapted from Knapp et al. (2016))

Google Ventures has developed a structured approach to design that supports rapid ideation and testing of potential solutions to a design challenge. This is called the *Google Design Sprint*. A sprint is divided into five phases, and each phase is completed in a day. This means that in five days, you can go from a design challenge to a solution that has been tested with customers. As the authors say, “You won’t finish with a complete, detailed, ready-to-ship product. But you will make rapid progress, and know for sure if you’re headed in the right direction” (Knapp et al., 2016, p16–17). Teams are encouraged to iterate on the last two phases and to develop and re-test prototypes. If necessary, the first idea can be thrown away and the process started again at Phase 1. There is preparation to be done before the sprint begins. This preparation and the five phases are described next (see Figure 2.6).




Figure 2.6 The five phases of the Google Design Sprint

Source: www.agilemarketing.net/google-design-sprints. Used courtesy of Agile Marketing

Setting the Stage

This time is used to choose the right design challenge, gather the right team, and organize the time and space to run the sprint (that is, full-time for everyone for five days). The sprint can help in high-stake challenges, when you're running out of time, or if you're just stuck. The team composition depends on the product, but it has about seven people including a decider (who chooses the design to show to the customer), customer expert, technical expert, and anyone who will bring a disruptive perspective.

Unpack

Day 1 focuses on making a map of the challenge and choosing a target, that is, a part of the challenge that can be achieved in a week.

Sketch Competing Solutions

Day 2 focuses on generating solutions, with an emphasis on sketching and individual creativity rather than group brainstorming.

Decide on the Best

Day 3 focuses on critiquing the solutions generated on Day 1, choosing the one most likely to meet the sprint's challenge, and producing a storyboard. Whichever solution is chosen, the decider needs to support the design.

Build a Realistic Prototype

Day 4 focuses on turning the storyboard into a realistic prototype, that is, something on which customers can provide feedback. We discuss prototyping further in Chapter 12.

Test with Target Customers

Day 5 focuses on getting feedback from five customers and learning from their reactions.

The Google Design Sprint is a process for answering critical business questions through design, prototyping, and testing ideas with customers. Marta Rey-Babarro, who works at Google as a staff UX researcher and was the cofounder of Google's internal Sprint Academy, describes how they used a sprint to improve the experience of traveling for business.

We wanted to see if we could improve the business travel experience. We started by doing research with Googlers to find out what experiences and what needs they had when they traveled. We discovered that there were some Googlers who traveled over 300 days a year and others who traveled only once or twice a year. Their travel experiences and needs were very different. After this research, some of us did a sprint in which we explored the whole travel experience, from the planning phase to coming back home and submitting receipts. Within five days we came up with a vision of what that experience could be. On the fifth day of the sprint, we presented that vision to higher-level execs. They loved it and sponsored the creation of a new team at Google that has developed new

tools and experiences for the traveling Googler. Some of those internal online experiences made it also to our external products and services outside of Google.

Marta Rey-Babarro ■

To see a more detailed description of the Google Design Sprint and to access a set of five videos that describe what happens on each day of the sprint, go to **www.gv.com/sprint/#book**.

BOX 2.4

Research in the Wild (Adapted from Rogers and Marshall (2017))

Research in the wild (RITW) develops technology solutions in everyday living by creating and evaluating new technologies and experiences *in situ*. The approach supports designing prototypes in which researchers often experiment with new technological possibilities that can change and even disrupt behavior, rather than ones that fit in with existing practices. The results of RITW studies can be used to challenge assumptions about technology and human behavior in the real world and to inform the re-thinking of HCI theories. The perspective taken by RITW studies is to observe how people react to technology and how they change and integrate it into their everyday lives.

Figure 2.7 shows the framework for RITW studies. In terms of the four activities introduced earlier, this framework focuses on designing, prototyping, and evaluating technology and ideas and is one way in which requirements may be discovered. It also considers relevant theory since often the purpose of an RITW study is to investigate a theory, idea, concept, or observation. Any one RITW study may emphasize the elements of the framework to a different degree.

Technology: Concerned with appropriating existing infrastructures/devices (e.g., Internet of Things toolkit, mobile app) *in situ* or developing new ones for a given setting (e.g., a novel public display).

Design: Covers the design space of an experience (e.g., iteratively creating a collaborative travel planning tool for families to use or an augmented reality game for playing outdoors).

In situ study: Concerned with evaluating *in situ* an existing device/tool/service or novel research-based prototype when placed in various settings or given to someone to use over a period of time.

Theory: Investigating a theory, idea, concept or observation about a behavior, setting or other phenomenon using existing ones or developing a new one or extending an existing one. ■




Figure 2.7 A framework for research in the wild studies

Source: Rogers and Marshall (2017), p. 6. Used courtesy of Morgan & Claypool

2.3 Some Practical Issues

The discussion so far has highlighted some issues about the practical application of user-centered design and the simple lifecycle of interaction design introduced earlier. These issues are listed here:

- Who are the users?
- What are the users' needs?
- How to generate alternative designs
- How to choose among alternatives
- How to integrate interaction design activities with other lifecycle models

2.3.1 Who Are the Users?

Identifying users may seem like a straightforward activity, but it can be harder than you think. For example, Sha Zhao et al. (2016) found a more diverse set of users for smartphones than most manufacturers recognize. Based on an analysis of one month's smartphone app

usage, they discovered 382 distinct types of users, including Screen Checkers and Young Parents. Charlie Wilson et al. (2015) found that little is understood about who the users of smart homes in general are expected to be, beyond those focused on health-related conditions. In part, this is because many products nowadays are being developed for use by large sections of the population, and so it can be difficult to determine a clear description. Some products (such as a system to schedule work shifts) have more constrained user communities, for example a specific role (shop assistant) within a particular industrial sector (retail). In this case, there may be a range of users with different roles who relate to the product in different ways. Examples are those who manage direct users, those who receive outputs from the system, those who test the system, those who make the purchasing decision, and those who use competitive products (Holzblatt and Jones, 1993).

There is a surprisingly wide collection of people who all have a stake in the development of a successful product. These people are called *stakeholders*. Stakeholders are the individuals or groups that can influence or be influenced by the success or failure of a project. Alan Dix et al. (2004) observed that is pertinent to a user-centered view of development: “It will frequently be the case that the formal ‘client’ who orders the system falls very low on the list of those affected. Be very wary of changes which take power, influence or control from some stakeholders without returning something tangible in its place.”

The group of stakeholders for a particular product will be larger than the group of users. It will include customers who pay for it; users who interact with it; developers who design, build, and maintain it; legislators who impose rules on the development and operation of it; people who may lose their jobs because of its introduction; and so on (Sharp et al., 1999).

Identifying the stakeholders for a project helps to decide who to involve as users and to what degree, but identifying relevant stakeholders can be tricky. Ian Alexander and Suzanne Robertson (2004) suggest using an onion diagram to model stakeholders and their involvement. This diagram shows concentric circles of stakeholder zones with the product being developed sitting in the middle. Soo Ling Lim and Anthony Finkelstein (2012) developed a method called StakeRare and supporting tool called StakeNet that relies on social networks and collaborative filtering to identify and prioritize relevant stakeholders.

ACTIVITY 2.3

Who are the stakeholders for an electricity smart meter for use in the home to help households control their energy consumption?

Comment

First, there are the people who live in the house, such as older adults and young children, with a range of abilities and backgrounds. To varying degrees, they will be users of the meter, and their stake in its success and usability is fairly clear and direct. Householders want to make sure that their bills are controlled, that they can easily access suppliers if they want to, and that their electricity supply is not interrupted. On the other hand, the entire family will

want to continue to live in the house in comfort, for example, with enough heat and light. Then there are the people who install and maintain the meter. They make sure that the meter is installed correctly and that it continues to work effectively. Installers and maintainers want the meter to be straightforward to install and to be robust and reliable to reduce the need for return visits or maintenance calls. Outside of these groups are electricity suppliers and distributors who also want to provide a competitive service so that the householders are satisfied and to minimize maintenance costs. They also don't want to lose customers and money because the meters are faulty or are providing inaccurate information. Other people who will be affected by the success of the meter include those who work on the powerlines and at electricity generation plants, those who work in other energy industries, and ultimately the government of the country that will want to maintain steady supply for its industry and population. ■

2.3.2 What Are the Users' Needs?

If you had asked someone in the street in the late 1990s what they needed, their answer probably wouldn't have included a smart TV, a ski jacket with an integrated smartphone, or a robot pet. If you presented the same person with these possibilities and asked whether they would buy them if they were available, then the answer may have been more positive. Determining what product to build is not simply a question of asking people "What do you need?" and then supplying it, because people don't necessarily know what is possible. Suzanne and James Robertson (2013) refer to "un-dreamed-of" needs, which are those that users are unaware they might have. Instead of asking users, this is approached by exploring the problem space, investigating the users and their activities to see what can be improved, or trying out ideas with potential users to see whether the ideas are successful. In practice, a mixture of these approaches is often taken—trying ideas in order to discover requirements and decide what to build, but with knowledge of the problem space, potential users, and their activities.

If a product is a new invention, then identifying the users and representative tasks for them may be harder. This is where in-the-wild studies or rapid design sprints that provide authentic user feedback on early ideas are valuable. Rather than imagining who might want to use a product and what they might want to do with it, it's more effective to put it out there and find out—the results might be surprising!

It may be tempting for designers simply to design what they would like to use themselves, but their ideas would not necessarily coincide with those of the target user group, because they have different experiences and expectations. Several practitioners and commentators have observed that it's an "eye-opening experience" when developers or designers see a user struggling to complete a task that seemed so clear to them (Ratcliffe and McNeill, 2012, p. 125).

Focusing on people's goals, usability goals and user experience goals is a more promising approach to interaction design than simply expecting stakeholders to be able to articulate the requirements for a product.

2.3.3 How to Generate Alternative Designs

A common human tendency is to stick with something that works. While recognizing that a better solution may exist, it is easy to accept the one that works as being “good enough.” Settling for a solution that is good enough may be undesirable because better alternatives may never be considered, and considering alternative solutions is a crucial step in the process of design. But where do these alternative ideas come from?

One answer to this question is that they come from the individual designer’s flair and creativity (the genius design described in Box 2.1). Although it is certainly true that some people are able to produce wonderfully inspired designs while others struggle to come up with any ideas at all, very little in this world is completely new. For example, the steam engine, commonly regarded as an invention, was inspired by the observation that steam from a kettle boiling on the stove lifted the lid. An amount of creativity and engineering was needed to make the jump from a boiling kettle to a steam engine, but the kettle provided inspiration to translate this experience into a set of principles that could be applied in a different context. Innovations often arise through cross-fertilization of ideas from different perspectives, individuals, and contexts; the evolution of an existing product through use and observation; or straightforward copying of other, similar products.

Cross-fertilization may result from discussing ideas with other designers, while Bill Buxton (2007) reports that different perspectives from users generated original ideas about alternative designs. As an example of evolution, consider the cell phone and its descendant, the smartphone. The capabilities of the phone in your pocket have increased from the time they first appeared. Initially, the cell phone simply made and received phone calls and texts, but now the smartphone supports a myriad of interactions, can take photos and record audio, play movies and games, and record your exercise routine.

Creativity and invention are often wrapped in mystique, but a lot has been uncovered about the process and of how creativity can be enhanced or inspired (for example, see Rogers, 2014). For instance, browsing a collection of designs will inspire designers to consider alternative perspectives and hence alternative solutions. As Roger Schank (1982, p. 22) puts it, “An expert is someone who gets reminded of just the right prior experience to help him in processing his current experiences.” And while those experiences may be the designer’s own, they can equally well be others’.

Another approach to creativity has been adopted by Neil Maiden et al. (2007). They ran creativity workshops to generate innovative requirements in an air traffic management (ATM) application domain. Their idea was to introduce experts in different fields into the workshop and then invite stakeholders to identify analogies between their own field and this new one. For example, they have invited an Indian textile expert, a musician, a TV program scheduler, and a museum exhibit designer. Although not all obviously analogical domains, they sparked creative ideas for the air traffic management application. For example, participants reported that one textile design was elegant, that is, simple, beautiful, and symmetrical. They then transferred these properties to a key area of the ATM domain—that of aircraft conflict resolution. They explored the meaning of elegance within this context and realized that elegance is perceived differently by different controllers. From this they generated the requirement that the system should be able to accommodate different air traffic controller styles during conflict resolution.

A more pragmatic answer to this question, then, is that alternatives come from seeking different perspectives and looking at other designs. The process of inspiration and creativity can be enhanced by prompting a designer's own experience and studying others' ideas and suggestions. Deliberately seeking out suitable sources of inspiration is a valuable step in any design process. These sources may be very close to the intended new product, such as competitors' products; they may be earlier versions of similar systems; or they may be from a completely different domain.

Under some circumstances, the scope to consider alternative designs is limited. Design is a process of balancing constraints and trading off one set of requirements with another, and the constraints may mean that there are few viable alternatives available. For example, when designing software to run under the Windows operating system, the design must conform to the Windows look and feel and to other constraints intended to make Windows programs consistent for the user. When producing an upgrade to an existing system, keeping familiar elements of it to retain the same user experience may be prioritized.

ACTIVITY 2.4

Consider the product introduced in Activity 2.1. Reflecting on the process again, what inspired your initial design? Are there any innovative aspects to it?

Comment

For our design, existing sources of information and their flaws were influential. For example, there is so much information available about travel, destinations, hotel comparisons, and so forth, that it can be overwhelming. However, travel blogs contain useful and practical insights, and websites that compare alternative options are informative. We were also influenced by some favorite mobile and desktop applications such as the United Kingdom's National Rail smartphone app for its real-time updating and by the Airbnb website for its mixture of simplicity and detail.

Perhaps you were inspired by something that you use regularly, like a particularly enjoyable game or a device that you like to use? I'm not sure how innovative our ideas were, but the main goal was for the application to tailor its advice for the user's preferences. There are probably other aspects that make your design unique and that may be innovative to a greater or lesser degree. ■

2.3.4 How to Choose Among Alternative Designs

Choosing among alternatives is mostly about making design decisions: Will the device use keyboard entry or a touch screen? Will the product provide an automatic memory function or not? These decisions will be informed by the information gathered about users and their tasks and by the technical feasibility of an idea. Broadly speaking, though, the decisions fall into two categories: those that are about externally visible and measurable features

BOX 2.5

A Box Full of Ideas

The innovative product design company IDEO was mentioned in Chapter 1. Underlying some of its creative flair is a collection of weird and wonderful engineering housed in a large flatbed filing cabinet called the TechBox. The TechBox holds hundreds of gizmos and interesting materials, divided into categories such as Amazing Materials, Cool Mechanisms, Interesting Manufacturing Processes, Electronic Technologies, and Thermal and Optical. Each item has been placed in the box because it represents a neat idea or a new process. The staff at IDEO take along a selection of items from the TechBox to brainstorming meetings. The items may be chosen because they provide useful visual props or possible solutions to a particular issue or simply to provide some light relief.

Each item is clearly labeled with its name and category, but further information can be found by accessing the TechBox's online catalog. Each item has its own page detailing what the item is, why it is interesting, where it came from, and who has used it or knows more about it. Items in the box include an example of metal-coated wood and materials with and without holes that stretch, bend, and change shape or color at different temperatures.

Each of IDEO's offices has a TechBox, and each TechBox has its own curator who is responsible for maintaining and cataloging the items and for promoting its use within the office. Anyone can submit a new item for consideration. As items become commonplace, they are removed from the TechBox to make way for the next generation of fascinating curios. ■

and those that are about characteristics internal to the system that cannot be observed or measured without dissecting it. For example, in a photocopier, externally visible and measurable factors include the physical size of the machine, the speed and quality of copying, the different sizes of paper it can use, and so on. Underlying each of these factors are

DILEMMA

Copying for Inspiration: Is It Legal?

Designers draw on their experience of design when approaching a new project. This includes the use of previous designs that they know work—both designs that they have created themselves and those that others have created. Others' creations often spark inspiration that also leads to new ideas and innovation. This is well known and understood. However, the expression of an idea is protected by copyright, and people who infringe on that copyright can be taken to court and prosecuted. Note that copyright covers the expression of an idea and not the idea itself. This means, for example, that while there are numerous smartphones all with similar functionality, this does not represent an infringement of copyright as the idea has been expressed in different ways and it is the expression that has been copyrighted. Copyright is free and is automatically invested in the author, for instance, the writer of a book or a programmer who develops a program, unless they sign the copyright over to someone else. Employment contracts often include a statement that the copyright relating to anything produced in the course of that employment is automatically assigned to the employer and does not remain with the employee.

Patenting is an alternative to copyright that does protect the idea rather than the expression of the idea. There are various forms of patenting, each of which is designed to allow the inventor to capitalize on their idea. For example, Amazon patented its one-click purchasing process, which allows regular users simply to choose a purchase and buy it with one mouse click (US Patent No. 5960411, September 29, 1999). This is possible because the system stores its customers' details and recognizes them when they access the Amazon site again.

In recent years, the creative commons community (<https://creativecommons.org/>) has suggested more flexible licensing arrangements that allow others to reuse and extend a piece of created work, thereby supporting collaboration. In the open source software development movement, for example, software code is freely distributed and can be modified, incorporated into other software, and redistributed under the same open source conditions. No royalty fees are payable on any use of open source code. These movements do not replace copyright or patent law, but they provide an alternative route for the dissemination of ideas.

So, the dilemma comes in knowing when it is OK to use someone else's work as a source of inspiration and when you are infringing copyright or patent law. The issues are complex and detailed and well beyond the scope of this book, but Bainbridge (2014) is a good resource to understand this area better. ■

other considerations that cannot be observed or studied without dissecting the machine. For example, the choice of materials used in a photocopier may depend on its friction rating and how much it deforms under certain conditions. In interaction design, the user experience is the driving force behind the design and so externally visible and measurable behavior is the main focus. Detailed internal workings are still important to the extent that they affect external behavior or features.

One answer to the previous question is that choosing between alternative designs is informed by letting users and stakeholders interact with them and by discussing their experiences, preferences, and suggestions for improvement. To do this, the designs must be in a form that can be reasonably evaluated by users, not in technical jargon or notation that seems impenetrable to them. Documentation is one traditional way to communicate a design, for example, a diagram showing the product's components or a description of how it works. But a static description cannot easily capture the dynamics of behavior, and for an interactive product this needs to be communicated so that users can see what it will be like to operate it.

Prototyping is often used to overcome potential client misunderstandings and to test the technical feasibility of a suggested design and its production. It involves producing a limited version of the product with the purpose of answering specific questions about the design's feasibility or appropriateness. Prototypes give a better impression of the user experience than simple descriptions; different kinds of prototyping are suitable for different stages of development and for eliciting different kinds of feedback. When a deployable version of the product is available, another way to choose between alternative designs is to deploy two different variations and collect data from actual use that is then used to inform the choice. This is called *A/B testing*, and it is often used for alternative website designs (see Box 2.5).

Another basis on how to choose between alternatives is quality, but that requires a clear understanding of what quality means, and people's views of quality vary. Everyone has a notion of the level of quality that is expected, wanted, or needed from a product. Whether this is expressed formally, informally, or not at all, it exists and informs the choice between alternatives. For example, one smartphone design might make it easy to access a popular music channel but restrict sound settings, while another requires more complicated key sequences to access the channel but has a range of sophisticated sound settings. One user's view of quality may lean toward ease of use, while another may lean toward sophisticated sound settings.

Most projects involve a range of different stakeholder groups, and it is common for each of them to define quality differently and to have different acceptable limits for it. For example, although all stakeholders may agree on goals for a video game such as "characters will be appealing" or "graphics will be realistic," the meaning of these statements can vary between different groups. Disputes will arise if, later in development, it transpires that "realistic" to a stakeholder group of teenage players is different from "realistic" to a group of parent stakeholders or to developers. Capturing these different views clearly clarifies expectations, provides a benchmark against which products and prototypes can be compared, and forms a basis on which to choose among alternatives.

The process of writing down formal, verifiable—and hence measurable—usability criteria is a key characteristic of an approach to interaction design called usability engineering. This has emerged over many years and with various proponents (Whiteside et al., 1988; Nielsen, 1993). Most recently, it is often applied in health informatics (for example, see Kushniruk et al., 2015). Usability engineering involves specifying quantifiable measures of product performance, documenting them in a usability specification, and assessing the product against them.

BOX 2.5

A/B Testing

A/B testing is an online method to inform the choice between two alternatives. It is most commonly used for comparing different versions of web pages or apps, but the principles and mathematics behind it came about in the 1920s (Gallo, 2017). In an interaction design context, different versions of web pages or apps are released for use by users performing their everyday tasks. Typically, users are unaware that they are contributing to an evaluation. This is a powerful way to involve users in choosing between alternatives because a huge number of users can be involved and the situations are authentic.

On the one hand, it's a simple idea—give one set of users one version and a second set the other version and see which set scores more highly against the success criteria. But dividing up the sets, choosing the success criteria, and working out the metrics to use are nontrivial (for example, see Deng and Shi, 2016). Pushing this idea further, it is common to have “multivariate” testing in which several options are tried at once, so you end up doing A/B/C testing or even A/B/C/D testing. ■

ACTIVITY 2.5

Consider your product from Activity 2.1. Suggest some usability criteria that could be applied to determine its quality. Use the usability goals introduced in Chapter 1—effectiveness, efficiency, safety, utility, learnability, and memorability. Be as specific as possible. Check the criteria by considering exactly what to measure and how to measure its performance.

Then try to do the same thing for some of the user experience goals introduced in Chapter 1. (These relate to whether a system is satisfying, enjoyable, motivating, rewarding, and so on.)

Comment

Finding measurable characteristics for some of these is not easy. Here are some suggestions, but there are others. Where possible, criteria that are measurable and specific are preferable.

- *Effectiveness*: Identifying measurable criteria for this goal is particularly difficult since it is a combination of the other goals. For example, does the system support travel organization, choosing transport routes, booking accommodation, and so on? In other words, is the product used?
- *Efficiency*: Is it clear how to ask for recommendations from the product? How quickly does it identify a suitable route or destination details?

(Continued)

- *Safety*: How often does data get lost or is the wrong option chosen? This may be measured, for example, as the number of times this happens per trip.
- *Utility*: How many functions offered are used for every trip, how many every other trip, and how many are not used at all? How many tasks are difficult to complete in a reasonable time because functionality is missing or the right subtasks aren't supported?
- *Learnability*: How long does it take for a novice user to be able to do a series of set tasks, for example, to book a hotel room in Paris near the meeting venue for the meeting dates, identify appropriate flights from Sydney to Wellington, or find out whether a visa is needed to go to China?
- *Memorability*: If the product isn't used for a month, how many functions can the user remember how to perform? How long does it take to remember how to perform the most frequent task?

Finding measurable characteristics for the user experience criteria is harder. How do you measure satisfaction, fun, motivation, or aesthetics? What is entertaining to one person may be boring to another; these kinds of criteria are subjective and so cannot be measured as objectively. ■

2.3.5 How to Integrate Interaction Design Activities Within Other Lifecycle Models

As illustrated in Chapter 1 (Figure 1.4), many other disciplines contribute to interaction design, and some of these disciplines have lifecycles of their own. Prominent among them are those associated with software development, and integrating interaction design activities within software development has been discussed for many years; for example, see Carmelo Ardito et al. (2014) and Ahmed Seffah et al. (2005).

The latest attempts to integrate these practices focus on agile software development. Agile methods began to emerge in the late 1990s. The most well-known of these are eXtreme Programming (Beck and Andres, 2005), Scrum (Schwaber and Beedle, 2002), and Kanban (Anderson, 2010). The Dynamic Systems Development Method (DSDM) (DSDM, 2014), although established before the current agile movement, also belongs to the agile family as it adheres to the agile manifesto. These methods differ, but they all stress the importance of iteration, early and repeated user feedback, being able to handle emergent requirements, and striking a good balance between flexibility and structure. They also all emphasize collaboration, face-to-face communication, streamlined processes to avoid unnecessary activities, and the importance of practice over process, that is, of getting work done.

The opening statement for the *Manifesto for Agile Software Development* (www.agilemanifesto.org/) reads as follows:

We are uncovering better ways of developing software by doing it and helping others do it. Through this work we have come to value:

- *Individuals and interactions over processes and tools*
- *Working software over comprehensive documentation*
- *Customer collaboration over contract negotiation*
- *Responding to change over following a plan*

This manifesto is underpinned by a series of principles, which range from communication with the business to excellence of coding and maximizing the amount of work done. The agile approach to development is particularly interesting from the point of view of interaction design because it incorporates tight iterations and feedback and collaboration with the customer. For example, in Scrum, each sprint is between one and four weeks, with a product of value being delivered at the end of each sprint. Also, eXtreme¹ Programming (XP) stipulates that the customer should be on-site with developers. In practice, the customer role is usually taken by a team rather than by one person (Martin et al., 2009), and integration is far from straightforward (Ferreira et al., 2012). Many companies have integrated agile methods with interaction design practices to produce a better user experience and business value (Loranger and Laubheimer, 2017), but it is not necessarily easy, as discussed in Chapter 13, “Interaction Design in Practice.”

In-Depth Activity

These days, timepieces (such as clocks, wristwatches, and so on) have a variety of functions. Not only do they tell the time and date, but they can speak to you, remind you when it’s time to do something, and record your exercise habits among other things. The interface for these devices, however, shows the time in one of two basic ways: as a digital number such as 11:40 or through an analog display with two or three hands—one to represent the hour, one for the minutes, and one for the seconds.

This in-depth activity is to design an innovative timepiece. This could be in the form of a wristwatch, a mantelpiece clock, a sculpture for a garden or balcony, or any other kind of timepiece you prefer. The goal is to be inventive and exploratory by following these steps:

- (a) Think about the interactive product that you are designing: What do you want it to do? Find three to five potential users, and ask them what they would want. Write a list of requirements for the clock, together with some usability criteria and user experience criteria based on the definitions in Chapter 1.
- (b) Look around for similar devices and seek out other sources of inspiration that you might find helpful. Make a note of any findings that are interesting, useful, or insightful.
- (c) Sketch some initial designs for the timepiece. Try to develop at least two distinct alternatives that meet your set of requirements.
- (d) Evaluate the two designs by using your usability criteria and by role-playing an interaction with your sketches. Involve potential users in the evaluation, if possible. Does it do what you want? Is the time or other information being displayed always clear? Design is iterative, so you may want to return to earlier elements of the process before you choose one of your alternatives.

¹The method is called *extreme* because it pushes a key set of good practices to the limit; that is, it is good practice to test often, so in XP the development is test-driven, and a complete set of tests is executed many times a day. It is good practice to talk to people about their requirements, so rather than having weighty documentation, XP reduces documentation to a minimum, thus forcing communication, and so on.

Summary

In this chapter, we looked at user-centered design and the process of interaction design. That is, what is user-centered design, what activities are required in order to design an interactive product, and how are these activities related? A simple interaction design lifecycle model consisting of four activities was introduced, and issues surrounding the involvement and identification of users, generating alternative designs, evaluating designs, and integrating user-centered concerns with other lifecycles were discussed.

Key Points

- Different design disciplines follow different approaches, but they have commonalities that are captured in the double diamond of design.
- It is important to have a good understanding of the problem space before trying to build anything.
- The interaction design process consists of four basic activities: discover requirements, design alternatives that meet those requirements, prototype the designs so that they can be communicated and assessed, and evaluate them.
- User-centered design rests on three principles: early focus on users and tasks, empirical measurement, and iterative design. These principles are also key for interaction design.
- Involving users in the design process assists with expectation management and feelings of ownership, but how and when to involve users requires careful planning.
- There are many ways to understand who users are and what their goals are in using a product, including rapid iterations of working prototypes.
- Looking at others' designs and involving other people in design provides useful inspiration and encourages designers to consider alternative design solutions, which is key to effective design.
- Usability criteria, technical feasibility, and users' feedback on prototypes can all be used to choose among alternatives.
- Prototyping is a useful technique for facilitating user feedback on designs at all stages.
- Interaction design activities are becoming better integrated with lifecycle models from other related disciplines such as software engineering.

Further Reading

ASHMORE, S. and RUNYAN, K. (2015) *Introduction to Agile Methods*, Addison Wesley. This book introduces the basics of agile software development and the most popular agile methods in an accessible way. It touches on usability issues and the relationship between agile and marketing. It is a good place to start for someone new to the agile way of working.

KELLEY, T., with LITTMAN, J. (2016) *The Art of Innovation*, Profile Books. Tom Kelley is a partner at IDEO. In this book, Kelley explains some of the innovative techniques used at

IDEO, but more importantly he talks about the culture and philosophy underlying IDEO's success. There are some useful practical hints in here as well as an informative story about building and maintaining a successful design company.

PRESSMAN, R.S. and MAXIM, B.R. (2014) *Software Engineering: A Practitioner's Approach (Int'l Ed)*, McGraw-Hill Education. If you are interested in pursuing the software engineering aspects of the lifecycle models section, then this book provides a useful overview of the main models and their purpose.

SIROKER, D. and KOOMEN, P. (2015) *A/B Testing: The Most Powerful Way to Turn Clicks into Customers*, John Wiley. This book is written by two experienced practitioners who have been using A/B testing with a range of organizations. It is particularly interesting because of the example cases that show the impact that applying A/B testing successfully can have.

ROGERS, Y. (2014) *Secrets of Creative People* (PDF available from www.id-book.com/). This short book summarizes the findings from a two-year research project into creativity. It emphasizes the importance of different perspectives to creativity and describes how successful creativity arises from sharing, constraining, narrating, connecting, and even sparring with others.

Chapter 3

CONCEPTUALIZING INTERACTION

3.1 Introduction

3.2 Conceptualizing Interaction

3.3 Conceptual Models

3.4 Interface Metaphors

3.5 Interaction Types

3.6 Paradigms, Visions, Theories, Models, and Frameworks

Objectives

The main goals of this chapter are to accomplish the following:

- Explain how to conceptualize interaction.
- Describe what a conceptual model is and how to begin to formulate one.
- Discuss the use of interface metaphors as part of a conceptual model.
- Outline the core interaction types for informing the development of a conceptual model.
- Introduce paradigms, visions, theories, models, and frameworks informing interaction design.

3.1 Introduction

When coming up with new ideas as part of a design project, it is important to conceptualize them in terms of what the proposed product will do. Sometimes, this is referred to as creating a *proof of concept*. In relation to the double diamond framework, it can be viewed as an initial pass to help define the area and also when exploring solutions. One reason for needing to do this is as a reality check where fuzzy ideas and assumptions about the benefits of the proposed product are scrutinized in terms of their feasibility: How realistic is it to develop what they have suggested, and how desirable and useful will it actually be? Another reason is to enable designers to begin articulating what the basic building blocks will be when developing the product. From a user experience (UX) perspective, it can lead to better clarity, forcing designers to explain how users will understand, learn about, and interact with the product.

For example, consider the bright idea that a designer has of creating a voice-assisted mobile robot that can help waiters in a restaurant take orders and deliver meals to customers (see Figure 3.1). The first question to ask is: why? What problem would this address? The designer might say that the robot could help take orders and entertain customers by having a conversation with them at the table. They could also make recommendations that can be customized to different customers, such as restless children or fussy eaters. However, none of these addresses an actual problem. Rather, they are couched in terms of the putative benefits of the new solution. In contrast, an actual problem identified might be the following: “It is difficult to recruit good wait staff who provide the level of customer service to which we have become accustomed.”



Figure 3.1 A nonspeaking robot waiter in Shanghai. What would be gained if it could also talk with customers?

Source: ZUMA Press / Alamy Stock Photo

Having worked through a problem space, it is important to generate a set of research questions that need to be addressed, when considering how to design a robot voice interface to wait on customers. These might include the following: How intelligent does it have to be? How would it need to move to appear to be talking? What would the customers think of it? Would they think it is too gimmicky and get easily tired of it? Or, would it always be a pleasure for them to engage with the robot, not knowing what it would say on each new visit to the restaurant? Could it be designed to be a grumpy extrovert or a funny waiter? What might be the limitations of this voice-assisted approach?

Many unknowns need to be considered in the initial stages of a design project, especially if it is a new product that is being proposed. As part of this process, it can be useful to show where your novel ideas came from. What sources of inspiration were used? Is there any theory or research that can be used to inform and support the nascent ideas?

Asking questions, reconsidering one’s assumptions, and articulating one’s concerns and standpoints are central aspects of the early ideation process. Expressing ideas as a set of concepts greatly helps to transform blue-sky and wishful thinking into more concrete models of

how a product will work, what design features to include, and the amount of functionality that is needed. In this chapter, we describe how to achieve this through considering the different ways of conceptualizing interaction.

3.2 Conceptualizing Interaction

When beginning a design project, it is important to be clear about the underlying assumptions and claims. By an assumption, we mean taking something for granted that requires further investigation; for example, people now want an entertainment and navigation system in their cars. By a claim, we mean stating something to be true when it is still open to question. For instance, a multimodal style of interaction for controlling this system—one that involves speaking or gesturing while driving—is perfectly safe.

Writing down your assumptions and claims and then trying to defend and support them can highlight those that are vague or wanting. In so doing, poorly constructed design ideas can be reformulated. In many projects, this process involves identifying human activities and interactivities that are problematic and working out how they might be improved through being supported with a different set of functions. In others, it can be more speculative, requiring thinking through how to design for an engaging user experience that does not exist.

Box 3.1 presents a hypothetical scenario of a team working through their assumptions and claims; this shows how, in so doing, problems are explained and explored and leads to a specific avenue of investigation agreed on by the team.

BOX 3.1

Working Through Assumptions and Claims

This is a hypothetical scenario of early design highlighting the assumptions and claims (*italicized*) made by different members of a design team.

A large software company has decided that it needs to develop an upgrade of its web browser for smartphones because its marketing team has discovered that many of the company's customers have switched over to using another mobile browser. The marketing people *assume* that something is wrong with their browser and that their rivals have a better product. But they don't know what the problem is with their browser.

The design team put in charge of this project *assumes* that they need to improve the usability of a number of the browser's functions. They *claim* that this will win back users by making features of the interface simpler, more attractive, and more flexible to use.

The user researchers on the design team conduct an initial user study investigating how people use the company's web browser on a variety of smartphones. They also look at other mobile web browsers on the market and compare their functionality and usability. They observe and talk to many different users. They discover several things about the usability of their web browser, some of which they were not expecting. One revelation is that many of their customers have never actually used the bookmarking tool. They present their findings to the rest of the team and have a long discussion about why each of them thinks it is not being used.

One member *claims* that the web browser's function for organizing bookmarks is tricky and error-prone, and she *assumes* that this is the reason why many users do not use it. Another member backs her up, saying how awkward it is to use this method when wanting to move bookmarks between folders. One of the user experience architects agrees, noting how several of the users with whom he spoke mentioned how difficult and time-consuming they found it when trying to move bookmarks between folders and how they often ended up accidentally putting them into the wrong folders.

A software engineer reflects on what has been said, and he makes the *claim* that the bookmark function is no longer needed since he *assumes* that most people do what he does, which is to revisit a website by flicking through their history of previously visited pages. Another member of the team disagrees with him, *claiming* that many users do not like to leave a trail of the sites they have visited and would prefer to be able to save only the sites that they think they might want to revisit. The bookmark function provides them with this option. Another option discussed is whether to include most-frequently visited sites as thumbnail images or as tabs. The software engineer agrees that providing all of the options could be a solution but worries how this might clutter the small screen interface.

After much discussion on the pros and cons of bookmarking versus history lists, the team decides to investigate further how to support effectively the saving, ordering, and retrieving of websites using a mobile web browser. All agree that the format of the existing web browser's structure is too rigid and that one of their priorities is to see how they can create a simpler way of revisiting websites on a smartphone. ■

Explaining people's assumptions and claims about why they think something might be a good idea (or not) enables the design team as a whole to view multiple perspectives on the problem space and, in so doing, reveals conflicting and problematic ones. The following framework is intended to provide a set of core questions to aid design teams in this process:

- Are there problems with an existing product or user experience? If so, what are they?
- Why do you think there are problems?
- What evidence do you have to support the existence of these problems?
- How do you think your proposed design ideas might overcome these problems?

ACTIVITY 3.1

Use the framework in the previous list to guess what the main assumptions and claims were behind 3D TV. Then do the same for curved TV, which was designed to be bendy so as to make the viewing experience more immersive. Are the assumptions similar? Why were they problematic?

Comment

There was much hype and fanfare about the enhanced user experience 3D and curved TVs would offer, especially when watching movies, sports events, and dramas (see Figure 3.2).

However, both never really took off. Why was this? One *assumption* for 3D TV was that people would not mind wearing the glasses that were needed to see in 3D, nor would they mind paying a lot more for a new 3D-enabled TV screen. A *claim* was that people would really enjoy the enhanced clarity and color detail provided by 3D, based on the favorable feedback received worldwide when viewing 3D films, such as *Avatar*, at a cinema. Similarly, an *assumption* made about curved TV was that it would provide more flexibility for viewers to optimize the viewing angles in someone's living room.



Figure 3.2 A family watching 3D TV

Source: Andrey Popov/Shutterstock

The unanswered question for both concepts was this: Could the enhanced cinema viewing experience that both *claimed* become an actual desired living room experience? There was no existing problem to overcome—what was being proposed was a new way of experiencing TV. The problem they might have *assumed* existed was that the experience of viewing TV at home was inferior to that of the cinema. The *claim* could have been that people would be prepared to pay more for a better-quality viewing experience more akin to that of the cinema.

But were people prepared to pay extra for a new TV because of this enhancement? A number of people did. However, a fundamental usability problem was overlooked—many people complained of motion sickness when watching 3D TV. The glasses were also easily lost. Moreover, wearing them made it difficult to do other things such as flicking through multiple channels, texting, and tweeting. (Many people simultaneously use additional devices, such as smartphones and tablets while watching TV.) Most people who bought 3D TVs stopped watching them after a while because of these usability problems. While curved TV didn't require viewers to wear special glasses, it also failed because the actual benefits were not that significant relative to the cost. While for some the curve provided a cool aesthetic look and an improved viewing angle, for others it was simply an inconvenience. ■

Making clear what one's assumptions are about a problem and the claims being made about potential solutions should be carried out early on and throughout a project. Design teams also need to work out how best to conceptualize the design space. Primarily, this involves articulating the proposed solution as a conceptual model with respect to the user experience. The benefits of conceptualizing the design space in this way are as follows:

Orientation Enabling the design team to ask specific kinds of questions about how the conceptual model will be understood by the targeted users.

Open-Mindedness Allowing the team to explore a range of different ideas to address the problems identified.

Common Ground Allowing the design team to establish a set of common terms that all can understand and agree upon, reducing the chance of misunderstandings and confusion arising later.

Once formulated and agreed upon, a conceptual model can then become a shared blueprint leading to a testable proof of concept. It can be represented as a textual description and/or in a diagrammatic form, depending on the preferred *lingua franca* used by the design team. It can be used not just by user experience designers but also to communicate ideas to business, engineering, finance, product, and marketing units. The conceptual model is used by the design team as the basis from which they can develop more detailed and concrete aspects of the design. In doing so, design teams can produce simpler designs that match up with users' tasks, allow for faster development time, result in improved customer uptake, and need less training and customer support (Johnson and Henderson, 2012).

3.3 Conceptual Models

A *model* is a simplified description of a system or process that helps describe how it works. In this section, we look at a particular kind of model used in interaction design intended to articulate the problem and design space—the *conceptual model*. In a later section, we describe more generally how models have been developed to explain phenomena in human-computer interaction.

Jeff Johnson and Austin Henderson (2002) define a conceptual model as “a high-level description of how a system is organized and operates” (p. 26). In this sense, it is an abstraction outlining what people can do with a product and what concepts are needed to understand how to interact with it. A key benefit of conceptualizing a design at this level is that it enables “designers to straighten out their thinking before they start laying out their widgets” (p. 28).

In a nutshell, a conceptual model provides a working strategy and a framework of general concepts and their interrelations. The core components are as follows:

- Metaphors and analogies that convey to people how to understand what a product is used for and how to use it for an activity (for example browsing and bookmarking).
- The concepts to which people are exposed through the product, including the task-domain objects they create and manipulate, their attributes, and the operations that can be performed on them (such as saving, revisiting, and organizing).
- The relationships between those concepts (for instance, whether one object contains another).

- The mappings between the concepts and the user experience the product is designed to support or invoke (for example, one can revisit a page through looking at a list of visited sites, most-frequently visited, or saved websites).

How the various metaphors, concepts, and their relationships are organized determines the user experience. By explaining these, the design team can debate the merits of providing different methods and how they support the main concepts, for example, saving, revisiting, categorizing, reorganizing, and their mapping to the task domain. They can also begin discussing whether a new overall metaphor may be preferable that combines the activities of browsing, searching, and revisiting. In turn, this can lead the design team to articulate the kinds of relationships between them, such as containership. For example, what is the best way to sort and revisit saved pages, and how many and what types of containers should be used (for example, folders, bars, or panes)? The same enumeration of concepts can be repeated for other functions of the web browser—both current and new. In so doing, the design team can begin to work out systematically what will be the simplest and most effective and memorable way of supporting users while browsing the Internet.

The best conceptual models are often those that appear obvious and simple; that is, the operations they support are intuitive to use. However, sometimes applications can end up being based on overly complex conceptual models, especially if they are the result of a series of upgrades, where more and more functions and ways of doing something are added to the original conceptual model. While tech companies often provide videos showing what new features are included in an upgrade, users may not pay much attention to them or skip them entirely. Furthermore, many people prefer to stick to the methods they have always used and trusted and, not surprisingly, become annoyed when they find one or more have been removed or changed. For example, when Facebook rolled out its revised newsfeed a few years back, many users were unhappy, as evidenced by their postings and tweets, preferring the old interface that they had gotten used to. A challenge for software companies, therefore, is how best to introduce new features that they have added to an upgrade—and explain their assumed benefits to users—while also justifying why they removed others.

BOX 3.2

Design Concept

Another term that is sometimes used is a *design concept*. Essentially, it is a set of ideas for a design. Typically, it is composed of scenarios, images, mood boards, or text-based documents. For example, Figure 3.3 shows the first page of a design concept developed for an ambient display that was aimed at changing people's behavior in a building, that is, to take the stairs instead of the elevator. Part of the design concept was envisioned as an animated pattern of twinkly lights that would be embedded in the carpet near the entrance of the building with the intention of luring people toward the stairs (Hazlewood et al., 2010). ■

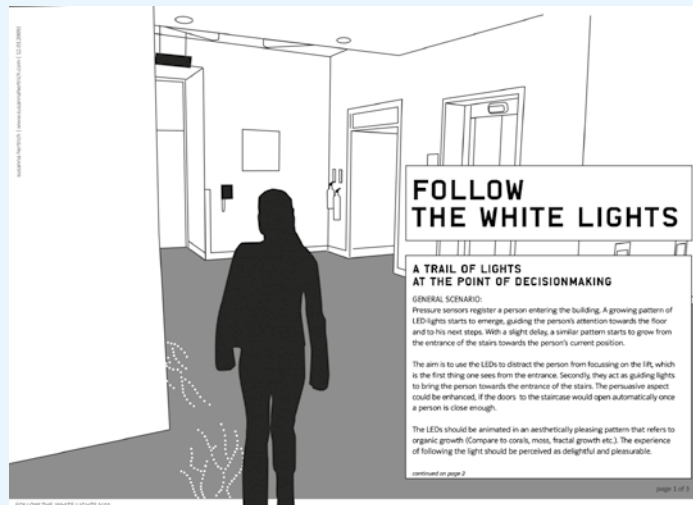


Figure 3.3 The first page of a design concept for an ambient display

Most interface applications are actually based on well-established conceptual models. For example, a conceptual model based on the core aspects of the customer experience when at a shopping mall underlies most online shopping websites. These include the placement of items that a customer wants to purchase into a shopping cart or basket and proceeding to checkout when they're ready to make the purchase. Collections of patterns are now readily available to help design the interface for these core transactional processes, together with many other aspects of a user experience, meaning interaction designers do not have to start from scratch every time they design or redesign an application. Examples include patterns for online forms and navigation on mobile phones.

It is rare for completely new conceptual models to emerge that transform the way daily and work activities are carried out at an interface. Those that did fall into this category include the following three classics: the desktop (developed by Xerox in the late 1970s), the digital spreadsheet (developed by Dan Bricklin and Bob Frankston in the late 1970s), and the World Wide Web (developed by Tim Berners Lee in the early 1980s). All of these innovations made what was previously limited to a few skilled people accessible to all, while greatly expanding what is possible. The graphical desktop dramatically changed how office tasks could be performed (including creating, editing, and printing documents). Performing these tasks using the computers prevalent at the time was significantly more arduous, having to learn and use a command language (such as DOS or UNIX). Digital spreadsheets made accounting highly flexible and easier to accomplish, enabling a diversity of new computations to be performed simply through filling in interactive boxes. The World Wide Web allowed anyone to browse a network of information remotely. Since then, e-readers and digital authoring tools have introduced new ways of reading documents and books online, supporting associated activities such as annotating, highlighting, linking, commenting, copying,

and tracking. The web has also enabled and made many other kinds of activities easier, such as browsing for news, weather, sports, and financial information, as well as banking, shopping, and learning online among other tasks. Importantly, all of these conceptual models were based on familiar activities.

BOX 3.3

A Classic Conceptual Model: The Xerox Star

The Star interface, developed by Xerox in 1981 (see Figure 3.4), revolutionized the way that interfaces were designed for personal computing (Smith et al., 1982; Miller and Johnson, 1996) and is viewed as the forerunner of today's Mac and Windows desktop interfaces. Originally, it was designed as an office system, targeted at workers not interested in computing *per se*, and it was based on a conceptual model that included the familiar knowledge of an office. Paper, folders, filing cabinets, and mailboxes were represented as icons on the screen and were designed to possess some of the properties of their physical counterparts. Dragging a document icon across the desktop screen was seen as equivalent to picking up a piece of paper in the physical world and moving it (but this, of course, is a very different action). Similarly, dragging a digital document into a digital folder was seen as being analogous to placing a physical document into a physical cabinet. In addition, new concepts that were incorporated as part of the desktop metaphor were operations that could not be performed in the physical world. For example, digital files could be placed onto an icon of a printer on the desktop, resulting in the computer printing them out. ■



Figure 3.4 The Xerox Star

Source: Used courtesy of Xerox

Video The history of the Xerox Star at <http://youtu.be/Cn4vC80Pv6Q>.

3.4 Interface Metaphors

Metaphors are considered to be a central component of a conceptual model. They provide a structure that is similar in some way to aspects of a familiar entity (or entities), but they also have their own behaviors and properties. More specifically, an *interface metaphor* is one that is instantiated in some way as part of the user interface, such as the desktop metaphor. Another well-known one is the *search engine*, originally coined in the early 1990s to refer to a software tool that indexed and retrieved files remotely from the Internet using various algorithms to match terms selected by the user. The metaphor invites comparisons between a mechanical engine, which has several working parts, and the everyday action of looking in different places to find something. The functions supported by a search engine also include other features besides those belonging to an engine that searches, such as listing and prioritizing the results of a search. It also does these actions in quite different ways from how a mechanical engine works or how a human being might search a library for books on a given topic. The similarities implied by the use of the term *search engine*, therefore, are at a general level. They are meant to conjure up the essence of the process of finding relevant information, enabling the user to link these to less familiar aspects of the functionality provided.

ACTIVITY 3.2

Go to a few online stores and see how the interface has been designed to enable the customer to order and pay for an item. How many use the “add to shopping cart/basket” followed by the “checkout” metaphor? Does this make it straightforward and intuitive to make a purchase?

Comment

Making a purchase online usually involves spending money by inputting one’s credit/debit card details. People want to feel reassured that they are doing this correctly and do not get frustrated with lots of forms to fill in. Designing the interface to have a familiar metaphor (with an icon of a shopping cart/basket, not a cash register) makes it easier for people to know what to do at the different stages of making a purchase. Most important, placing an item in the basket does not commit the customer to purchase it there and then. It also enables them to browse further and select other items, as they might in a physical store. ■

Interface metaphors are intended to provide familiar entities that enable people readily to understand the underlying conceptual model and know what to do at the interface. However, they can also contravene people’s expectations about how things should be, such

as the recycle bin (trash can) that sits on the desktop. Logically and culturally (meaning, in the real world), it should be placed under the desk. But users would not have been able to see it because it would have been hidden by the desktop surface. So, it needed to go on the desktop. While some users found this irksome, most did not find it to be a problem. Once they understood why the recycle bin icon was on the desktop, they simply accepted it being there.

An interface metaphor that has become popular in the last few years is the card. Many of the social media apps, such as Facebook, Twitter, and Pinterest, present their content on cards. Cards have a familiar form, having been around for a long time. Just think of how many kinds there are: playing cards, business cards, birthday cards, credit cards, and postcards to name a few. They have strong associations, providing an intuitive way of organizing limited content that is “card sized.” They can easily be flicked through, sorted, and themed. They structure content into meaningful chunks, similar to how paragraphs are used to chunk a set of related sentences into distinct sections (Babich, 2016). In the context of the smartphone interface, the Google Now card provides short snippets of useful information. This appears on and moves across the screen in the way people would expect a real card to do—in a lightweight, paper-based sort of way. The elements are also structured to appear as if they were on a card of a fixed size, rather than, say, in a scrolling web page (see Figure 3.5).

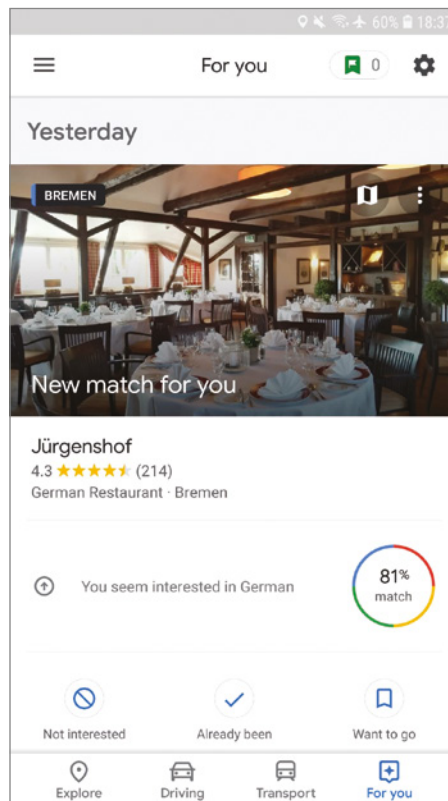


Figure 3.5 Google Now card for restaurant recommendation in Germany

Source: Used courtesy of Johannes Schöning

In many cases, new interface metaphors rapidly become integrated into common parlance, as witnessed by the way people talk about them. For example, parents talk about how much screen time children are allowed each day in the same way they talk more generally about spending time. As such, the interface metaphors are no longer talked about as familiar terms to describe less familiar computer-based actions; they have become everyday terms in their own right. Moreover, it is hard not to use metaphorical terms when talking about technology use, as they have become so ingrained in the language that we use to express ourselves. Just ask yourself or someone else to describe Twitter and Facebook and how people use them. Then try doing it without using a single metaphor.

Albrecht Schmidt (2017) suggests a pair of glasses as a good metaphor for thinking about future technologies, helping us think more about how to amplify human cognition. Just as they are seen as an extension of ourselves that we are not aware of most of the time (except when they steam up!), he asks can we design new technologies that enable users to do things without having to think about how to use them? He contrasts this “amplify” metaphor with the “tool” metaphor of a pair of binoculars that is used for a specific task—where someone consciously has to hold them up against their eyes while adjusting the lens to bring what they are looking at into focus. Current devices, like mobile phones, are designed more like binoculars, where people have to interact with them explicitly to perform tasks.

BOX 3.4

Why Are Metaphors So Popular?

People frequently use metaphors and analogies (here we use the terms interchangeably) as a source of inspiration for understanding and explaining to others what they are doing, or trying to do, in terms that are familiar to them. They are an integral part of human language (Lakoff and Johnson, 1980). Metaphors are commonly used to explain something that is unfamiliar or hard to grasp by way of comparison with something that is familiar and easy to grasp. For example, they are frequently employed in education, where teachers use them to introduce something new to students by comparing the new material with something they already understand. An example is the comparison of human evolution with a game. We are all familiar with the properties of a game: there are rules, each player has a goal to win (or lose), there are heuristics to deal with situations where there are no rules, there is the propensity to cheat when the other players are not looking, and so on. By conjuring up these properties, the analogy helps us begin to understand the more difficult concept of evolution—how it happens, what rules govern it, who cheats, and so on.

It is not surprising, therefore, to see how widely metaphors have been used in interaction design to conceptualize abstract, hard-to-imagine, and difficult-to-articulate computer-based concepts and interactions in more concrete and familiar terms and as graphical visualizations at the interface level. Metaphors and analogies are used in these three main ways:

- As a way of conceptualizing what we are doing (for instance, surfing the web)
- As a conceptual model instantiated at the interface level (for example, the card metaphor)
- As a way of visualizing an operation (such as an icon of a shopping cart into which items are placed that users want to purchase on an online shopping site) ■

3.5 Interaction Types

Another way of conceptualizing the design space is in terms of the interaction types that will underlie the user experience. Essentially, these are the ways a person interacts with a product or application. Originally, we identified four main types: instructing, conversing, manipulating, and exploring (Preece et al., 2002). A fifth type has since been proposed by Christopher Lueg et al. (2019) that we have added to ours, which they call *responding*. This refers to proactive systems that initiate a request in situations to which a user can respond, for example, when Netflix pauses a person’s viewing to ask them whether they would like to continue watching.

Deciding upon which of the interaction types to use, and why, can help designers formulate a conceptual model before committing to a particular interface in which to implement them, such as speech-based, gesture-based, touch-based, menu-based, and so on. Note that we are distinguishing here between interaction types (which we discuss in this section) and interface types (which are discussed in Chapter 7, “Interfaces”). While cost and other product constraints will often dictate which interface style can be used for a given application, considering the interaction type that will best support a user experience can highlight the potential trade-offs, dilemmas, and pros and cons.

Here, we describe in more detail each of the five types of interaction. It should be noted that they are not meant to be mutually exclusive (for example, someone can interact with a system based on different kinds of activities); nor are they meant to be definitive. Also, the label used for each type refers to the user’s action even though the system may be the active partner in initiating the interaction.

- *Instructing*: Where users issue instructions to a system. This can be done in a number of ways, including typing in commands, selecting options from menus in a windows environment or on a multitouch screen, speaking aloud commands, gesturing, pressing buttons, or using a combination of function keys.
- *Conversing*: Where users have a dialog with a system. Users can speak via an interface or type in questions to which the system replies via text or speech output.
- *Manipulating*: Where users interact with objects in a virtual or physical space by manipulating them (for instance, opening, holding, closing, and placing). Users can hone their familiar knowledge of how to interact with objects.
- *Exploring*: Where users move through a virtual environment or a physical space. Virtual environments include 3D worlds and augmented and virtual reality systems. They enable users to hone their familiar knowledge by physically moving around. Physical spaces that use sensor-based technologies include smart rooms and ambient environments, also enabling people to capitalize on familiarity.
- *Responding*: Where the system initiates the interaction and the user chooses whether to respond. For example, proactive mobile location-based technology can alert people to points of interest. They can choose to look at the information popping up on their phone or ignore it. An example is the Google Now Card, shown in Figure 3.5, which pops up a restaurant recommendation for the user to contemplate when they are walking nearby.

Besides these core activities of instructing, conversing, manipulating, exploring, and responding, it is possible to describe the specific domain and context-based activities in which users engage, such as learning, working, socializing, playing, browsing, writing, problem-solving, decision-making, and searching—just to name but a few. Malcolm McCullough (2004)

suggests describing them as situated activities, organized by work (for example, presenting to groups), home (such as resting), in town (for instance, eating), and on the road (for example, walking). The rationale for classifying activities in this way is to help designers be more systematic when thinking about the usability of technology-modified places in the environment. In the following sections we illustrate in more detail the five core interaction types and how to design applications for them.

3.5.1 Instructing

This type of interaction describes how users carry out their tasks by telling the system what to do. Examples include giving instructions to a system to perform operations such as tell the time, print a file, and remind the user of an appointment. A diverse range of products has been designed based on this model, including home entertainment systems, consumer electronics, and computers. The way in which the user issues instructions can vary from pressing buttons to typing in strings of characters. Many activities are readily supported by giving instructions.

In Windows and other *graphical user interfaces (GUIs)*, control keys or the selection of menu options via a mouse, touch pad, or touch screen are used. Typically, a wide range of functions are provided from which users have to select when they want to do something to the object on which they are working. For example, a user writing a report using a word processor will want to format the document, count the number of words typed, and check the spelling. The user instructs the system to do these operations by issuing appropriate commands. Typically, commands are carried out in a sequence, with the system responding appropriately (or not) as instructed.

One of the main benefits of designing an interaction based on issuing instructions is that the interaction is quick and efficient. It is particularly fitting where there is a frequent need to repeat actions performed on multiple objects. Examples include the repetitive actions of saving, deleting, and organizing files.

ACTIVITY 3.3

There are many different kinds of vending machines in the world. Each offers a range of goods, requiring users to part with some of their money. Figure 3.6 shows photos of two different types of vending machines: one that provides soft drinks and the other that delivers a range of snacks. Both machines use an instructional mode of interaction. However, the way they do so is quite different.

What instructions must be issued to obtain a soda from the first machine and a bar of chocolate from the second? Why has it been necessary to design a more complex mode of interaction for the second vending machine? What problems can arise with this mode of interaction?

Comment

The first vending machine has been designed using simple instructions. There is a small number of drinks from which to choose, and each is represented by a large button displaying the label of each drink. The user simply has to press one button, and this will have the effect of delivering the selected drink. The second machine is more complex, offering a wider range of snacks. The trade-off for providing more options, however, is that the user can no longer




Figure 3.6 Two different types of vending machine

instruct the machine using a simple one-press action but is required to follow a more complex process involving (1) reading off the code (for example, C12) under the item chosen, then (2) keying this into the number pad adjacent to the displayed items, and finally (3) checking the price of the selected option and ensuring that the amount of money inserted is the same or greater (depending on whether the machine provides change). Problems that can arise from this type of interaction are the customer misreading the code and/or incorrectly keying the code, resulting in the machine not issuing the snack or providing the wrong item.

A better way of designing an interface for a large number of options of variable cost might be to continue to use direct mapping but use buttons that show miniature versions of the snacks placed in a large matrix (rather than showing actual versions). This would use the available space at the front of the vending machine more economically. The customer would need only to press the button of the object chosen and put in the correct amount of money. There is a lower chance of error resulting from pressing the wrong code or keys. The trade-off for the vending company, however, is that the machine is less flexible in terms of which snacks it can sell. If a new product line comes out, they will also need to replace part of the physical interface to the machine, which would be costly. ■

3.5.2 Conversing

This form of interaction is based on the idea of a person having a conversation with a system, where the system acts as a dialogue partner. In particular, the system is designed to respond in a way that another human being might when having a conversation. It differs from the activity of instructing insofar as it encompasses a two-way communication process, with the system acting like a partner rather than a machine that obeys orders. It has been most commonly used for applications where the user needs to find out specific kinds of information or wants to discuss issues. Examples include advisory systems, help facilities, chatbots, and robots.

The kinds of conversation that are currently supported range from simple voice-recognition, menu-driven systems, to more complex natural language–based systems that involve the system parsing and responding to queries typed in or spoken by the user. Examples of the former include banking, ticket booking, and train-time inquiries, where the user talks to the system in single-word phrases and numbers, that is, yes, no, three, and so on, in response to prompts from the system. Examples of the latter include help systems, where the user types in a specific query, such as “How do I change the margin widths?” to which the system responds by giving various answers. Advances in AI during the last few years have resulted in a significant improvement in speech recognition to the extent that many companies now routinely employ speech-based and chatbot-based interaction for their customer queries.

A main benefit of developing a conceptual model that uses a conversational style of interaction is that it allows people to interact with a system in a way that is familiar to them. For example, Apple’s speech system, Siri, lets you talk to it as if it were another person. You can ask it to do tasks for you, such as make a phone call, schedule a meeting, or send a message. You can also ask it indirect questions that it knows how to answer, such as “Do I need an umbrella today?” It will look up the weather for where you are and then answer with something like, “I don’t believe it’s raining” while also providing a weather forecast (see Figure 3.7).



Figure 3.7 Siri’s response to the question “Do I need an umbrella today?”

A problem that can arise from using a conversational-based interaction type is that certain kinds of tasks are transformed into cumbersome and one-sided interactions. This is especially true for automated phone-based systems that use auditory menus to advance the interaction. Users have to listen to a voice providing several options, then make a selection, and repeat through further layers of menus before accomplishing their goal, for example, reaching a real human or paying a bill. Here is the beginning of a dialogue between a user who wants to find out about car insurance and an insurance company's phone reception system:

<user dials an insurance company>

"Welcome to St. Paul's Insurance Company. Press 1 if you are a new customer; 2 if you are an existing customer."

<user presses 1>

"Thank you for calling St. Paul's Insurance Company. If you require house insurance, say 1; car insurance, say 2; travel insurance, say 3; health insurance, say 4; other, say 5."

<user says 2>

"You have reached the car insurance division. If you require information about fully comprehensive insurance, say 1; third-party insurance, say 2. ..."



**"If you'd like to press 1, press 3.
If you'd like to press 3, press 8.
If you'd like to press 8, press 5..."**

Source: © Glasbergen. Reproduced with permission of Glasbergen Cartoon Service

3.5.3 Manipulating

This form of interaction involves manipulating objects, and it capitalizes on users' knowledge of how they do so in the physical world. For example, digital objects can be manipulated by moving, selecting, opening, and closing. Extensions to these actions include zooming in and out, stretching, and shrinking—actions that are not possible with objects in the real world. Human actions can be imitated through the use of physical controllers (for example, the Wii) or gestures made in the air, such as the gesture control technology now used in some cars. Physical toys and robots have also been embedded with technology that enable them to act and react in ways depending on whether they are squeezed, touched, or moved. Tagged physical objects (such as balls, bricks, or blocks) that are manipulated in a physical world (for example, placed on a surface) can result in other physical and digital events occurring, such as a lever moving or a sound or animation being played.

A framework that has been highly influential (originating from the early days of HCI) in guiding the design of GUI applications is direct manipulation (Shneiderman, 1983). It proposes that digital objects be designed at the interface level so that they can be interacted with in ways that are analogous to how physical objects in the physical world are manipulated.

In so doing, direct manipulation interfaces are assumed to enable users to feel that they are directly controlling the digital objects represented by the computer. The three core principles are as follows:

- Continuous representation of the objects and actions of interest
- Rapid reversible incremental actions with immediate feedback about the object of interest
- Physical actions and button pressing instead of issuing commands with complex syntax

According to these principles, an object on the screen remains visible while a user performs physical actions on it, and any actions performed on it are immediately visible. For example, a user can move a file by dragging an icon that represents it from one part of the desktop to another. The benefits of direct manipulation include the following:

- Helping beginners learn basic functionality rapidly
- Enabling experienced users to work rapidly on a wide range of tasks
- Allowing infrequent users to remember how to carry out operations over time
- Preventing the need for error messages, except rarely
- Showing users immediately how their actions are furthering their goals
- Reducing users' experiences of anxiety
- Helping users gain confidence and mastery and feel in control

Many apps have been developed based on some form of direct manipulation, including word processors, video games, learning tools, and image editing tools. However, while direct manipulation interfaces provide a versatile mode of interaction, they do have their drawbacks. In particular, not all tasks can be described by objects, and not all actions can be undertaken directly. Some tasks are also better achieved through issuing commands. For example, consider how you edit a report using a word processor. Suppose that you had referenced work by Ben Shneiderman but had spelled his name as Schneiderman throughout. How would you correct this error using a direct manipulation interface? You would need to read the report and manually select the *c* in every *Schneiderman*, highlight it, and then delete it. This would be tedious, and it would be easy to miss one or two. By contrast, this operation is relatively effortless and also likely to be more accurate when using a command-based interaction. All you need to do is instruct the word processor to find every *Schneiderman* and replace it with *Shneiderman*. This can be done by selecting a menu option or using a combination of command keys and then typing the changes required into the dialog box that pops up.

3.5.4 Exploring

This mode of interaction involves users moving through virtual or physical environments. For example, users can explore aspects of a virtual 3D environment, such as the interior of a building. Physical environments can also be embedded with sensing technologies that, when they detect the presence of someone or certain body movements, respond by triggering certain digital or physical events. The basic idea is to enable people to explore and interact with an environment, be it physical or digital, by exploiting their knowledge of how they move and navigate through existing spaces.

Many 3D virtual environments have been built that comprise digital worlds designed for people to move between various spaces to learn (for example, virtual campuses) and fantasy worlds where people wander around different places to socialize (for instance, virtual parties)

or play video games (such as Fortnite). Many virtual landscapes depicting cities, parks, buildings, rooms, and datasets have also been built, both realistic and abstract, that enable users to fly over them and zoom in and out of different parts. Other virtual environments that have been built include worlds that are larger than life, enabling people to move around them, experiencing things that are normally impossible or invisible to the eye (see Figure 3.8a); highly realistic representations of architectural designs, allowing clients and customers to imagine how they will use and move through planned buildings and public spaces; and visualizations of complex datasets that scientists can virtually climb inside and experience (see Figure 3.8b).




Figure 3.8 (a) A CAVE that enables the user to stand near a huge insect, for example, a beetle, be swallowed, and end up in its abdomen; and (b) NCSA's CAVE being used by a scientist to move through 3D visualizations of the datasets

Source: (a) Used courtesy of Alexei Sharov (b) Used courtesy of Kalev Leetaru, National Center for Supercomputing Applications, University of Illinois.

3.5.5 Responding

This mode of interaction involves the system taking the initiative to alert, describe, or show the user something that it “thinks” is of interest or relevance to the context the user is presently in. It can do this through detecting the location and/or presence of someone in a vicinity (for instance, a nearby coffee bar where friends are meeting) and notifying them about it on their phone or watch. Smartphones and wearable devices are becoming increasingly proactive in initiating user interaction in this way, rather than waiting for the user to ask, command, explore, or manipulate. An example is a fitness tracker that notifies the user of a milestone they have reached for a given activity, for example, having walked 10,000 steps in a day. The fitness tracker does this automatically without any requests made by the user; the user responds by looking at the notification on their screen or listening to an audio announcement that is made. Another example is when the system automatically provides some funny or useful information for the user, based on what it has learned from their repeated behaviors when carrying out particular actions in a given context. For example, after taking a photo

of a friend's cute dog in the park, Google Lens will automatically pop up information that identifies the breed of the dog (see Figure 3.9).




Figure 3.9 Google Lens in action, providing pop-up information about Pembroke Welsh Corgi having recognized the image as one

Source: <https://lens.google.com>

For some people, this kind of system-initiated interaction—where additional information is provided which has not been requested—might get a bit tiresome or frustrating, especially if the system gets it wrong. The challenge is knowing when the user will find it useful and interesting and how much and what kind of contextual information to provide without overwhelming or annoying them. Also, it needs to know what to do when it gets it wrong. For example, if it thinks the dog is a teddy bear, will it apologize? Will the user be able to correct it and tell it what the photo actually is? Or will the system be given a second chance?

3.6 Paradigms, Visions, Theories, Models, and Frameworks

Other sources of conceptual inspiration and knowledge that are used to inform design and guide research are paradigms, visions, theories, models, and frameworks (Carroll, 2003). These vary in terms of their scale and specificity to a particular problem space. A *paradigm* refers to a general approach that has been adopted by a community of researchers and designers for carrying out their work in terms of shared assumptions, concepts, values, and practices. A *vision* is a future scenario that frames research and development in interaction design—often depicted in the form of a film or a narrative. A *theory* is a well-substantiated explanation of some aspect of a phenomenon; for example, the theory of information processing that explains how the mind, or some aspect of it, is assumed to work. A *model* is a simplification of some aspect of human-computer interaction intended to make it easier for designers to predict and evaluate alternative designs. A *framework* is a set of interrelated concepts and/or

a set of specific questions that are intended to inform a particular domain area (for example, collaborative learning), or an analytic method (for instance, ethnographic studies).

3.6.1 Paradigms

Following a particular paradigm means adopting a set of practices upon which a community has agreed. These include the following:

- The questions to be asked and how they should be framed
- The phenomena to be observed
- The way in which findings from studies are to be analyzed and interpreted (Kuhn, 1972)

In the 1980s, the prevailing paradigm in human-computer interaction was how to design user-centered applications for the desktop computer. Questions about what and how to design were framed in terms of specifying the requirements for a single user interacting with a screen-based interface. Task analytic and usability methods were developed based on an individual user's cognitive capabilities. Windows, Icons, Menus, and Pointers (WIMP) was used as a way of characterizing the core features of an interface for a single user. This was later superseded by the graphical user interface (GUI). Now many interfaces have touch screens that users tap, press and hold, pinch, swipe, slide, and stretch.

A big influence on the paradigm shift that took place in HCI in the 1990s was Mark Weiser's (1991) vision of ubiquitous technology. He proposed that computers would become part of the environment, embedded in a variety of everyday objects, devices, and displays. He envisioned a world of serenity, comfort, and awareness, where people were kept perpetually informed of what was happening around them, what was going to happen, and what had just happened. *Ubiquitous computing* devices would enter a person's center of attention when needed and move to the periphery of their attention when not, enabling the person to switch calmly and effortlessly between activities without having to figure out how to use a computer when performing their tasks. In essence, the technology would be unobtrusive and largely disappear into the background. People would be able to get on with their everyday and working lives, interacting with information and communicating and collaborating with others without being distracted or becoming frustrated with technology.

This vision was successful at influencing the computing community's thinking; inspiring them especially regarding what technologies to develop and problems to research (Abowd, 2012). Many HCI researchers began to think beyond the desktop and design mobile and pervasive technologies. An array of technologies was developed that could extend what people could do in their everyday and working lives, such as smart glasses, tablets, and smartphones.

The next big paradigm shift that took place in the 2000s was the emergence of *Big Data* and the *Internet of Things* (IoT). New and affordable sensor technologies enabled masses of data to be collected about people's health, well-being, and real-time changes happening in the environment (for example, air quality, traffic congestion, and business). Smart buildings were also built, where an assortment of sensors were embedded and experimented with in homes, hospitals, and other public buildings. Data science and machine-learning algorithms were developed to analyze the amassed data to draw new inferences about what actions to take on behalf of people to optimize and improve their lives. This included introducing variable speed limits on highways, notifying people via apps of dangerous pollution levels, crowds at an airport, and so on. In addition, it became the norm for sensed data to be used to automate mundane operations and actions—such as turning lights or faucets on and off or flushing toilets automatically—replacing conventional knobs, buttons, and other physical controls.

Video IBM's Internet of Things: <http://youtu.be/sfEbMV295Kk>.

3.6.2 Visions

Visions of the future, like Mark Weiser's vision of ubiquitous technology, provide a powerful driving force that can lead to a paradigm shift in terms of what research and development is carried out in companies and universities. A number of tech companies have produced videos about the future of technology and society, inviting audiences to imagine what life will be like in 10, 15, or 20 years' time. One of the earliest was Apple's 1987 Knowledge Navigator, which presented a scenario of a professor using a touchscreen tablet with a speech-based intelligent assistant reminding him of what he needed to do that day while answering the phone and helping him prepare his lectures. It was 25 years ahead of its time—set in 2011—the actual year that Apple launched its speech system, Siri. It was much viewed and discussed, inspiring widespread research into and development of future interfaces.

You can watch a video about the Apple Knowledge Navigator here: <http://youtu.be/HGYFEI6uLy0>.

A current vision that has become pervasive is AI. Both utopian and dystopian visions are being bandied about on how AI will make our lives easier on the one hand and how it will take our jobs away on the other. This time, it is not just computer scientists who are extolling the benefits or dangers of AI advances for society but also journalists, social commentators, policy-makers, and bloggers. AI is now replacing the user interface for an increasing number of applications where the user had to make choices, for example, smartphones learning your music preferences and home heating systems deciding when to turn the heating on and off and what temperature you prefer. One objective is to reduce the stress of people having to make decisions; another is to improve upon what they would choose. For example, in the future instead of having to agonize over which clothes to buy, or vacation to select, a personal assistant will be able to choose on your behalf. Another example depicts what a driverless car will be like in a few years, where the focus is not so much on current concerns with safety and convenience but more on improving comfort and life quality in terms of the ultimate personalized passenger experience (for example, see VW's video). More and more everyday tasks will be transformed through AI learning what choices are best in a given situation.

VW's vision of its future car can be seen in this video: <https://youtu.be/AyihacflLto>.

While there are many benefits of letting machines make decisions for us, we may feel a loss of control. Moreover, we may not understand why the AI system chose to drive the car along a particular route or why our voice-assisted home robot keeps ordering too much milk. There are increasing expectations that AI researchers find ways of explaining the rationale behind the decisions that AI systems make on the user's behalf. This need is often referred to as *transparency* and *accountability*—which we discuss further in Chapter 10. It is an area that is of central concern to interaction design researchers, who have started conducting user studies on transparency and developing explanations that are meaningful and reassuring to the user (e.g., Radar et al., 2018).

Another challenge is to develop new kinds of interfaces and conceptual models that can support the synergy of humans and AI systems, which will amplify and extend what they can do currently. This could include novel ways of enhancing group collaboration, creative problem-solving, forward planning, policy-making, and other areas that can become intractable, complex, and messy, such as divorce settlements.

Science fiction has also become a source of inspiration in interaction design. By this, we mean in movies, writing, plays, and games that envision what role technology may play in the future. Dan Russell and Svetlana Yarosh (2018) discuss the pros and cons of using different kinds of science fiction for inspiration in HCI design, arguing that they can provide a good grounding for debate but are often not a great source of accurate predictions of future technologies. They point out how, although the visions can be impressively futuristic, their embellishments and what they actually look like are often limited by the author's ability to extend and build upon the ideas and the cultural expectations associated with the current era. For example, the holodeck portrayed in the *Star Trek* TV series had 3D-bubble indicator lights and push-button designs on its bridge with the sound of a teletype in the background. This is the case to such an extent that Russell and Yarosh even argue that the priorities and concerns of the author's time and their cultural upbringing can bias the science fiction toward telling narratives from the perspective of the present, rather than providing new insights and paving the way to future designs.

The different kinds of future visions provide concrete scenarios of how society can use the next generation of imagined technologies to make their lives more comfortable, safe, informative, and efficient. Furthermore, they also raise many questions concerning privacy, trust, and what we want as a society. They provide much food for thought for researchers, policy-makers, and developers, challenging them to consider both positive and negative implications.

Many new challenges, themes, and questions have been articulated through such visions (see, for example, Rogers, 2006; Harper et al., 2008; Abowd, 2012), including the following:

- How to enable people to access and interact with information in their work, social, and everyday lives using an assortment of technologies
- How to design user experiences for people using interfaces that are part of the environment but where there are no obvious controlling devices
- How and in what form to provide contextually relevant information to people at appropriate times and places to support them while on the move
- How to ensure that information that is passed around via interconnected displays, devices, and objects is secure and trustworthy

3.6.3 Theories

Over the past 30 years, numerous theories have been imported into human-computer interaction, providing a means of analyzing and predicting the performance of users carrying out tasks for specific types of computer interfaces and systems (Rogers, 2012). These have been primarily cognitive, social, affective, and organizational in origin. For example, cognitive theories about human memory were used in the 1980s to determine the best ways of representing operations, given people's memory limitations. One of the main benefits of applying such theories in interaction design is to help identify factors (cognitive, social, and affective) relevant to the design and evaluation of interactive products. Some of the most influential theories in HCI, including distributed cognition, will be covered in the next chapter.

3.6.4 Models

We discussed earlier why a conceptual model is important and how to generate one when designing a new product. The term *model* has also been used more generally in interaction design to describe, in a simplified way, some aspect of human behavior or human-computer interaction. Typically, it depicts how the core features and processes underlying a phenomenon are structured and related to one another. It is usually abstracted from a theory coming from a contributing discipline, like psychology. For example, Don Norman (1988) developed a number of models of user interaction based on theories of cognitive processing, arising out of cognitive science, which were intended to explain the way users interacted with interactive technologies. These include the seven stages of action model that describes how users move from their plans to executing physical actions that they need to perform to achieve them to evaluating the outcome of their actions with respect to their goals. More recent models developed in interaction design are user models, which predict what information users want in their interactions and models that characterize core components of the user experience, such as Marc Hassenzahl's (2010) model of experience design.

3.6.5 Frameworks

Numerous frameworks have been introduced in interaction design to help designers constrain and scope the user experience for which they are designing. In contrast to a model, a framework offers advice to designers as to what to design or look for. This can come in a variety of forms, including steps, questions, concepts, challenges, principles, tactics, and dimensions. Frameworks, like models, have traditionally been based on theories of human behavior, but they are increasingly being developed from the experiences of actual design practice and the findings arising from user studies.

Many frameworks have been published in the HCI/interaction design literature, covering different aspects of the user experience and a diversity of application areas. For example, there are frameworks for helping designers think about how to conceptualize learning, working, socializing, fun, emotion, and so on, and others that focus on how to design particular kinds of technologies to evoke certain responses, for example, persuasive technologies (see Chapter 6, "Emotional Interaction"). There are others that have been specifically developed to help researchers analyze the qualitative data they collect in a user study, such as Distributed Cognition (Rogers, 2012). One framework, called DiCoT (Furniss and Blandford, 2006), was developed to analyze qualitative data at the system level, allowing researchers to understand how technologies are used by teams of people in work or home settings. (Chapter 9, "Data Analysis," describes DiCoT in more detail.)

A classic example of a conceptual framework that has been highly influential in HCI is Don Norman's (1988) explanation of the relationship between the design of a conceptual model and a user's understanding of it. The framework comprises three interacting components: the designer, the user, and the system. Behind each of these are the following:

Designer's Model The model the designer has of how the system should work

System Image How the system actually works, which is portrayed to the user through the interface, manuals, help facilities, and so on

User's Model How the user understands how the system works

The framework makes explicit the relationship between how a system should function, how it is presented to users, and how it is understood by them. In an ideal world, users should be able to carry out activities in the way intended by the designer by interacting with the system image that makes it obvious what to do. If the system image does not make the designer's model clear to the users, it is likely that they will end up with an incorrect understanding of the system, which in turn will increase the likelihood of their using the system ineffectively and making errors. This has been found to happen often in the real world. By drawing attention to this potential discrepancy, designers can be made aware of the importance of trying to bridge the gap more effectively.

To summarize, paradigms, visions, theories, models, and frameworks are not mutually exclusive, but rather they overlap in their way of conceptualizing the problem and design space, varying in their level of rigor, abstraction, and purpose. *Paradigms* are overarching approaches that comprise a set of accepted practices and framing of questions and phenomena to observe; *visions* are scenarios of the future that set up challenges, inspirations, and questions for interaction design research and technology development; *theories* tend to be comprehensive, explaining human-computer interactions; *models* are an abstraction that simplify some aspect of human-computer interaction, providing a basis for designing and evaluating systems; and *frameworks* provide a set of core concepts, questions, or principles to consider when designing for a user experience or analyzing data from a user study.

DILEMMA

Who Is in Control?

A recurrent theme in interaction design, especially in the current era of AI-based systems, is who should be in control at the interface level. The different interaction types vary in terms of how much control a user has and how much the computer has. While users are primarily in control for instructing direct manipulation interfaces, they are less so in responding type interfaces, such as sensor-based and context-aware environments where the system takes the initiative to act. User-controlled interaction is based on the premise that people enjoy mastery and being in control. It assumes that people like to know what is going on, be involved in the action, and have a sense of power over the computer.

In contrast, autonomous and context-aware control assumes that having the environment monitor, recognize, and detect deviations in a person's behavior can enable timely, helpful, and even critical information to be provided when considered appropriate (Abowd and Mynatt, 2000). For example, elderly people's movements can be detected in the home and emergency or care services alerted if something untoward happens to them that might otherwise go unnoticed, for instance, if they fall over and are unable to sound the alarm. But what happens if a person chooses to take a rest in an unexpected area (on the carpet), which the system detects as a fall? Will the emergency services be called out unnecessarily and cause care givers undue worry? Will the person who triggered the alarm be mortified at triggering a false alarm? And how will it affect their sense of privacy, knowing that their every move is constantly being monitored?

Another concern is what happens when the locus of control switches between user and system. For example, consider who is in control when using a GPS for vehicle navigation. At the beginning, the driver is very much in control, issuing instructions to the system as to where to go and what to include, such as highways, gas stations, and traffic alerts. However, once on the road, the system takes over and is in control. People often find themselves slavishly following what the GPS tells them to do, even though common sense suggests otherwise.



Source: Cluff / Cartoon Stock

To what extent do you need to be in control in your everyday and working life? Are you happy to let technology monitor and decide what you need or might be interested in knowing, or do you prefer to tell it what you want to do? What do you think of autonomous cars that drive for you? While it might be safer and more fuel-efficient, will it take the pleasure out of driving?

A tongue-in-cheek video made by Superflux, called *Uninvited Guests*, shows who is very much in control when a man is given lots of smart gadgets by his children for his birthday to help him live more healthily: <https://vimeo.com/128873380>. ■

In-Depth Activity

The aim of this in-depth activity is for you to think about the appropriateness of different kinds of conceptual models that have been designed for similar physical and digital information artifacts.

Compare the following:

- A paperback book and an ebook
- A paper-based map and a smartphone map app

What are the main concepts and metaphors that have been used for each? (Think about the way time is conceptualized for each of them.) How do they differ? What aspects of the paper-based artifact have informed the digital app? What is the new functionality? Are any aspects of the conceptual model confusing? What are the pros and cons?

Summary

This chapter explained the importance of conceptualizing the problem and design spaces before trying to build anything. It stressed throughout the need to be explicit about the claims and assumptions behind design decisions that are suggested. It described an approach to formulating a conceptual model and explained the evolution of interface metaphors that have been designed as part of the conceptual model. Finally, it considered other ways of conceptualizing interaction in terms of interaction types, paradigms, visions, theories, models, and frameworks.

Key Points

- A fundamental aspect of interaction design is to develop a conceptual model.
- A conceptual model is a high-level description of a product in terms of what users can do with it and the concepts they need to understand how to interact with it.
- Conceptualizing the problem space in this way helps designers specify what it is they are doing, why, and how it will support users in the way intended.
- Decisions about conceptual design should be made before commencing physical design (such as choosing menus, icons, dialog boxes).
- Interface metaphors are commonly used as part of a conceptual model.
- Interaction types (for example, conversing or instructing) provide a way of thinking about how best to support the activities users will be doing when using a product or service.
- Paradigms, visions, theories, models, and frameworks provide different ways of framing and informing design and research.

Further Reading

Here we recommend a few seminal readings on interaction design and the user experience (in alphabetical order).

DOURISH, P. (2001) *Where the Action Is*. MIT Press. This book presents an approach for thinking about the design of user interfaces and user experiences based on the notion of embodied interaction. The idea of embodied interaction reflects a number of trends that have emerged in HCI, offering new sorts of metaphors.

JOHNSON, J. and HENDERSON, A. (2012) *Conceptual Models: Core to Good Design*. Morgan and Claypool Publishers. This short book, in the form of a lecture, provides a comprehensive overview of what is a conceptual model using detailed examples. It outlines how to construct one and why it is necessary to do so. It is cogently argued and shows how and where this design activity can be integrated into interaction design.

JU, W. (2015) *The Design of Implicit Interactions*. Morgan and Claypool Publishers. This short book, in the form of a lecture, provides a new theoretical framework to help design smart, automatic, and interactive devices by examining the small interactions that we engage in our everyday lives, often without any explicit communication. It puts forward the idea of implicit interaction as a central concept for designing future interfaces.



INTERVIEW with Albrecht Schmidt

Albrecht Schmidt is professor of human-centered ubiquitous media in the computer science department of the Ludwig-Maximilians-Universität München in Germany. He studied computer science in Ulm and Manchester and received a PhD from Lancaster University, United Kingdom, in 2003. He held several prior academic positions at different universities, including Stuttgart, Cambridge, Duisburg-Essen, and Bonn. He also worked as a researcher at the Fraunhofer Institute for Intelligent Analysis and Information Systems (IAIS) and at Microsoft Research in Cambridge. In his research, he investigates the inherent complexity of human-computer interaction in ubiquitous computing environments, particularly in view of increasing computer intelligence and system autonomy. Albrecht has actively contributed to the scientific discourse in human-computer interaction through the development, deployment, and study of functional prototypes of interactive systems and interface technologies in different real world domains. Most recently, he focuses on how information technology can provide cognitive and perceptual support to amplify the human mind.

How do you think future visions inspire research in interaction design? Can you give an example from your own work?

Envisioning the future is key to research in human-computer interaction. In contrast to traditional fields that discover phenomena (such as physics or sociology), research in interaction design is constructive and creates new things that potentially change our world. Research in interaction design also analyzes the world and aims to understand phenomena but mainly as a means to inspire and guide innovations. A major aspect of research is then to create concrete designs, build concepts and prototypes, and evaluate them.

Future visions are an excellent way to describe the big picture of a future where we still have to invent and implement the details. A vision enables us to communicate the overall goal for which we are aiming. In formulating the vision, we have to contextualize our ideas, link them to practices in our lives, and describe the anticipated impact on individuals and society. A prerequisite for formulating a coherent future vision is a good understanding of the problems that we want to address. If formulated well, the vision shows a clear direction, but it leaves room for researchers in the community to make their own interpretation. A well-formulated future vision also leaves room for individuals to align their research efforts with the goal or to criticize it fundamentally through their research.

We have proposed the vision of *amplifying human perception and cognition through digital technologies* (see Schmidt, 2017a; 2017b). This vision emerged from our various concrete research prototypes over the last 10 years. We realized that many of the prototypes and technologies we developed were pointed in a similar direction: enabling superhuman abilities through devices and applications. At the same time, we demonstrate why amplifying human abilities is a timely endeavor, in particular given the current advances in artificial intelligence, in sensing technologies, as well as in personal display devices. For our group and for colleagues with whom we work, this vision has become a means for inspiring new ideas, for investigating relevant areas for potential innovation systematically, and for assessing ideas early on.

Why do metaphors persist in HCI?

Good metaphors allow people to transfer their understanding and skills from another domain in the real world to interaction with computers and data. Good metaphors are abstract enough to persist over time, but concrete enough to simplify the use of computers. Early metaphors included computers as an advanced typewriter and computers as intelligent assistants. Such metaphors help in the design process to create understandable interaction concepts and user interfaces. A designer can take their idea for a user interface or an interaction concept and evaluate it in the light of the metaphor. They can assess if the interaction is understandable for people familiar with the concept or technologies on which the metaphor is based. Metaphors often suggest interaction styles and hence can help to create interfaces that are more consistent and interaction designs

that can be used without explanation using intuition (which in this case is the implicit understanding of the metaphor).

Metaphors for which the underlying concept has disappeared from everyday usage may still persist. In many cases, users will not know the original concept from their own experience (for example, a typewriter), but have grown up with technologies using the metaphor. For a metaphor to persist, it must remain conducive and helpful over time for new users as well as for experienced ones.

What do you think of the rise of AI and automation? Do you think there is a role for HCI, and if so, what?

Advancements in AI and in automation are exciting. They have the potential to empower humans to do things, to think things, and to experience things we cannot even imagine right now. However, the key to unlocking this potential is to create efficient ways for interacting with artificial intelligence. Meaningful automation and intelligent systems always have boundaries and intersections with human action. For example, an autonomous car will transport a human, a drone will deliver a parcel to a person, an automated kitchen will prepare a meal for a family, and large-scale data analytics in companies will lead to better services for their customers. With intelligent systems and smart services taking a more active role through artificial intelligence, the way in which interaction and interfaces are designed becomes even more crucial. Creating a positive user experience in the presence of artificial intelligence is a challenge where new visions and metaphors are required.

One concept that we suggested is the notion of intervention user interfaces (Schmidt and Herrmann, 2017). The basic

expectation is that in the future many intelligent systems in our environment will work just fine, without any human interaction. However, to stay in control and to tailor the system to current and unforeseeable needs, as well as to customize the user experience, human interventions should be easily possible. Designing interaction concepts for interventions and user interfaces that empower humans to make the most of a system driven by artificial intelligence is a huge challenge and includes many basic research questions. Getting the interaction with AI right, basically finding ways for humans to harness the power of AI for what they want to do, is as important as developing the underlying algorithms. One without the other is of very limited value.

What do you see are the challenges ahead for HCI at scale?

There are many challenges ahead in the context of autonomous systems and artificial intelligence as outlined earlier. Closely related to this is human data interaction beyond interactive visualization. How can we empower humans to work with big and unstructured data? Here is a concrete example: I had a discussion with medical professionals this morning. For a specific cancer type, there are several thousand publications readily available. Many of them may have similar results, and others may have conflicting ones. Reading all of the publications in their current form is an intractable problem for a human reader, as it would take too long and would overload a person's working memory. The simple question that resulted from

the discussion is: what would a system and interface look like that uses AI to preprocess 10,000 papers, allows interactive presentation of relevant content, and enables humans to make sense of the state of the art and come up with its own hypotheses? Preferably, the interface would support the person to do this in a few hours, rather than in their entire lifetime.

Another challenge at the societal scale is to understand the long-term impact of interactive systems that we create. So far, this was very much trial and error. Providing unlimited and easy-to-use mass communication to individuals without journalistic training has changed how we read news. Personal communication devices and instant messaging have altered communication patterns in families and classrooms. Working in the office using a computer in order to create texts is reducing our physical movements. The way that we design interactive systems, the things we make easy or hard to use, and the modalities that we choose in our interaction design have inevitably resulted in long-term impacts on people. With the current methods and tools in HCI, we are well equipped to do a great job in developing easy-to-use systems with an amazing short-term user experience for the individual. However, looking at upcoming major innovations in mobility and healthcare technologies, the interfaces we design may have many more consequences. One major challenge at scale is to design for a longer-term user experience (months to years) on a societal scale. Here we still first have to research and invent methods and tools.

Chapter 4

COGNITIVE ASPECTS

4.1 Introduction

4.2 What Is Cognition?

4.3 Cognitive Frameworks

Objectives

The main goals of the chapter are to accomplish the following:

- Explain what cognition is and why it is important for interaction design.
- Discuss what attention is and its effects on our ability to multitask.
- Describe how memory can be enhanced through technology aids.
- Show the difference between various cognitive frameworks that have been applied to HCI.
- Explain what are mental models.
- Enable you to elicit a mental model and understand what it means.

4.1 Introduction

Imagine that it is getting late in the evening and you are sitting in front of your laptop. You have a report to complete by tomorrow morning, but you are not getting very far with it. You begin to panic and start biting your nails. You see two text messages flash up on your smartphone. You instantly abandon your report and cradle your smartphone to read them. One is from your mother, and the other is from your friend asking if you want to go out for a drink. You reply immediately to both of them. Before you know it, you're back on Facebook to see whether any of your friends have posted anything about the party you wanted to go to but had to say no. Your phone rings, and you see that it's your dad calling. You answer it, and he asks if you have been watching the football game. You say that you are too busy working toward a deadline, and he tells you that your team has just scored. You chat with him for a while and then say you have to get back to work. You realize 30 minutes have passed, and you return your attention to your report. But before you realize it, you click your favorite sports site to check the latest score of the football game and discover that your team has just scored again. Your phone starts buzzing. Two new WhatsApp messages are waiting for you. And on it goes. You glance at the time on your laptop. It is midnight. You really are in a panic now and finally close everything down except your word processor.

In the past 10–15 years, it has become increasingly common for people to be switching their attention constantly among multiple tasks. The study of human cognition can help us understand the impact of multitasking on human behavior. It can also provide insights into other types of digital behaviors, such as decision-making, searching, and designing when using computer technologies by examining human abilities and limitations.

This chapter covers these aspects by examining the cognitive aspects of interaction design. It considers what humans are good and bad at, and it shows how this knowledge can inform the design of technologies that both extend human capabilities and compensate for human weaknesses. Finally, relevant cognitive theories, which have been applied in HCI to inform technology design, are described. (Other ways of conceptualizing human behavior that focus on the social and emotional aspects of interaction are presented in the following two chapters.)

4.2 What Is Cognition?

There are many different kinds of cognition, such as thinking, remembering, learning, day-dreaming, decision-making, seeing, reading, writing, and talking. A well-known way of distinguishing between different modes of cognition is in terms of whether it is experiential or reflective (Norman, 1993). *Experiential cognition* is a state of mind where people perceive, act, and react to events around them intuitively and effortlessly. It requires reaching a certain level of expertise and engagement. Examples include driving a car, reading a book, having a conversation, and watching a video. In contrast, *reflective cognition* involves mental effort, attention, judgment, and decision-making, which can lead to new ideas and creativity. Examples include designing, learning, and writing a report. Both modes are essential for everyday life. Another popular way of describing cognition is in terms of fast and slow thinking (Kahneman, 2011). Fast thinking is similar to Don Norman’s experiential mode insofar as it is instinctive, reflexive, and effortless, and it has no sense of voluntary control. Slow thinking, as the name suggests, takes more time and is considered to be more logical and demanding, and it requires greater concentration. The difference between the two modes is easy to see when asking someone to give answers to the following two arithmetic equations:

$$2 + 2 =$$

$$21 \times 19 =$$

The former can be done by most adults in a split second without thinking, while the latter requires much mental effort; many people need to externalize the task to be able to complete it by writing it down on paper and using the long multiplication method. Nowadays, many people simply resort to fast thinking by typing the numbers to be added or multiplied into a calculator app on a smartphone or computer.

Other ways of describing cognition are in terms of the context in which it takes place, the tools that are employed, the artifacts and interfaces that are used, and the people involved (Rogers, 2012). Depending on when, where, and how it happens, cognition can be distributed, situated, extended, and embodied. Cognition has also been described in terms of specific kinds of processes (Eysenck and Brysbaert, 2018). These include the following:

- Attention
- Perception

- Memory
- Learning
- Reading, speaking, and listening
- Problem-solving, planning, reasoning, and decision-making

It is important to note that many of these cognitive processes are interdependent: several may be involved for a given activity. It is rare for one to occur in isolation. For example, when reading a book one has to attend to the text, perceive and recognize the letters and words, and try to make sense of the sentences that have been written.

In the following sections we describe the main kinds of cognitive processes in more detail, followed by a summary box highlighting the core design implications for each. The most relevant for interaction design are attention and memory, which we describe in the greatest detail.

4.2.1 Attention

Attention is central to everyday life. It enables us to cross the road without being hit by a car or bicycle, notice when someone is calling our name, and be able to text while at the same time watching TV. It involves selecting things on which to concentrate, at a point in time, from the range of possibilities available, allowing us to focus on information that is relevant to what we are doing. The extent to which this process is easy or difficult depends on (1) whether someone has clear goals and (2) whether the information they need is salient in the environment.

4.2.1.1 Clear Goals

If someone knows exactly what they want to find out, they try to match this with the information that is available. For example, when someone has just landed at an airport after a long flight, which did not have Wi-Fi onboard, and they want to find out who won the World Cup, they might scan the headlines on their smartphone or look at breaking news on a public TV display inside the airport. When someone is not sure exactly what they are looking for, they may browse through information, allowing it to guide their attention to interesting or salient items. For example, when going to a restaurant, someone may have the general goal of eating a meal but only a vague idea of what they want to eat. They peruse the menu to find things that whet their appetite, letting their attention be drawn to the imaginative descriptions of various dishes. After scanning through the possibilities and imagining what each dish might be like, as well as considering other factors, such as cost, who they are with, what are the specials, what the waiter recommends, and whether they want a two- or three-course meal, and so on), they then decide.

4.2.1.2 Information Presentation

The way information is displayed can also greatly influence how easy or difficult it is to comprehend appropriate pieces of information. Look at Figure 4.1, and try the activity (based on Tullis, 1997). Here, the information-searching tasks are precise, requiring specific answers.

| South Carolina | | | | | |
|----------------|-----------------|-----------|----------|--------|--------|
| City | Motel/Hotel | Area code | Phone | Rates | |
| | | | | Single | Double |
| Charleston | Best Western | 803 | 747-0961 | \$126 | \$130 |
| Charleston | Days Inn | 803 | 881-1000 | \$118 | \$124 |
| Charleston | Holiday Inn N | 803 | 744-1621 | \$136 | \$146 |
| Charleston | Holiday Inn SW | 803 | 556-7100 | \$133 | \$147 |
| Charleston | Howard Johnsons | 803 | 524-4148 | \$131 | \$136 |
| Charleston | Ramada Inn | 803 | 774-8281 | \$133 | \$140 |
| Charleston | Sheraton Inn | 803 | 744-2401 | \$134 | \$142 |
| | | | | | |
| Columbia | Best Western | 803 | 796-9400 | \$129 | \$134 |
| Columbia | Carolina Inn | 803 | 799-8200 | \$142 | \$148 |
| Columbia | Days Inn | 803 | 736-0000 | \$123 | \$127 |
| Columbia | Holiday Inn NW | 803 | 794-9440 | \$132 | \$139 |
| Columbia | Howard Johnsons | 803 | 772-7200 | \$125 | \$127 |
| Columbia | Quality Inn | 803 | 772-0270 | \$134 | \$141 |
| Columbia | Ramada Inn | 803 | 796-2700 | \$136 | \$144 |
| Columbia | Vagabond Inn | 803 | 796-6240 | \$127 | \$130 |

(a)

| Pennsylvania | |
|--|----------------------------------|
| Bedford Motel/Hotel: Crinaline Courts | (814) 623-9511 S: \$118 D: \$120 |
| Bedford Motel/Hotel: Holiday Inn | (814) 623-9006 S: \$129 D: \$136 |
| Bedford Motel/Hotel: Midway | (814) 623-8107 S: \$121 D: \$126 |
| Bedford Motel/Hotel: Penn Manor | (814) 623-8177 S: \$119 D: \$125 |
| Bedford Motel/Hotel: Quality Inn | (814) 623-5189 S: \$123 D: \$128 |
| Bedford Motel/Hotel: Terrace | (814) 623-5111 S: \$122 D: \$124 |
| Bradley Motel/Hotel: De Soto | (814) 362-3567 S: \$120 D: \$124 |
| Bradley Motel/Hotel: Holiday House | (814) 362-4511 S: \$122 D: \$125 |
| Bradley Motel/Hotel: Holiday Inn | (814) 362-4501 S: \$132 D: \$140 |
| Breezewood Motel/Hotel: Best Western Plaza | (814) 735-4352 S: \$120 D: \$127 |
| Breezewood Motel/Hotel: Motel 70 | (814) 735-4385 S: \$116 D: \$118 |

(b)

Figure 4.1 Two different ways of structuring the same information at the interface level. One makes it much easier to find information than the other.

Source: Used courtesy of Dr. Tom Tullis

ACTIVITY 4.1

Look at the top screen of Figure 4.1 and (1) find the price for a double room at the Quality Inn in Columbia, South Carolina, and (2) find the phone number of the Days Inn in Charleston, South Carolina. Then look at the bottom screen in Figure 4.1 and (1) find the price of a double room at the Holiday Inn in Bradley, Pennsylvania, and (2) find the phone number of the Quality Inn in Bedford, Pennsylvania. Which took longer to do?

In an early study, Tullis found that the two screens produced quite different results: It took an average of 3.2 seconds to search the top screen, while it took an average of 5.5 seconds to find the same kind of information in the bottom screen. Why is this so, considering that both displays have the same density of information relative to the background?

Comment

The primary reason for the disparity is the way that the characters are grouped in the display. In the top screen, they are grouped into vertical categories of information (that is, place, type of accommodation, phone number, and rates), and this screen has space in between the columns of information. In the bottom screen, the information is bunched together, making it much more difficult to search. ■

4.2.1.3 *Multitasking and Attention*

As mentioned in the introduction to this chapter, many people now multitask, frequently switching their attention among different tasks. For example, in a study of teenage multitasking, it was found that the majority of teenagers were found to multitask most or some of the time while listening to music, watching TV, using a computer, or reading (Rideout et al., 2010). It is probably even higher now, considering their use of smartphones while walking, talking, and studying. While attending a presentation at a conference, we witnessed someone deftly switch between four ongoing instant message chats (one at the conference, one at school, one with friends, and one at her part-time job), read, answer, delete, and place all new messages in various folders of her two email accounts, and check and scan her Facebook and her Twitter feeds, all while appearing to listen to the talk, take some notes, conduct a search on the speaker's background, and open up their publications. When she had a spare moment, she played the game *Patience*. It was exhausting just watching her for a few minutes. It was as if she were capable of living in multiple worlds simultaneously while not letting a moment go to waste. But how much did she really take in of the presentation?

Is it possible to perform multiple tasks without one or more of them being detrimentally affected? There has been much research on the effects of multitasking on memory and attention (Burgess, 2015). The general finding is that it depends on the nature of the tasks and how much attention each demands. For example, listening to gentle music while working can help people tune out background noise, such as traffic or other people talking, and help them concentrate on what they are doing. However, if the music is loud, like heavy metal, it can be distracting.

Individual differences have also been found. For example, the results of a series of experiments comparing heavy with light multitaskers showed that heavy media multitaskers (such as the person described above) were more prone to being distracted by the multiple streams of media they are viewing than those who infrequently multitask. The latter were found to be better at allocating their attention when faced with competing distractions

(Ophir et al., 2009). This suggests that people who are heavy multitaskers are likely to be those who are easily distracted and find it difficult to filter out irrelevant information. However, a more recent study by Danielle Lotteridge et al. (2015) found that it may be more complex. They found that while heavy multitaskers are easily distracted, they can also put this to good use if the distracting sources are relevant to the task in hand. Lotteridge et al. conducted a study that involved writing an essay under two conditions—either with relevant or irrelevant information. They found that if the information sources are relevant, they don't affect the essay writing. The condition where irrelevant information was provided was found to negatively impact task performance. In summary, they found that multitasking can be both good and bad—it depends on what you are distracted by and how relevant it is to the task at hand.

The reason why multitasking is thought to be detrimental for human performance is that it overloads people's capacity to focus. Having switched attention from what someone is working on to another piece of information requires additional effort to get back into the other task and to remember where they were in the ongoing activity. Thus, the time to complete a task can be significantly increased. A study of completion rates of coursework found that students who were involved in instant messaging took up to 50 percent longer to read a passage from a textbook compared with those who did not instant message while reading (Bowman et al., 2010). Multitasking can also result in people losing their train of thought, making errors, and needing to start over.

Nevertheless, many people are expected to multitask in the workplace nowadays, such as in hospitals, as a result of the introduction of ever more technology (for example, multiple screens in an operating room). The technology is often introduced to provide new kinds of real-time and changing information. However, this usually requires the constant attention of clinicians to check whether any of the data is unusual or unexpected. Managing the ever-increasing information load requires professionals, like clinicians, to develop new attention and scanning strategies, looking out for anomalies in data visualizations and listening for audio alarms alerting them to potential dangers. Interaction designers have tried to make this easier by including the use of ambient displays that come on when something needs attention—flashing arrows to direct attention to a particular type of data or history logs of recent actions that can be quickly examined to refresh one's memory of what has just happened on a given screen. However, how well clinicians manage to switch and divide their attention among different tasks in tech-rich environments has barely been researched (Douglas et al., 2017).



"This project calls for real concentration.
Are you still able to monotask?"

Source: Chris Wildt / Cartoon Stock

DILEMMA

Is It OK to Use a Phone While Driving?

There has been considerable debate about whether drivers should be able to talk or text on their phones at the same time as driving (see Figure 4.2). People talk on their phones while walking, so why not be able to do the same thing when driving? The main reasons are that driving is more demanding, drivers are more prone to being distracted, and there is a greater chance of causing accidents (however, it is also the case that some people, when using their phones, walk out into a road without looking to see whether any cars are coming).



Figure 4.2 How distracting is it to be texting the phone while driving?

Source: Tetra Images / Alamy Stock Photo

A meta-review of research that has investigated mobile phone use in cars has found that drivers' reaction times are longer to external events when engaged in phone conversations (Caird et al., 2018). Drivers who use phones have also been found to be much poorer at staying in their lane and maintaining the correct speed (Stavrinos et al., 2013). The reason for this is that drivers on a phone rely more on their expectations about what is likely to happen next and, as a result, respond much more slowly to unexpected events, such as the car in front of them stopping (Briggs et al., 2018). Moreover, phone conversations cause the driver visually to imagine what is being talked about. The driver may also imagine the facial expression of the person to whom they are speaking. The visual imagery involved competes for the processing resources also needed to enable the driver to notice and react to what is in front of them on the road. The idea that using a hands-free device is safer than actually holding the phone to carry out a conversation is false, as the same type of cognitive processing takes place both ways.

(Continued)

It has also been found that drivers who engage in conversation with their passengers experienced similar negative effects. However, there is a difference between having a conversation with a passenger sitting next to the driver and one with a person located remotely. The driver and front-seat passenger can observe jointly what is happening in front of them on the road and will moderate or cease their conversation in order to switch their full attention to a potential or actual hazard. Someone on the other end of a phone, however, is not privy to what the driver is seeing and will carry on the conversation. They might have just asked “Where did you leave the spare set of keys?” and caused the driver mentally to search for them in their home, making it more difficult for them to switch their full attention back to what is happening on the road.

Because of these hazardous problems, many countries have banned the use of phones while driving. To help drivers resist the temptation to answer a phone that rings or glance at an incoming notification that pings, smartphone device manufacturers have been asked by some governments to introduce a driver mode akin to the airplane mode that could automatically lock down a smartphone, preventing access to apps, while disabling the phone’s keyboard when it detects a person who is driving. For example, the iPhone now has implemented this option. ■

In several contexts, therefore, multitasking can be detrimental to performance, such as texting or speaking on the phone when driving. The cost of switching attention varies from person to person and which information resources are being switched between. When developing new technology to provide more information for people in their work settings, it is important to consider how best to support them so that they can easily switch their attention back and forth

Design Implications

Attention

- Consider context. Make information salient when it requires attention at a given stage of a task.
- Use techniques to achieve this when designing visual interfaces, such as animated graphics, color, underlining, ordering of items, sequencing of different information, and spacing of items.
- Avoid cluttering visual interfaces with too much information. This applies especially to the use of color and graphics: It is tempting to use lots of these attributes, which results in a mishmash of media that is distracting and annoying rather than helping the user attend to relevant information.
- Consider designing different ways of supporting effective switching and returning to a particular interface. This could be done subtly, such as the use of pulsing lights gradually getting brighter, or abruptly, such as the use of alerting sounds or voice. How much competing visual information or ambient sound is present also needs to be considered. ■

among the multiple displays or devices and be able to return readily to what they were doing after an interruption (for instance, the phone ringing or people entering their space to ask questions).

4.2.2 Perception

Perception refers to how information is acquired from the environment via the five sense organs (vision, hearing, taste, smell, and touch) and transformed into experiences of objects, events, sounds, and tastes (Roth, 1986). In addition, we have the additional sense of kinaesthesia, which relates to the awareness of the position and movement of the parts of the body through internal sensory organs (known as *proprioceptors*) located in the muscles and joints. Perception is complex, involving other cognitive processes such as memory, attention, and language. Vision is the most dominant sense for sighted individuals, followed by hearing and touch. With respect to interaction design, it is important to present information in a way that can be readily perceived in the manner it was intended.

As was demonstrated in Activity 4.1, grouping items together and leaving spaces between them can aid attention because it breaks up the information. Having chunks of information makes it easier to scan, rather than one long list of text that is all the same. In addition, many designers recommend using blank space (more commonly known as *white space*) when grouping objects, as it helps users to perceive and locate items more easily and quickly (Malamed, 2009). In a study comparing web pages displaying the same amount of information but structured using different graphical methods (see Figure 4.3), it was found that people took less time to locate items from information that was grouped using a border than

Design Implications

Perception

Representations of information need to be designed to be perceptible and recognizable across different media.

- Design icons and other graphical representations so that users can readily distinguish between them.
- Obvious separators and white space are effective visual methods for grouping information that make it easier to perceive and locate items.
- Design audio sounds to be readily distinguishable from one another so that users can perceive how they differ and remember what each one represents.
- Research proper color contrast techniques when designing an interface, especially when choosing a color for text so that it stands out from the background. For example, it is okay to use yellow text on a black or blue background, but not on a white or green background.
- Haptic feedback should be used judiciously. The kinds of haptics used should be easily distinguishable so that, for example, the sensation of squeezing is represented in a tactile form that is different from the sensation of pushing. Overuse of haptics can cause confusion. Apple iOS suggests providing haptic feedback in response to user-initiated actions, such as when the action of unlocking a vehicle using a smartwatch has been completed. ■

| | | | |
|---|--|--|--|
| Black Hills Forest Cheyenne River Social Science South San Jose Badlands Park Juvenile Justice | Peters Landing Public Health San Bernardino Moreno Valley Altamonte Springs Peach Tree City | Jefferson Farms Psychophysics Political Science Game Schedule South Addition Cherry Hills Village | Devlin Hall Positions Hubard Hall Fernadino Beach Council Bluffs Classical Lit |
| Results and Stats Thousand Oaks Promotions North Palermo Credit Union Wilner Hall | Highland Park Machesney Park Vallecito Mts. Rock Falls Freeport Slaughter Beach | Creative Writing Lake Havasu City Engineering Bldg Sports Studies Lakewood Village Rock Island | Sociology Greek Wallace Hall Concert Tickets Public Radio FM Children's Museum |
| Performing Arts Italian Coaches Mckees Rocks Glenwood Springs Urban Affairs | Rocky Mountains Latin Pleasant Hills Observatory Public Affairs Heskett Center | Deerfield Beach Arlington Hill Preview Game Richland Hills Experts Guides Neff Hall | Writing Center Theater Auditions Delaware City Scholarships Hendricksville Knights Landing |
| McLeansboro Experimental Links Graduation Emory Lindquist Clinton Hall San Luis Obispo | Brunswick East Millinocket Women's Studies Vacant News Theatre Candlewood Isle | Grand Wash Cliffs Indian Well Valley Online Courses Lindquist Hall Fisk Hall Los Padres Forest | Modern Literature Studio Arts Hugher Complex Cumberland Flats Central Village Hoffman Estates |
| Webmaster Russian Athletics Go Shockers Degree Options Newsletter | Curriculum Emergency (EMS) Statistics Award Documents Language Center Future Shockers | Student Life Accountancy Mc Knight Center Council of Women Commute Small Business | Dance Gerontlogie Marketing College Bylaws Why Wichita? Tickets |
| Gelogy Manufacturing Management UCATS Alumni News Saso | Intercollegiate Bowling Wichita Gateway Transfer Day Job Openings Live Radio | Thinker & Movers Alumni Foundations Corbin Center Jardine Hall Hugo Wall School | Career Services Doers & Shockers Core Values Grace Wilkie Hall Strategic Plan Medical Tech |
| Educational Map Physical Plant Graphic Design Non Credit Class Media Relations Advertising | Beta Alpha Psi Liberal Arts Counseling Biological Science Duerksen Fine Art EMT Program | Staff Aerospace Choral Dept. Alberg Hall French Spanish | Softball, Men's McKinley Hall Email Dental Hygiene Tenure Personnel Policies |
| English Graduate Complex Music Education Advising Center Medical School Levitt Arena | Religion Art Composition Physics Entrepreneurship Koch Arena Roster | Parents Wrestling Philosopy Wichita Lyceum Fairmount Center Women's Museum | Instrmental Nrsing Opera Sports History Athletic Dept. Health Plan |

Figure 4.3 Two ways of structuring information on a web page

Source: Weller (2004)

when using color contrast (Weller, 2004). The findings suggest that using contrasting colors in this manner may not be a good way to group information on a screen, but that using borders is more effective (Galitz, 1997).

4.2.3 Memory

Memory involves recalling various kinds of knowledge that allow people to act appropriately. For example, it allows them to recognize someone's face, remember someone's name, recall when they last met them, and know what they said to them last.

It is not possible for us to remember everything that we see, hear, taste, smell, or touch, nor would we want to, as our brains would get overloaded. A filtering process is used to decide what information gets further processed and memorized. This filtering process, however, is not without its problems. Often, we forget things that we would like to remember and conversely remember things that we would like to forget. For example, we may find it difficult to remember everyday things, like people's names, or scientific knowledge such as mathematical formulae. On the other hand, we may effortlessly remember trivia or tunes that cycle endlessly through our heads.

How does this filtering process work? Initially, encoding takes place, determining which information is paid attention to in the environment and how it is interpreted. The extent to which it takes place affects people's ability to recall that information later. The more attention that is paid to something and the more it is processed in terms of thinking about it and comparing it with other knowledge, the more likely it is to be remembered. For example, when learning about a topic, it is much better to reflect on it, carry out exercises, have discussions with others about it, and write notes rather than passively reading a book or watching a video about it. Thus, how information is interpreted when it is encountered greatly affects how it is represented in memory and how easy it is to retrieve subsequently.

Another factor that affects the extent to which information can be subsequently retrieved is the context in which it is encoded. One outcome is that sometimes it can be difficult for people to recall information that was encoded in a different context from the one in which they are at present. Consider the following scenario:

You are on a train and someone comes up to you and says hello. You don't recognize this person for a few moments, but then you realize it is one of your neighbors. You are only used to seeing them in the hallway of your apartment building and seeing them out of context makes this person initially difficult to recognize.

Another well-known memory phenomenon is that people are much better at recognizing things than recalling things. Furthermore, certain kinds of information is easier to recognize than others. In particular, people are good at recognizing thousands of pictures even if they have only seen them briefly before. In contrast, people are not as good at remembering details about the things they photograph when visiting places, such as museums. It seems that they remember less about objects when they have photographed them than when they observe them with the naked eye (Henkel, 2014). The reason for this is that the study participants appeared to be focusing more on framing the photo and less on the details of the object being photographed. Consequently, people don't process as much information about an object when taking photos of it compared with when they are actually looking at it; hence, they are unable to remember as much about it later.

ACTIVITY 4.2

Try to remember the birthdays of all the members of your family and closest friends. How many can you remember? Then try to describe the image/graphic of the latest app you downloaded.

Comment

It is likely that you remembered the image, the colors, and the name of the app you downloaded much better than the birthdays of your family and friends—most people now rely on Facebook or other online app to remind them about such special dates. People are good at remembering visual cues about things, for example, the color of items, the location of objects (for example, a book being on the top shelf), and marks on an object (like a scratch on a watch, a chip on a cup, and so on). In contrast, people find other kinds of information persistently difficult to learn and remember, especially arbitrary material like phone numbers. ■

Increasingly, people rely on the Internet and their smartphones to act as cognitive prostheses. Smartphones with Internet access have become an indispensable extension of the mind. Sparrow et al. (2011) showed how expecting to have readily available Internet access reduces the need and hence the extent to which people attempt to remember the information itself, while enhancing their memory for knowing where to find it online. Many people will whip out a smartphone to find out who acted in a movie, the name of a book, or what year a pop song was first released, and so on. Besides search engines, there are a number of other cognitive prosthetic apps that instantly help people find out or remember something, such as Shazam.com, the popular music recognition app.

4.2.3.1 *Personal Information Management*

The number of documents written, images created, music files recorded, video clips downloaded, emails with attachments saved, URLs bookmarked, and so on, increases every day. A common practice is for people to store these files on a phone, on a computer, or in the cloud with a view to accessing them later. This is known as *personal information management* (PIM). The design challenge here is deciding which is the best way of helping users organize their content so that it can be easily searched, for example, via folders, albums, or lists. The solution should help users readily access specific items at a later date, for example, a particular image, video, or document. This can be difficult, however, especially when there are thousands or hundreds of thousands of pieces of information available. How does someone find that photo they took of their dog spectacularly jumping into the sea to chase a seagull, which they believe was taken two or three years ago? It can take them ages to wade through the hundreds of folders they have catalogued by date, name, or tag. Do they start by homing in on folders for a given year, looking for events, places, or faces, or typing in a search term to find the specific photo?

It can become frustrating if an item is not easy to locate, especially when users have to spend lots of time opening numerous folders when searching for a particular image or an old document, simply because they can't remember what they called it or where they stored it. How can we improve upon this cognitive process of remembering?

Naming is the most common means of encoding content, but trying to remember a name someone created some time back can be difficult, especially if they have tens of thousands of named files, images, videos, emails, and so forth. How might such a process be facilitated, considering individual memory abilities? Ofer Bergman and Steve Whittaker (2016) have proposed a model for helping people manage their “digital stuff” based on curation. The model involves three interdependent processes: how to decide what personal information to keep, how to organize that information when storing it, and which strategies to use to retrieve it later. The first stage can be assisted by the system they use. For example, email, texts, music, and photos are stored as default by many devices. Users have to decide whether to place these in folders or delete them. In contrast, when browsing the web, they have to make a conscious decision as to whether a site they are visiting is worth bookmarking as one they might want to revisit later.

A number of ways of adding metadata to documents have been developed, including time stamping, categorizing, tagging, and attribution (for example color, text, icon, sound, or image). Surprisingly, however, the majority of people still prefer the old-fashioned way of using folders for holding their files and other digital content. One reason is that folders provide a powerful metaphor (see Chapter 3, “Conceptualizing Interaction”) that people can readily understand—placing things that have something in common into a container.

A folder that is often seen on many users' desktop is one simply labeled “stuff.” This is where documents, images, and so forth, that don't have an obvious place to go are often placed but that people still want to keep somewhere. It has also been found that there is a strong preference for scanning across and within folders when looking for something rather than simply typing a term into a search engine (Ofer and Whittaker, 2016). Part of the problem with using search engines is that it can be difficult to recall the name of the file someone is seeking. This process requires more cognitive effort than navigating through a set of folders.

To help users with searching, a number of search and find tools, such as Apple's Spotlight, now enable them to type a partial name or even the first letter of a file that it then searches for throughout the entire system, including the content inside documents, apps, games, emails, contacts, images, calendars, and applications. Figure 4.4 shows a partial list of files that Spotlight matched to the word *cognition*, categorized in terms of documents, mail and text messages, PDF documents, and so on.

4.2.3.2 *Memory Load and Passwords*

Phone, online, and mobile banking allow customers to carry out financial transactions, such as paying bills and checking the balance of their accounts, at their convenience. One of the problems confronting banks that provide these capabilities, however, is how to manage security concerns, especially preventing fraudulent transactions.

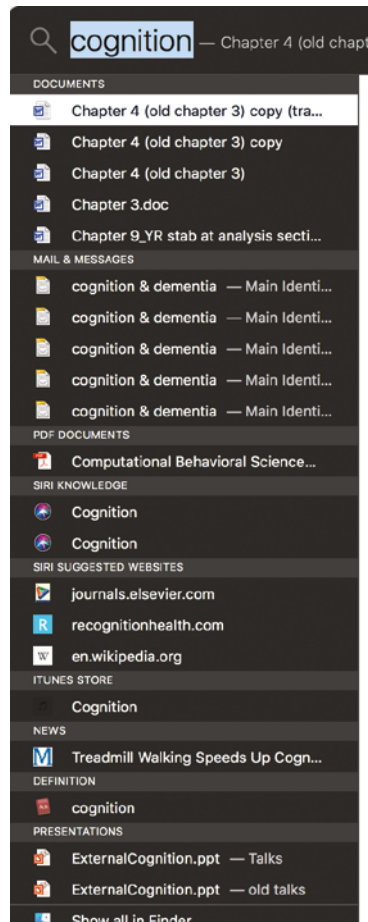


Figure 4.4 Apple's Spotlight search tool

BOX 4.1

The Problem with the Magical Number Seven, Plus or Minus Two

Perhaps the best-known finding in psychology (certainly the one that nearly all students remember many years after they have finished their studies) is George Miller's (1956) theory that seven, plus or minus two, chunks of information can be held in short-term memory at any one time. However, it is also one that has been misapplied in interaction design because several designers assume that it means they should design user interfaces only to have seven, plus or minus two, widgets on a screen, such as menus. In fact, however, this is a misapplication of the phenomenon, as explained here.

By short-term memory, Miller meant a memory store in which information was assumed to be processed when first perceived. By chunks of information, Miller meant a range of items such as numbers, letters, or words. According to Miller's theory, therefore, people's immediate memory capacity is very limited. They are able to remember only a few words or numbers that they have heard or seen. If you are not familiar with this phenomenon, try the following exercise:

Read the first set of numbers here (or get someone to read them to you), cover it up, and then try to recall as many of the items as possible. Repeat this for the other sets.

- 3, 12, 6, 20, 9, 4, 0, 1, 19, 8, 97, 13, 84
- cat, house, paper, laugh, people, red, yes, number, shadow, broom, rain, plant, lamp, chocolate, radio, one, coin, jet
- t, k, s, y, r, q, x, p, a, z, l, b, m, e

How many did you correctly remember for each set? Between five and nine, as suggested by Miller's theory?

Chunks of information can also be combined items that are meaningful. For example, it is possible to remember the same number of two-word phrases like hot chocolate, banana split, cream cracker, rock music, cheddar cheese, leather belt, laser printer, tree fern, fluffy duckling, or cold rain. When these are all jumbled up (that is, split belt, fern crackers, banana laser, printer cream, cheddar tree, rain duckling, or hot rock), however, it is much harder to remember as many chunks. This is mainly because the first set contains all meaningful two-word phrases that have been heard before and that require less time to be processed in short-term memory, whereas the second set is made up of completely novel phrases that don't exist in the real world. You need to spend time linking the two parts of the phrase together while trying to memorize them. This takes more time and effort to achieve. Of course, it is possible to do if you have time to spend rehearsing them, but if you are asked to do it having heard them only once in quick succession, it is most likely that you will remember only a few.

So, how might people's ability to remember only 7 ± 2 chunks of information that they have just read or heard be usefully applied to interaction design? According to a survey by Bob Bailey (2000), several designers have been led to believe the following guidelines and have created interfaces based on them:

- Have only seven options on a menu.
- Display only seven icons on a menu bar.
- Never have more than seven bullets in a list.
- Place only seven tabs at the top of a website page.
- Place only seven items on a pull-down menu.

He points out how this is not how the principle should be applied. The reason is that these are all items that can be scanned and rescanned visually and hence *do not* have to be recalled from short-term memory. They don't just flash up on the screen and disappear, requiring the user to remember them before deciding which one to select. If you were asked to find an item of food most people crave in the set of single words listed earlier, would you have any problem? No, you would just scan the list until you recognized the one (chocolate) that matched the task and then select it—just as people do when interacting with menus, lists, and tabs, regardless of whether they consist of three or 30 items. What users are required to do here is not remember as many items as possible, having only heard or seen them once in a sequence, but instead scan through a set of items until they recognize the one they want. This is a quite different task. ■

One solution has been to develop rigorous security measures whereby customers must provide multiple pieces of information before gaining access to their accounts. This is called *multifactor authentication* (MFA). The method requires a user to provide two or more pieces of evidence that only they know, such as the following:

- Their ZIP code or postal code
- Their mother's maiden name
- Their birthplace
- The last school they attended
- The first school they attended
- A password of between five and ten letters
- A memorable address (not their home)
- A memorable date (not their birthday)

Many of these are relatively easy to remember and recall since they are familiar to the specific user. But consider the last two. How easy is it for someone to come up with such memorable information and then be able to recall it readily? Perhaps the customer can give the address and birthday of another member of their family as a memorable address and date. But what about the request for a password? Suppose a customer selects the word *interaction* as a password—fairly easy to remember, yes? The problem is that banks do not ask for the full password because of the danger that someone in the vicinity might overhear or oversee. Instead, they ask the customer to provide specific letters or numbers from it, like the seventh followed by the fifth. Certainly, such information does not spring readily to mind. Instead, it requires mentally counting each alphanumeric character of the password until the desired one is reached. How long does it take you to determine the seventh letter of the password *interaction*? How did you do it?

To make things harder, banks also randomize the questions they ask. Again, this is to prevent someone else who is nearby from memorizing the sequence of information. However, it also means that the customers themselves cannot learn the sequence of information required, meaning that they have to generate different information each time.

This requirement to remember and recall such information puts a big memory load on customers. Some people find such a procedure quite nerve-racking and are prone to forget certain pieces of information. As a coping strategy, they write down their details on a sheet of paper. Having such an external representation at hand makes it much easier for them to read off the necessary information rather than having to recall it from memory. However, it also makes them vulnerable to the fraud the banks are trying to prevent should anyone else get ahold of that piece of paper! Software companies have also developed password managers to help reduce memory load. An example is LastPass (<https://www.lastpass.com/>), which is designed to remember all of your passwords, meaning that you only have to remember one master password.

ACTIVITY 4.3

How can banks overcome the problem of providing a secure system while making the memory load easier for people wanting to use online and mobile phone banking?

Comment

Advances in computer vision and biometrics technology means that it is now possible to replace the need for passwords to be typed in each time. For example, facial and touch ID can be configured on newer smartphones to enable password-free mobile banking. Once these are set up, a user simply needs to put their face in front of their phone's camera or their finger on the fingerprint sensor. These alternative approaches put the onus on the phone to recognize and authenticate the person rather than the person having to learn and remember a password. ■

BOX 4.2**Digital Forgetting**

Much of the research on memory and interaction design has focused on developing cognitive aids that help people to remember, for example, reminders, to-do lists, and digital photo collections. However, there are times when we want to forget a memory. For example, when someone breaks up with their partner, it can be emotionally painful to be reminded of them through shared digital images, videos, and Facebook friends. How can technology be designed to help people forget such memories? How could social media, such as Facebook, be designed to support this process?

Corina Sas and Steve Whittaker (2013) suggest designing new ways of harvesting digital materials connected to a broken relationship through using various automatic methods, such as facial recognition, which dispose of them without the person needing to go through them personally and be confronted with painful memories. They also suggest that during a separation, people could create a collage of their digital content connected to the ex, so as to transform them into something more abstract, thereby providing a means for closure and helping with the process of moving on. ■

Much research has been conducted into how to design technology to help people suffering from memory loss (for instance those with Alzheimer's disease). An early example was the SenseCam, which was originally developed by Microsoft Research Labs in Cambridge (UK) to enable people to remember everyday events. The device they developed was a wearable camera that intermittently took photos, without any user intervention, while it was worn (see Figure 4.5). The camera could be set to take pictures at particular times, for example, every 30 seconds, or based on what it sensed (for example, acceleration). The camera employed a fish-eye lens, enabling nearly everything in front of the wearer to be captured. The digital images for each day were stored, providing a record of the events that a person experienced. Several studies were conducted on patients with various forms of memory loss using the device. For example, Steve Hodges et al. (2006) describe how a patient, Mrs. B, who had amnesia

was given a SenseCam to wear. The images that were collected were uploaded to a computer at the end of each day. For the next two weeks, Mrs. B and her husband looked through these and talked about them. During this period, Mrs. B's recall of an event nearly tripled, to a point where she could remember nearly everything about that event. Prior to using the SenseCam, Mrs. B would have typically forgotten the little that she could initially remember about an event within a few days.



Figure 4.5 The SenseCam device and a digital image taken with it

Source: Used courtesy of Microsoft Research Cambridge

Since this seminal research, a number of digital memory apps have been developed for people with dementia. For example, RemArc has been designed to trigger long-term memories in people with dementia using BBC Archive material such as old photos, videos, and sound clips.

Design Implications

Memory

- Reduce cognitive load by avoiding long and complicated procedures for carrying out tasks.
- Design interfaces that promote recognition rather than recall by using familiar interaction patterns, menus, icons, and consistently placed objects.
- Provide users with a variety of ways of labeling digital information (for example files, emails, and images) to help them easily identify it again through the use of folders, categories, color, tagging, time stamping, and icons. ■

4.2.4 Learning

Learning is closely connected with memory. It involves the accumulation of skills and knowledge that would be impossible to achieve without memory. Likewise, people would not be able to remember things unless they had learned them. Within cognitive psychology, learning is thought to be either incidental or intentional. *Incidental learning* occurs without any intention to learn. Examples include learning about the world such as recognizing faces, streets, and objects, and what you did today. In contrast, *intentional learning* is goal-directed with the goal of being able to remember it. Examples include studying for an exam, learning a foreign language, and learning to cook. This is much harder to achieve. Software developers, therefore, cannot assume that users will simply be able to learn how to use an app or a product. It often requires much conscious effort.

Moreover, it is well known that people find it hard to learn by reading a set of instructions in a manual. Instead, they much prefer to learn through doing. GUIs and direct manipulation interfaces are good environments for supporting this kind of active learning by supporting exploratory interaction and, importantly, allowing users to undo their actions, that is, return to a previous state if they make a mistake by clicking the wrong option.

There have been numerous attempts to harness the capabilities of different technologies to support intentional learning. Examples include online learning, multimedia, and virtual reality. They are assumed to provide alternative ways of learning through interacting with information that is not possible with traditional technologies, for example, books. In so doing, they have the potential of offering learners the ability to explore ideas and concepts in different ways. For example, multimedia simulations, wearables, and augmented reality (see Chapter 7, “Interfaces”) have been designed to help teach abstract concepts (such as mathematical formulae, notations, laws of physics, biological processes) that students find difficult to grasp. Different representations of the same process (for instance, a graph, formula, sound, or simulation) are displayed and interacted with in ways that make their relationship with each other clearer to the learner.

People often learn effectively when collaborating together. Novel technologies have also been designed to support sharing, turn-taking, and working on the same documents. How these can enhance learning is covered in the next chapter.

Design Implications

Learning

- Design interfaces that encourage exploration.
- Design interfaces that constrain and guide users to select appropriate actions when initially learning. ■

4.2.5 Reading, Speaking, and Listening

Reading, speaking, and listening are three forms of language processing that have some similar and some different properties. One similarity is that the meaning of sentences or phrases is the same regardless of the mode in which it is conveyed. For example, the sentence “Computers are a wonderful invention.” essentially has the same meaning whether one reads it, speaks it, or hears it. However, the ease with which people can read, listen, or speak differs depending on the person, task, and context. For example, many people find listening easier than reading. Specific differences between the three modes include the following:

- Written language is permanent while listening is transient. It is possible to re-read information if not understood the first time around. This is not possible with spoken information that is being broadcast unless it is recorded.
- Reading can be quicker than speaking or listening, as written text can be rapidly scanned in ways not possible when listening to serially presented spoken words.
- Listening requires less cognitive effort than reading or speaking. Children often prefer to listen to narratives provided in multimedia or web-based learning material rather than to read the equivalent text online. The popularity of audiobooks suggests adults also enjoy listening to novels, and so forth.
- Written language tends to be grammatical, while spoken language is often ungrammatical. For example, people often start talking and stop in midsentence, letting someone else start speaking.
- Dyslexics have difficulties understanding and recognizing written words, making it hard for them to write grammatical sentences and spell correctly.

Many applications have been developed either to capitalize on people’s reading, writing, and listening skills, or to support or replace them where they lack or have difficulty with them. These include the following:

- Interactive books and apps that help people to read or learn foreign languages.
- Speech-recognition systems that allow people to interact with them by using spoken commands (for example, Dragon Home, Google Voice Search, and home devices, such as Amazon Echo, Google Home, and Home Aware that respond to vocalized requests).
- Speech-output systems that use artificially generated speech (for instance, written text-to-speech systems for the blind).
- Natural-language interfaces that enable people to type in questions and get written responses (for example, chatbots).
- Interactive apps that are designed to help people who find it difficult to read, write, or speak. Customized input and output devices that allow people with various disabilities to have access to the web and use word processors and other software packages.
- Tactile interfaces that allow people who are visually impaired to read graphs (for example, Designboom’s braille maps for the iPhone).

Design Implications

Reading, Speaking, and Listening

- Keep the length of speech-based menus and instructions to a minimum. Research has shown that people find it hard to follow spoken menus with more than three or four options. Likewise, they are bad at remembering sets of instructions and directions that have more than a few parts.
- Accentuate the intonation of artificially generated speech voices, as they are harder to understand than human voices.
- Provide opportunities for making text large on a screen, without affecting the formatting, for people who find it hard to read small text. ■

4.2.6 Problem-Solving, Planning, Reasoning, and Decision-Making

Problem-solving, planning, reasoning, and decision-making are processes involving reflective cognition. They include thinking about what to do, what the available options are, and what the consequences might be of carrying out a given action. They often involve conscious processes (being aware of what one is thinking about), discussion with others (or oneself), and the use of various kinds of artifacts (for example, maps, books, pens, and paper). Reasoning involves working through different scenarios and deciding which is the best option or solution to a given problem. For example, when deciding on where to go on a vacation, people may weigh the pros and cons of different locations, including cost, weather at the location, availability and type of accommodation, time of flights, proximity to a beach, the size of the local town, whether there is nightlife, and so forth. When weighing all of the options, they reason through the advantages and disadvantages of each before deciding on the best one.

There has been a growing interest in how people make decisions when confronted with information overload, such as when shopping on the web or at a store (Todd et al., 2011). How easy is it to decide when confronted with an overwhelming choice? Classical rational theories of decision-making (for instance, von Neumann and Morgenstern, 1944) posit that making a choice involves weighing up the costs and benefits of different courses of action. This is assumed to involve exhaustively processing the information and making trade-offs between features. Such strategies are very costly in computational and informational terms—not the least because they require the decision-maker to find a way of comparing the different options. In contrast, research in cognitive psychology has shown how people tend to use simple heuristics when making decisions (Gigerenzer et al., 1999). A theoretical explanation is that human minds have evolved to act quickly, making just good enough decisions by using fast and frugal heuristics. We typically ignore most of the available information and rely only on a few important cues. For example, in the supermarket, shoppers make snap judgments based on a paucity of information, such as buying brands that they recognize, that are low-priced, or that offer attractive packaging—seldom reading other package information. This suggests that an effective design strategy is to make key information about a product

highly salient. However, what exactly is salient will vary from person to person. It may depend on the user's preferences, allergies, or interests. For example, one person might have a nut allergy and be interested in food miles, while another may be more concerned about the farming methods used (such as organic, FairTrade, and so on) and a product's sugar content.

Thus, instead of providing ever more information to enable people to compare products when making a choice, a better strategy is to design technological interventions that provide just enough information, and in the right form, to facilitate good choices. One solution is to exploit new forms of augmented reality and wearable technology that enable information-frugal decision-making and that have glanceable displays that can represent key information in an easy-to-digest form (Rogers et al., 2010b). The interface for an AR or wearable app could be designed to provide certain “food” or other information filters, which could be switched on or off by the user to match their preferences.

DILEMMA

Can You Make Up Your Mind Without an App?

In their book *The App Generation* (Yale University Press, 2014), Howard Gardner and Katie Davis note how some young people find it hard to make their own decisions because they are becoming more and more risk averse. The reason for this is that they now rely on using an increasing number of mobile apps to help them in their decision-making, removing the risk of having to decide for themselves. Often, they will first read what others have said on social media sites, blogs, and recommender apps before choosing where to eat or go, what to do or listen to, and so on. However, relying on a multitude of apps means that young people are becoming increasingly unable to make decisions by themselves. For many, their first big decision is choosing which college or university to attend. This has become an agonizing and prolonged experience where both parents and apps play a central role in helping them out. They will read countless reviews, go on numerous visits to colleges and universities with their parents over several months, study university rankings that apply different measures, read up on what others say on social networking sites, and so on. In the end, however, they may finally choose the institution where their friends attend or the one they liked the look of in the first place. ■

Design Implications

Problem-Solving, Planning, Reasoning, and Decision-Making

- Provide information and help pages that are easy to access for people who want to understand more about how to carry out an activity more effectively (for example, web searching).
- Use simple and memorable functions to support rapid decision-making and planning. Enable users to set or save their own criteria or preferences. ■

4.3 Cognitive Frameworks

A number of conceptual frameworks have been developed to explain and predict user behavior based on theories of cognition. In this section, we outline three that focus primarily on mental processes and three others that explain how humans interact and use technologies in the context in which they occur. These are mental models, gulfs of execution and evaluation, information processing, distributed cognition, external cognition, and embodied interaction.

4.3.1 Mental Models

Mental models are used by people when needing to reason about a technology, in particular, to try to fathom what to do when something unexpected happens with it or when encountering unfamiliar products for the first time. The more someone learns about a product and how it functions, the more their mental model develops. For example, broadband engineers have a deep mental model of how Wi-Fi networks work that allows them to work out how to set them up and fix them. In contrast, an average citizen is likely to have a reasonably good mental model of how to use the Wi-Fi network in their home but a shallow mental model of how it works.

Within cognitive psychology, mental models have been postulated as internal constructions of some aspect of the external world that are manipulated, enabling predictions and inferences to be made (Craik, 1943). This process is thought to involve the fleshing out and the running of a mental model (Johnson-Laird, 1983). This can involve both unconscious and conscious mental processes, where images and analogies are activated.

ACTIVITY 4.4

To illustrate how we use mental models in our everyday reasoning, imagine the following two scenarios:

- You arrive home from a vacation on a cold winter's night to a cold house. You have a small baby, and you need to get the house warm as quickly as possible. Your house is centrally heated, but it does not have a smart thermostat that can be controlled remotely. Do you set the thermostat as high as possible or turn it to the desired temperature (for instance, 70°F)?
- You arrive home after being out all night and you're starving hungry. You look in the freezer and find all that is left is a frozen pizza. The instructions on the package say heat the oven to 375°F and then place the pizza in the oven for 20 minutes. Your oven is electric. How do you heat it up? Do you turn it to the specified temperature or higher?

Comment

Most people when asked the first question imagine the scenario in terms of what they would do in their own house and choose the first option. A typical explanation is that setting the temperature to be as high as possible increases the rate at which the room warms up. While many people may believe this, it is incorrect. Thermostats work by switching on the heat and keeping it going at a constant speed until the desired set temperature is reached, at which point it cuts out. They cannot control the rate at which heat is given out from a heating system. Left at a given setting, thermostats will turn the heat on and off as necessary to maintain the desired temperature.

(Continued)

When asked the second question, most people say they would turn the oven to the specified temperature and put the pizza in when they think it is at the right temperature. Some people answer that they would turn the oven to a higher temperature in order to warm it up more quickly. Electric ovens work on the same principle as central heating, so turning the heat up higher will not warm it up any quicker. There is also the problem of the pizza burning if the oven is too hot! ■

Why do people use erroneous mental models? It seems that in the previous two scenarios, they are using a mental model based on a general valve theory of the way something works (Kempton, 1986). This assumes the underlying principle of more is more: the more you turn or push something, the more it causes the desired effect. This principle holds for a range of physical devices, such as faucets, where the more you turn them, the more water that comes out. However, it does not hold for thermostats, which instead function based on the principle of an on-off switch. What seems to happen is that in everyday life, people develop a core set of abstractions about how things work and apply these to a range of devices, irrespective of whether they are appropriate.

Using incorrect mental models to guide behavior is surprisingly common. Just watch people at a pedestrian crossing or waiting for an elevator. How many times do they press the button? A lot of people will press it at least twice. When asked why, a common reason is that they think it will make the lights change faster or ensure the elevator arrives.

Many people's understanding of how technologies and services work is poor, for instance, the Internet, wireless networking, broadband, search engines, computer viruses, the cloud, or AI. Their mental models are often incomplete, easily confusable, and based on inappropriate analogies and superstition (Norman, 1983). As a consequence, they find it difficult to identify, describe, or solve a problem, and they lack the words or concepts to explain what is happening.

How can user experience (UX) designers help people to develop better mental models? A major obstacle is that people are resistant to spending much time learning about how things work, especially if it involves reading manuals or other documentation. An alternative approach is to design technologies to be more transparent, which makes them easier to understand in terms of how they work and what to do when they don't. This includes providing the following:

- Clear and easy-to-follow instructions
- Appropriate online help, tutorials, and context-sensitive guidance for users in the form of online videos and chatbot windows, where users can ask how to do something
- Background information that can be accessed to let people know how something works and how to make the most of the functionality provided
- Affordances of what actions an interface allows (for example, swiping, clicking, or selecting).

The concept of transparency has been used to refer to making interfaces intuitive to use so that people can simply get on with their tasks, such as taking photos, sending messages, or talking to someone remotely without having to worry about long sequences of buttons to

press or options to select. An ideal form of transparency is where the interface simply disappears from the focus of someone’s attention. Imagine if every time you had to give a presentation that all you had to do was say, “Upload and start my slides for the talk I prepared today,” and they would simply appear on the screen for all to see. That would be bliss! Instead, many AV projector systems persist in being far from transparent, requiring many counterintuitive steps for someone to get their slides to show. This can include trying to find the right dongle, setting up the system, typing in a password, setting up audio controls, and so forth, all of which seems to take forever, especially when there is an audience waiting.

4.3.2 Gulfs of Execution and Evaluation

The *gulf of execution* and the *gulf of evaluation* describe the gaps that exist between the user and the interface (Norman, 1986; Hutchins et al., 1986). The gulfs are intended to show how to design the latter to enable the user to cope with them. The first one, the gulf of execution, describes the distance from the user to the physical system while the second one, the gulf of evaluation, is the distance from the physical system to the user (see Figure 4.6). Don Norman and his colleagues suggest that designers and users need to concern themselves with how to bridge the gulfs to reduce the cognitive effort required to perform a task. This can be achieved, on the one hand, by designing usable interfaces that match the psychological characteristics of the user (for example, taking into account their memory limitations) and, on the other hand, by the user learning to create goals, plans, and action sequences that fit with how the interface works.

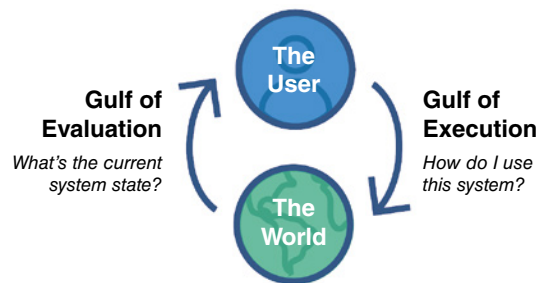


Figure 4.6 Bridging the gulfs of execution and evaluation

Source: <https://www.nngroup.com/articles/two-ux-gulfs-evaluation-execution>. Used courtesy of the Nielsen Norman Group

The conceptual framework of the gulfs is still considered useful today, as it can help designers consider whether their proposed interface design is increasing or decreasing cognitive load and whether it makes it obvious as to which steps to take for a given task. For example, Kathryn Whinton (2018), who is a digital strategy manager, describes how the gulfs prevented her from understanding and why she could not get her Bluetooth headset to connect with her computer despite following the steps in the manual. She wasted a whole hour repeating the steps and getting more and more frustrated and not making any progress. Eventually, she discovered that the system she thought was toggled “on” was actually showing her that it was “off” (see Figure 4.7). She found this out by searching the web to see

whether someone else could help her. She found a site that showed a screenshot of what the settings switch looks like when turned on. There was an inconsistency between the labels of two similar-looking switches, one showing the current status of the interaction (off) and the other showing what would happen if the interaction were engaged (Add Bluetooth Or Other Device).

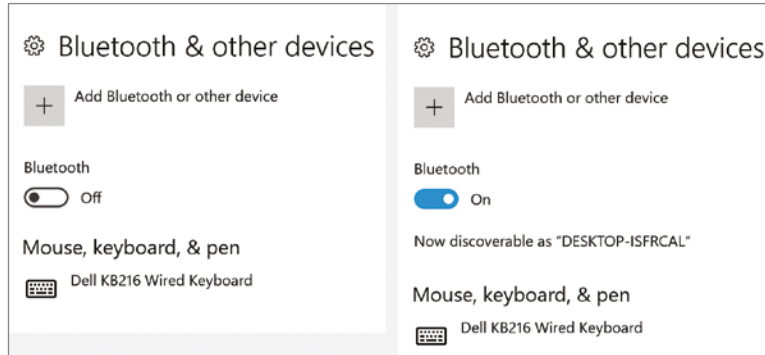


Figure 4.7 An example where the gulfs helped explain how a seemingly trivial design decision led to much user frustration

Source: <https://www.nngroup.com/articles/two-ux-gulfs-evaluation-execution>. Used courtesy of the Nielsen Norman Group

This inconsistency of similar functions illustrated how the gulfs of execution and evaluation were poorly bridged, making it confusing and difficult for the user to know what the problem was or why they could not get their headset to connect with their computer despite many attempts. In the article, she explains how the gulfs could be easily bridged by designing all sliders to give the same information as to what happens when they are moved from one side to the other. For more details about this situation, see <https://www.nngroup.com/articles/two-ux-gulfs-evaluation-execution/>.

4.3.3. Information Processing

Another approach to conceptualizing how the mind works has been to use metaphors and analogies to describe cognitive processes. Numerous comparisons have been made, including conceptualizing the mind as a reservoir, a telephone network, a digital computer, and a deep learning network. One prevalent metaphor from cognitive psychology is the idea that the mind is an information processor. Information is thought to enter and exit the mind through a series of ordered processing stages (see Figure 4.8). Within these stages, various processes are assumed to act upon mental representations. Processes include comparing and matching. Mental representations are assumed to comprise images, mental models, rules, and other forms of knowledge.

The *information processing model* provides a basis from which to make predictions about human performance. Hypotheses can be made about how long someone will take to perceive and respond to a stimulus (also known as *reaction time*) and what bottlenecks occur if a person is overloaded with too much information. One of the first HCI models to




Figure 4.8 Human information processing model

Source: P. Barber (1998). *Applied Cognitive Psychology*. London: Methuen. Used courtesy of Taylor & Francis

be derived from the information processing theory was the human processor model, which modeled the cognitive processes of a user interacting with a computer (Card et al., 1983). Cognition was conceptualized as a series of processing stages, where perceptual, cognitive, and motor processors are organized in relation to one another. The model predicts which cognitive processes are involved when a user interacts with a computer, enabling calculations to be made of how long a user will take to carry out various tasks. In the 1980s, it was found to be a useful tool for comparing different word processors for a range of editing tasks. Even though it is not often used today to inform interaction design, it is considered to be a HCI classic.


The information processing approach was based on modeling mental activities that happen exclusively inside the head. Nowadays, it is more common to understand cognitive activities in the context in which they occur, analyzing cognition as it happens in the wild (Rogers, 2012). A central goal has been to look at how structures in the environment can both aid human cognition and reduce cognitive load. The three external approaches we consider next are distributed cognition, external cognition, and embodied cognition.

4.3.4 Distributed Cognition

Most cognitive activities involve people interacting with external kinds of representations, such as books, documents, and computers and also with each other. For example, when someone goes home from wherever they have been, they do not need to remember the details of the route because they rely on cues in the environment (for instance, they know to turn left at the red house, right when the road comes to a T-junction, and so on). Similarly, when they are at home, they do not have to remember where everything is because information is available as needed. They decide what to eat and drink by scanning the items in the fridge, look out the window to see whether it is raining or not, and so on. Likewise, they are always creating external representations for a number of reasons, not only to help reduce memory load and the cognitive cost of computational tasks, but also, importantly, to extend what they can do and allow people to think more powerfully (Kirsh, 2010).

The *distributed cognition* approach was developed to study the nature of cognitive phenomena across individuals, artifacts, and internal and external representations (Hutchins, 1995). Typically, it involves describing a cognitive system, which entails interactions among people, the artifacts they use, and the environment in which they are working. An example of a cognitive system is an airline cockpit, where the top-level goal is to fly the plane (see Figure 4.9). This involves all of the following:

- The pilot, captain, and air traffic controller interacting with one another
- The pilot and captain interacting with the instruments in the cockpit
- The pilot and captain interacting with the environment in which the plane is flying (that is, the sky, runway, and so on)



Propagation of representational states:

- 1 ATC gives clearance to pilot to fly to higher altitude (verbal)
- 2 Pilot changes altitude meter (mental and physical)
- 3 Captain observes pilot (visual)
- 4 Captain flies to higher altitude (mental and physical)

Figure 4.9 A cognitive system in which information is propagated through different media

A primary objective of the distributed cognition approach is to describe these interactions in terms of how information is propagated through different media. By this we mean how information is represented and re-represented as it moves across individuals and through the array of artifacts that are used (for example, maps, instrument readings, scribbles, and spoken word) during activities. These transformations of information are referred to as changes in representational state.

This way of describing and analyzing a cognitive activity contrasts with other cognitive approaches, such as the information processing model, in that it focuses not on what is happening inside the head of an individual but on what is happening across a system of individuals and artifacts. For example, in the cognitive system of the cockpit, a number of people and artifacts are involved in the activity of flying at a higher altitude. The air traffic controller initially tells the pilot when it is safe to ascend to a higher altitude. The pilot then alerts the captain, who is flying the plane, by moving a knob on the instrument panel in front of them, confirming that it is now safe to fly.

Hence, the information concerning this activity is transformed through different media (over the radio, through the pilot, and via a change in the position of an instrument). This kind of analysis can be used to derive design recommendations, suggesting how to change or

redesign an aspect of the cognitive system, such as a display or a socially mediated practice. In the previous example, distributed cognition could draw attention to the importance of any new design needing to keep shared awareness and redundancy in the system so that both the pilot and the captain can be kept aware and also know that the other is aware of the changes in altitude that are occurring. It is also the basis for the DiCOT analytic framework that has been developed specifically for understanding healthcare settings and has also been used for software team interactions (see Chapter 9, “Data Analysis”).

4.3.5 External Cognition

People interact with or create information by using a variety of external representations, including books, multimedia, newspapers, web pages, maps, diagrams, notes, drawings, and so on. Furthermore, an impressive range of tools has been developed throughout history to aid cognition, including pens, calculators, spreadsheets, and software workflows. The combination of external representations and physical tools has greatly extended and supported people’s ability to carry out cognitive activities (Norman, 2013). Indeed, they are such an integral part of our cognitive activities that it is difficult to imagine how we would go about much of our everyday life without them.

External cognition is concerned with explaining the cognitive processes involved when we interact with different external representations such as graphical images, multimedia, and virtual reality (Scaife and Rogers, 1996). A main goal is to explain the cognitive benefits of using different representations for different cognitive activities and the processes involved. The main ones include the following:

- Externalizing to reduce memory load
- Computational offloading
- Annotating and cognitive tracing

4.3.5.1 Externalizing to Reduce Memory Load

Numerous strategies have been developed for transforming knowledge into external representations to reduce memory load. One such strategy is externalizing things that we find difficult to remember, such as birthdays, appointments, and addresses. Diaries, personal reminders, and calendars are examples of cognitive artifacts that are commonly used for this purpose, acting as external reminders of what we need to do at a given time, such as buy a card for a relative’s birthday.

Other kinds of external representations that people frequently employ are notes, such as sticky notes, shopping lists, and to-do lists. Where these are placed in the environment can also be crucial. For example, people often place notes in prominent positions, such as on walls, on the side of computer screens, by the front door, and sometimes even on their hands in a deliberate attempt to ensure that they do remind them of what needs to be done or remembered. People also place things in piles in their offices and by the front door, indicating what needs to be done urgently versus what can wait for a while.

Externalizing, therefore, can empower people to trust that they will be reminded without having to remember themselves, thereby reducing their memory burden in the following ways:

- Reminding them to do something (for example, get something for mother’s birthday)
- Reminding them of what to do (such as buy a card)
- Reminding them of when to do something (for instance, send it by a certain date)

This is an obvious area where technology can be designed to help remind. Indeed, many apps have been developed to reduce the burden on people to remember things, including to-do and alarm-based lists. These can also be used to help improve people's time management and work-life balance.

4.3.5.2 *Computational Offloading*

Computational offloading occurs when we use a tool or device in conjunction with an external representation to help us carry out a computation. An example is using pen and paper to solve a math problem as mentioned in the introduction of the chapter where you were asked to multiply 21×19 in your head versus using a pen and paper. Now try doing the sum again but using roman numerals: XXI \times XVIII. It is much harder unless you are an expert in using roman numerals—even though the problem is equivalent under both conditions. The reason for this is that the two different representations transform the task into one that is easy and one that is more difficult, respectively. The kind of tool used also can change the nature of the task to being easier or more difficult.

4.3.5.3 *Annotating and Cognitive Tracing*

Another way in which we externalize our cognition is by modifying representations to reflect changes that are taking place that we want to mark. For example, people often cross things off a to-do list to indicate tasks that have been completed. They may also reorder objects in the environment by creating different piles as the nature of the work to be done changes. These two types of modification are called *annotating* and *cognitive tracing*.

- Annotating involves modifying external representations, such as crossing off or underlining items.
- Cognitive tracing involves externally manipulating items into different orders or structures.

Annotating is often used when people go shopping. People usually begin their shopping by planning what they are going to buy. This often involves looking in their cupboards and fridge to see what needs stocking up. However, many people are aware that they won't remember all this in their heads, so they often externalize it as a written shopping list. The act of writing may also remind them of other items that they need to buy, which they may not have noticed when looking through the cupboards. When they actually go shopping at the store, they may cross off items on the shopping list as they are placed in the shopping basket or cart. This provides them with an annotated externalization, allowing them to see at a glance what items are still left on the list that need to be bought.

There are a number of digital annotation tools that allow people to use pens, styluses, or their fingers to annotate documents, such as circling data or writing notes. The annotations can be stored with the document, enabling the users to revisit theirs or others' externalizations at a later date.

Cognitive tracing is useful in conditions where the current situation is in a state of flux and the person is trying to optimize their position. This typically happens when playing games, such as the following:

- In a card game, when the continuous rearrangement of a hand of cards into suits, in ascending order, or collecting same numbers together helps to determine what cards to keep and which to play as the game progresses and tactics change

- In Scrabble, where shuffling letters around in the tray helps a person work out the best word given the set of letters (Maglio et al., 1999)

Cognitive tracing has also been used as an interactive function, for example, letting students know what they have studied in an online learning package. An interactive diagram can be used to highlight all of the nodes visited, exercises completed, and units still to be studied.

A general cognitive principle for interaction design based on the external cognition approach is to provide external representations at an interface that reduce memory load, support creativity, and facilitate computational offloading. Different kinds of information visualizations can be developed that reduce the amount of effort required to make inferences about a given topic (for example, financial forecasting or identifying programming bugs). In so doing, they can extend or amplify cognition, allowing people to perceive and do activities that they couldn't do otherwise. For example, information visualizations (discussed in chapter 10) are used to represent big data in a visual form that can make it easier to make cross-comparisons across dimensions and see patterns and anomalies. Workflow and contextual dialog boxes can also pop up at appropriate times to guide users through their interactions, especially where there are potentially hundreds and sometimes thousands of options available. This reduces memory load significantly and frees up more cognitive capacity for enabling people to complete desired tasks.

4.3.6 Embodied Interaction

Another way of describing our interactions with technology and the world is to conceive of it as embodied. By this we mean the practical engagement with the social and physical environment (Dourish, 2001). This involves creating, manipulating, and making meaning through our engaged interaction with physical things, including mundane objects such as cups and spoons, and technological devices, such as smartphones and robots. Artifacts and technologies that indicate how they are coupled to the world make it clear how they should be used. For example, a physical artifact, like a book when left opened on someone's desk, can remind them to complete an unfinished task the next day (Marshall and Hornecker, 2013).

Eva Hornecker et al. (2017) further explain *embodied interaction* in terms of how our bodies and active experiences shape how we perceive, feel, and think. They describe how our ability to think abstractly is thought to be a result of our sensorimotor experiences with the world. This enables us to learn how to think and talk using abstract concepts, such as inside-outside, up-down, on top of, and behind. Our numerous experiences of moving through and manipulating the world since we were born (for example, climbing, walking, crawling, stepping into, holding, or placing) is what enables us to develop a sense of the world at both a concrete and abstract level.

Within HCI, the concept of embodied interaction has been used to describe how the body mediates our various interactions with technology (Klemmer et al., 2006) and also our emotional interactions (Höök, 2018). By theorizing about embodied interactions in these ways has helped researchers uncover problems that can arise in the use of existing technologies while also informing the design of new technologies in the context in which they will be used.

David Kirsh (2013) suggests that a theory of embodiment can provide HCI practitioners and theorists with new ideas about interaction and new principles for better designs. He explains how interacting with tools changes the way people think and perceive of their environments. He also argues that a lot of times we think with our bodies and not just with our brains. He studied choreographers and dancers and observed that they often partially model

a dance (known as *marking*) through using abbreviated moves and small gestures rather than doing a full workout or mentally simulating the dance in their heads. This kind of marking was found to be a better method of practice than the other two methods. The reason for doing it this way is not that it is saving energy or preventing dancers from getting exhausted emotionally, but that it enables them to review and explore particular aspects of a phrase or movement without the mental complexity involved in a full work out. The implication of how people use embodiment in their lives is that learning new procedures and skills might be better taught by a process like marking, where learners create little models of things or use their own bodies to act out. For example, rather than developing fully fledged virtual reality simulations for learning golf, tennis, skiing, and so on, it might be better to teach sets of abbreviated actions, using augmented reality, as a form of embodied marking.

In-Depth Activity

The aim of this in-depth activity is for you to try to elicit mental models from people. In particular, the goal is for you to understand the nature of people's knowledge about an interactive product in terms of how to use it and how it works.

1. First, elicit your own mental model. Write down how you think contactless cards (see Figure 4.10) work—where customers place their debit or credit card over a card reader. If you are not familiar with contactless cards, do the same for a smartphone app like Apple Pay or Google Pay. Then answer the following questions:



Figure 4.10 A contactless debit card indicated by symbol

- What information is sent between the card/smartphone and the card reader when it is placed in front of it?
- What is the maximum amount you can pay for something using a contactless card, or Apple/Google Pay?
- Why is there an upper limit?
- How many times can you use a contactless card or Apple/Google Pay in a day?
- What happens if you have two contactless cards in the same wallet/purse?
- What happens when your contactless card is stolen and you report it to the bank?

Next, ask two other people the same set of questions.

2. Now analyze your answers. Do you get the same or different explanations? What do the findings indicate? How accurate are people's mental models about the way contactless cards and smartphone Apple/Google Pay work?

Summary

This chapter explained the importance of understanding the cognitive aspects of interaction. It described relevant findings and theories about how people carry out their everyday activities and how to learn from these to help in designing interactive products. It provided illustrations of what happens when you design systems with the user in mind and what happens when you don't. It also presented a number of conceptual frameworks that allow ideas about cognition to be generalized across different situations.

Key points

- Cognition comprises many processes, including thinking, attention, memory, perception, learning, decision-making, planning, reading, speaking, and listening.
- The way in which an interface is designed can greatly affect how well people can perceive, attend, learn, and remember how to carry out their tasks.
- The main benefits of conceptual frameworks based on theories of cognition are that they can explain user interaction, inform design, and predict user performance.

Further Reading

BERGMAN, O. and WHITTAKER, S. (2016). *The Science of Managing Our Digital Stuff*. MIT Press. This very readable book provides a fascinating account of how we manage all of our digital stuff that increases by the bucket load each day. It explains why we persist with seemingly old-fashioned methods when there are alternative, seemingly better approaches that have been designed by software companies.

ERICKSON, T. D. and MCDONALD, D. W. (2008) *HCI Remixed: Reflections on Works That Have Influenced the HCI Community*. MIT Press. This collection of essays from more than 50 leading HCI researchers describes the accessible prose papers, books, and software that influenced their approach to HCI and shaped its history. They include some of the classic papers on cognitive theories, including the psychology of HCI and the power of external representations.

EYSENCK, M. and BRYBAERT, M. (2018) *Fundamentals of Cognition* (3rd ed.). Routledge. This introductory textbook about cognition provides a comprehensive overview of the fundamentals of cognition. In particular, it describes the processes that allow us to make sense of the world around us and to enable us to make decisions about how to manage our everyday lives. It also covers how technology can provide new insights into how the mind works, for example, revealing how CAPTCHAs tell us more about perception.

GIGERENZER, G. (2008) *Gut Feelings*. Penguin. This provocative paperback is written by a psychologist and behavioral expert in decision-making. When confronted with choice in a variety of contexts, he explains how often “less is more.” He explains why this is so in terms of how people rely on fast and frugal heuristics when making decisions, which are often unconscious rather than rational. These revelations have huge implications for interaction design that are only just beginning to be explored.

JACKO, J. (ed.) (2012) *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications* (3rd ed). CRC Press. Part 1 is about human aspects of HCI and includes in-depth chapters on information processing, mental models, decision-making, and perceptual motor interaction.

KAHNEMAN, D. (2011) *Thinking, Fast and Slow*. Penguin. This bestseller presents an overview of how the mind works, drawing on aspects of cognitive and social psychology. The focus is on how we make judgments and choices. It proposes that we use two ways of thinking: one that is quick and based on intuition and one that is slow and more deliberate and challenging. The book explores the many facets of life and how and when we use each.

Chapter 5

SOCIAL INTERACTION

5.1 Introduction

5.2 Being Social

5.3 Face-to-Face Conversations

5.4 Remote Conversations

5.5 Co-presence

5.6 Social Engagement

Objectives

The main goals of the chapter are to accomplish the following:

- Explain what is meant by social interaction.
- Describe the social mechanisms that people use to communicate and collaborate.
- Explain what social presence means.
- Give an overview of new technologies intended to facilitate collaboration and group participation.
- Discuss how social media has changed how we keep in touch, make contacts, and manage our social and working lives.
- Outline examples of new social phenomena that are a result of being able to connect online.

5.1 Introduction

People are inherently social: we live together, work together, learn together, play together, interact and talk with each other, and socialize. A number of technologies have been developed specifically to enable us to persist in being social when physically apart from one another, many of which have now become part of the fabric of society. These include the widespread use of smartphones, video chat, social media, gaming, messaging, and telepresence. Each of these afford different ways of supporting how people connect.

There are many ways to study what it means to be social. In this chapter, we focus on how people communicate and collaborate face-to-face and remotely in their social, work,

and everyday lives—with the goal of providing models, insights, and guidelines to inform the design of “social” technologies that can better support and extend them. A diversity of communication technologies is also examined that have changed the way people live—how they keep in touch, make friends, and coordinate their social and work networks. The conversation mechanisms that have conventionally been used in face-to-face interactions are described and discussed in relation to how they have been adapted for the various kinds of computer-based conversations that now take place at a distance. Examples of social phenomena that have emerged as a result of social engagement at scale are also presented.

5.2 Being Social

A fundamental aspect of everyday life is being social, and that entails interacting with each other. People continually update each other about news, changes, and developments on a given project, activity, person, or event. For example, friends and families keep each other posted on what’s happening at work, at school, at a restaurant or club, next door, in reality shows, and in the news. Similarly, people who work together keep each other informed about their social lives and everyday events, as well as what is happening at work, for instance when a project is about to be completed, plans for a new project, problems with meeting deadlines, rumors about closures, and so on.

While face-to-face conversations remain central to many social interactions, the use of social media has dramatically increased. People now spend several hours a day communicating with others online—texting, emailing, tweeting, Facebooking, Skyping, instant messaging, and so on. It is also common practice for people at work to keep in touch with each other via WhatsApp groups and other workplace communication tools, such as Slack, Yammer, or Teams.

The almost universal adoption of social media in mainstream life has resulted in most people now being connected in multiple ways over time and space—in ways that were unimaginable 25 or even 10 years ago. For example, adults average about 338 Facebook friends, while it is increasingly common for people to have more than 1,000 connections on LinkedIn—many more than those made through face-to-face networking. The way that people make contact, how they stay in touch, who they connect to, and how they maintain their social networks and family ties have irrevocably changed. During the last 20 or so years, social media, teleconferencing, and other social-based technologies (often referred to as *social computing*) have also transformed how people collaborate and work together globally—including the rise of flexible and remote working, the widespread use of shared calendars and collaboration tools (for example Slack, Webex, Trello, and Google Docs), and professional networking platforms (such as LinkedIn, Twitter, and WhatsApp).

A key question that the universal adoption of social media and other social computing tools in society raises is how it has affected people’s ability to connect, work, and interact with one another. Have the conventions, norms, and rules established in face-to-face interactions to maintain social order been adopted in social media interactions, or have new norms emerged? In particular, are the established conversational rules and etiquette, whose function it is to let people know how they should behave in social groups, also applicable to online social behavior? Or, have new conversational mechanisms evolved for the various kinds of social media? For example, do people greet each other in the same way, depending on whether they are chatting online, Skyping, or at a party? Do people take turns when online

chatting in the way they do when talking with each other face to face? How do they choose which technology or app to use from the variety available today for their various work and social activities, such as SnapChat, text messaging, Skype, or phone calls? Answering these questions can help us understand how existing tools support communication and collaborative work while helping to inform the design of new ones.

When planning and coordinating social activities, groups often switch from one mode to another. Most people send texts in preference to calling someone up, but they may switch to calling or mobile group messaging (such as WhatsApp, GroupMe) at different stages of planning to go out (Schuler et al., 2014). However, there can be a cost as conversations about what to do, where to meet, and who to invite multiply across people. Some people might get left off or others might not reply, and much time can be spent to-ing and fro-ing across the different apps and threads. Also, some people may not look at their notifications in a timely manner, while further developments in the group planning have evolved. This is compounded by the fact that often people don't want to commit until close to the time of the event, in case an invitation to do something from another friend appears that is more interesting to them. Teenagers, especially, often leave it until the last minute to micro-coordinate their arrangements with their friends before deciding on what to do. They will wait and see if a better offer comes their way rather than deciding for themselves a week in advance, say, to see a movie with a friend and sticking to it. This can make it frustrating for those who initiate the planning and are waiting to book tickets before they sell out.

A growing concern that is being raised within society is how much time people spend looking at their phones—whether interacting with others, playing games, tweeting, and so forth—and its consequences on people's well-being (see Ali et al., 2018). A report on the impact of the “decade of a smartphone” notes that on average a person in the United Kingdom spends more than a day a week online (Ofcom, 2018). Often, it is the first thing they do upon waking and the last thing they do before going to bed. Moreover, lots of people cannot go for long without checking their phone. Even when sitting together, they resort to being in their own digital bubbles (see Figure 5.1). Sherry Turkle (2015) bemoans the negative impact that this growing trend is having on modern life, especially how it is affecting everyday conversation. She points out that many people will admit to preferring texting to talking to others, as it is easier, requires less effort, and is more convenient. Furthermore, her research has shown that when children hear adults talking less, they likewise talk less. This in turn reduces opportunities to learn how to empathize. She argues that while online communication has its place in society, it is time to reclaim conversation, where people put down their phones more often and (re)learn the art and joy of spontaneously talking to each other.

On the other hand, it should be stressed that several technologies have been designed to encourage social interaction to good effect. For example, voice assistants that come with smart speakers, such as Amazon's Echo devices, provide a large number of “skills” intended to support multiple users taking part at the same time, offering the potential for families to play together. An example skill is “Open the Magic Door,” which allows group members (such as families) to choose their path in a story by selecting different options through the narrative. Social interaction may be further encouraged by the affordance of a smart speaker when placed on a surface in the home, such as a kitchen counter or mantelpiece. In particular, its physical presence in this shared location affords joint ownership and use—similar to other domestic devices, such as the radio or TV. This differs from other virtual voice assistants that are found on phones or laptops that support individual use.



Figure 5.1 A family sits together, but they are all in their own digital bubbles—including the dog!

Source: Helen Sharp

ACTIVITY 5.1

Think of a time where you enjoyed meeting up with friends to catch up in a cafe. Compare this social occasion with the experience that you have when texting with them on your smartphone. How are the two kinds of conversations different?

Comment

The nature of the conversations is likely to be very different with pros and cons for each. Face-to-face conversations ebb and flow unpredictably and spontaneously from one topic to the next. There can be much laughing, gesturing, and merriment among those taking part in the conversation. Those present pay attention to the person speaking, and then when someone else starts talking, all eyes move to them. There can be much intimacy through eye contact, facial expressions, and body language, in contrast to when texters send intermittent messages back and forth in bursts of time. Texting is also more premeditated; people decide what to say and can review what they have written. They can edit their message or decide even not to send it, although sometimes people press the Send button without much thought about its impact on the interlocutor that can lead to regrets afterward.

Emoticons are commonly used as a form of expressivity to compensate for nonverbal communication. While they can enrich a message by adding humor, affection, or a personal touch, they are nothing like a real smile or a wink shared at a key moment in a conversation. Another difference is that people say things and ask each other things in conversations that they would never do via text. On the one hand, such confiding and directness may be more engaging and enjoyable, but on the other hand, it can sometimes be embarrassing. It depends on the context as to whether conversing face-to-face versus texting is preferable. ■

5.3 Face-to-Face Conversations

Talking is something that is effortless and comes naturally to most people. And yet holding a conversation is a highly skilled collaborative achievement, having many of the qualities of a musical ensemble. In this section we examine what makes up a conversation. Understanding how conversations start, progress, and finish is useful when designing dialogues that take place with chatbots, voice assistants, and other communication tools. In particular, it helps researchers and developers understand how natural it is, how comfortable people are when conversing with digital agents, and the extent to which it is important to follow conversation mechanisms that are found in human conversations. We begin by examining what happens at the beginning.

A: Hi there.

B: Hi!

C: Hi.

A: All right?

C: Good. How's it going?

A: Fine, how are you?

C: Good.

B: OK. How's life treating you?

Such mutual greetings are typical. A dialogue may then ensue in which the participants take turns asking questions, giving replies, and making statements. Then, when one or more of the participants wants to draw the conversation to a close, they do so by using either implicit or explicit cues. An example of an implicit cue is when a participant looks at their watch, signaling indirectly to the other participants that they want the conversation to draw to a close. The other participants may choose to acknowledge this cue or carry on and ignore it. Either way, the first participant may then offer an explicit signal, by saying, “Well, I have to go now. I got a lot of work to do” or, “Oh dear, look at the time. I gotta run. I have to meet someone.” Following the acknowledgment by the other participants of such implicit and explicit signals, the conversation draws to a close, with a farewell ritual. The different participants take turns saying, “Goodbye,” “Bye,” “See you,” repeating themselves several times until they finally separate.

ACTIVITY 5.2

How do you start and end a conversation when (1) talking on the phone and (2) chatting online? Do you use the same conversational mechanisms that are used in face-to-face conversations?

Comment

The person answering the call will initiate the conversation by saying “hello” or, more formally, the name of their company/department. Most phones (landline and smartphones) have the facility to display the name of the caller (Caller ID) so the receiver can be more personal when answering, for example “Hello, John. How are you doing?” Phone conversations usually start with a mutual greeting and end with a mutual farewell. In contrast, conversations that take place when chatting online have evolved new conventions. The use of opening and ending greetings when joining and leaving is rare; instead, most people simply start their message with what they want to talk about and then stop when they have gotten an answer, as if in the middle of a conversation. ■

Many people are now overwhelmed by the number of emails they receive each day and find it difficult to reply to them all. This has raised the question of which conversational techniques to use to improve the chances of getting someone to reply. For example, can the way people compose their emails, especially the choice of opening and ending a conversation, increase the likelihood that the recipient will respond to it? A study by Boomerang (Brendan G, 2017) of 300,000 emails taken from mailing list archives of more than 20 different online communities examined whether the opening or closing phrase that was used affected the reply rate. They found that the most common opening phrase used “hey” (64 percent), followed by “hello” (63 percent), and then “hi” (62 percent) were the ones that got the highest rate of reply, in the region of 63–64 percent. This was found to be higher than emails that opened with more formal phrases, like “Dear” (57 percent) or “Greetings” (56 percent). The most popular form of sign-off was found to be “thanks” (66 percent), “regards” (63 percent), and “cheers” (58 percent), with “best” being used less (51 percent). Again, they found that emails that used closings with a form of “thank you” got the highest rate of responses. Hence, which conversational mechanism someone uses to address the recipient can determine whether they will reply to it.

Conversational mechanisms enable people to coordinate their talk with one another, allowing them to know how to start and stop. Throughout a conversation, further turn-taking rules are followed that enable people to know when to listen, when it is their cue to speak, and when it is time for them to stop again to allow the others to speak. Sacks et al. (1978), famous for their work on conversation analysis, describe these in terms of three basic rules.

Rule 1 The current speaker chooses the next speaker by asking a question, inviting an opinion, or making a request.

Rule 2 Another person decides to start speaking.

Rule 3 The current speaker continues talking.

The rules are assumed to be applied in this order so that whenever there is an opportunity for a change of speaker to occur, for instance, someone comes to the end of a sentence, rule 1 is applied. If the listener to whom the question or request is addressed does not accept the offer to take the floor, rule 2 is applied, and someone else taking part in the conversation may take up the opportunity and offer a view on the matter. If this does not happen, then rule 3 is applied, and the current speaker continues talking. The rules are cycled through recursively until someone speaks again.

To facilitate rule following, people use various ways of indicating how long they are going to talk and on what topic. For example, a speaker might say right at the beginning of his turn in the conversation that he has three things to say. A speaker may also explicitly request a change in speaker by saying to the listeners, “OK, that’s all I want to say on that matter. So, what do you think?” More subtle cues to let others know that their turn in the conversation is coming to an end include the lowering or raising of the voice to indicate the end of a question or the use of phrases like “You know what I mean?” or simply “OK?” Back channeling (uh-huh, mmm), body orientation (such as moving away from or closer to someone), gaze (staring straight at someone or glancing away), and gesturing (for example, raising of arms) are also used in different combinations when talking in order to signal to others when someone wants to hand over or take up a turn in the conversation.

Another way in which conversations are coordinated and given coherence is through the use of adjacency pairs (Schegloff and Sacks, 1973). Utterances are assumed to come in pairs in which the first part sets up an expectation of what is to come next and directs the way in which what does come next is heard. For example, A may ask a question to which B responds appropriately.

A: So, shall we meet at 8:00?

B: Um, can we make it a bit later, say 8:30?

Sometimes adjacency pairs get embedded in each other, so it may take some time for a person to get a reply to their initial request or statement.

A: So, shall we meet at 8:00?

B: Wow, look at them.

A: Yes, what a funny hairdo!

B: Um, can we make it a bit later, say 8:30?

For the most part, people are not aware of following conversational mechanisms and would be hard-pressed to articulate how they can carry on a conversation. Furthermore, people don't necessarily abide by the rules all the time. They may interrupt each other or talk over each other, even when the current speaker has clearly indicated a desire to hold the floor for the next two minutes to finish an argument. Alternatively, a listener may not take up a cue from a speaker to answer a question or take over the conversation but instead continue to say nothing even though the speaker may be making it glaringly obvious that it is the listener's turn to say something. Oftentimes, a teacher will try to hand over the conversation to a student in a seminar by staring at them and asking a specific question, only to see the student look at the floor and say nothing. The outcome is an embarrassing silence, followed by either the teacher or another student picking up the conversation again.

Other kinds of breakdowns in conversation arise when someone says something that is ambiguous, and the interlocutor misinterprets it to mean something else. In such situations, the participants will collaborate to overcome the misunderstanding by using repair mechanisms. Consider the following snippet of conversation between two people:

A: Can you tell me the way to get to the Multiplex Ranger cinema?

B: Yes, you go down here for two blocks and then take a right (pointing to the right), proceed until you get to the light, and then it's on the left.

A: Oh, so I go along here for a couple of blocks and then take a right, and the cinema is at the light (pointing ahead of him)?

B: No, you go down *this* street for a couple of blocks (gesturing more vigorously than before to the street to the right of him while emphasizing the word *this*).

A: Ahhhh! I thought you meant *that* one: so it's *this* one (pointing in the same direction as the other person).

B: Uh-hum, yes, that's right: *this* one.

Detecting breakdowns in conversation requires that the speaker and listener both pay attention to what the other says (or does not say). Once they have understood the nature of the failure, they can then go about repairing it. As shown in the previous example, when the listener misunderstands what has been communicated, the speaker repeats what they said earlier, using a stronger voice intonation and more exaggerated gestures. This allows the speaker to repair the mistake and be more explicit with the listener, allowing them to understand and follow better what they are saying. Listeners may also signal when they don't

understand something or want further clarification by using various tokens, like “Huh?” or “What?” (Schegloff, 1981), together with giving a puzzled look (usually frowning). This is especially the case when the speaker says something that is vague. For example, they might say “I want it” to their partner, without saying what *it* is they want. The partner may reply using a token or, alternatively, explicitly ask, “What do you mean by *it*?” Nonverbal communication also plays an important role in augmenting face-to-face conversation, involving the use of facial expressions, back channeling, voice intonation, gesturing, and other kinds of body language.

Taking turns also provides opportunities for the listener to initiate repair or request clarification or for the speaker to detect that there is a problem and initiate repair. The listener will usually wait for the next turn in the conversation before interrupting the speaker in order to give the speaker the chance to clarify what is being said by completing the utterance.

ACTIVITY 5.3

How do people repair breakdowns when conversing via email? Do they do the same when texting?

Comment

As people usually cannot see each other when communicating by email or text, they have to rely on other means of repairing the conversation when things are left unsaid or are unclear. For example, when someone proposes an ambiguous meeting time, where the date and day given don’t match up for the month, the person receiving the message may begin their reply by asking politely, “Did you mean this month or June?” rather than baldly stating the other person’s error, for example, “the 13th May is not a Wednesday!”

When someone does not reply to an email or text when the sender is expecting them to do so, it can put them in a quandary as to what to do next. If someone does not reply to an email within a few days, then the sender might send them a gentle nudge message that diminishes any blame, for example, “I am not sure if you got my last email as I was using a different account” rather than explicitly asking them why they have not answered the email they sent. When texting, it depends on whether it is a dating, family, or business-related text that has been sent. When starting to date, some people will deliberately wait a while before replying to a text as a form of “playing games” and trying not to appear to be overly keen. If they don’t reply at all, it is a generally accepted notion that they are not interested, and no further texts should be sent. In contrast, in other contexts, double-texting has become an acceptable social norm as a way of reminding someone, without sounding too rude, to reply. It implicitly implies that the sender understands that the recipient has overlooked the first text because they were too busy or doing something else at the time, thereby saving face.

Emails and texts can also appear ambiguous, especially when things are left unsaid. For example, the use of an ellipsis (...) at the end of a sentence can make it difficult to work out what the sender intended when using it. Was it to indicate something was best left unsaid, the sender is agreeing to something but their heart is not in it, or simply that the sender did not know what to say? This email or text convention puts the onus on the receiver to decide what is meant by the ellipsis and not on the sender to explain what they meant. ■

5.4 Remote Conversations

The telephone was invented in the nineteenth century by Alexander Graham Bell, enabling two people to talk to one another at a distance. Since then, a number of other technologies have been developed that support synchronous remote conversations, including videophones that were developed in the 1960s–1970s (see Figure 5.2). In the late 1980s and 1990s, a range of “media spaces” were the subjects of experimentation—audio, video, and computer systems were combined to extend the world of desks, chairs, walls, and ceilings (Harrison, 2009). The goal was to see whether it was possible for people, distributed over space and different time zones, could communicate and interact with one another as if they were actually physically present.



Figure 5.2 One of British Telecom’s early videophones

Source: British Telecommunications Plc

An example of an early media space was the VideoWindow (Bellcore, 1989) that was developed to enable people in different locations to carry on a conversation as they would do if they were drinking coffee together in the same room (see Figure 5.3). Two lounge areas that were 50 miles apart were connected via a 3-foot by 5-foot picture window onto which video images of each location were projected. The large size enabled viewers to see a room of people roughly the same size as themselves. A study of its use showed that many of the conversations that took place between the remote conversants were indeed indistinguishable from similar face-to-face interactions, with the difference being that they spoke a bit louder and constantly talked about the video system (Kraut et al., 1990). Other research on how people interact when using videoconferencing has shown that they tend to project themselves more, take longer conversational turns, and interrupt each other less (O’Connell et al., 1993).

Since this early research, videoconferencing has come of age. The availability of cheap webcams and cameras now embedded as a default in tablets, laptops, and phones has



Figure 5.3 Diagram of VideoWindow system in use

greatly helped make videoconferencing mainstream. There are now numerous platforms available from which to choose, both free and commercial. Many videoconferencing apps (for example, Zoom or Meeting Owl) also allow multiple people at different sites to connect synchronously. To indicate who has the floor, screen effects are often used, such as enlarging the person who is talking to take up most of the screen or highlighting their portal when they take the floor. The quality of the video has also improved, making it possible for people to appear more life-like in most setups. This is most noticeable in high-end telepresence rooms that use multiple high-definition cameras with eye-tracking features and directional microphones (see Figure 5.4). The effect can be to make remote people appear more present by projecting their body movements, actions, voice, and facial expressions to the other location.

Another way of describing this development is in terms of the degree of telepresence. By this we mean the perception of being there when physically remote. Robots, for example, have been built with telepresence in mind to enable people to attend events and communicate with others by controlling them remotely. Instead of sitting in front of a screen from their



Figure 5.4 A telepresence room

Source: Cisco Systems, Inc.

location and seeing the remote place solely through a fixed camera at the other place, they can look around the remote place by controlling the “camera’s” eyes, which are placed on the robot and are physically moving it around. For example, telepresence robots have been developed to enable children who are in a hospital to attend school by controlling their assigned robots to roam around the classroom (Rae et al., 2015).

Telepresence robots are also being investigated to determine whether they can help people who have developmental difficulties visit places remotely, such as museums. Currently, several of the activities that are involved in going on such a visit, such as buying a ticket or using public transport, are cognitively challenging, preventing them from going on such trips. Natalie Friedman and Alex Cabral (2018) conducted a study with six participants with developmental difficulties to see whether providing them each with a telepresence robot would increase their physical and social self-efficacy and well-being. The participants were taken on a remote tour of two museum exhibits and then asked to rate their experience afterward. Their responses were positive, suggesting that this kind of telepresence can open doors to social experiences that were previously denied to those with disabilities.

BOX 5.1

Facebook Spaces: How Natural Is It to Socialize in a 3D World?

Facebook’s vision of social networking is to immerse people in 3D, where they interact with their friends in virtual worlds. Figure 5.5 shows what it might look like: Two avatars (Jack and Diane) are talking at a virtual table beside a lake and with some mountains in the background. Users experience this by wearing virtual reality (VR) headsets. The goal is to provide users with a magical feeling of presence, one where they can feel as if they are together, even



Figure 5.5 Facebook’s vision of socializing in a 3D world

Source: Facebook

(Continued)

though they are apart in the physical world. To make the experience appear more life-like, users can move their avatar's arms through controls provided by the VR OculusTouch.

While big strides have been made toward improving social presence, there is still a way to go before the look and feel of socializing with virtual avatars becomes more like the real thing. For a start, the facial expressions and skin tone of virtual avatars still appear to be cartoon-like.

Similar to the term telepresence, *social presence* refers to the feeling of being there with a real person in virtual reality. Specifically, it refers to the degree of awareness, feeling, and reaction to other people who are virtually present in an online environment. The term differs from *telepresence*, which refers to one party being virtually present with another party who is present in a physical space, such as a meeting room (note that it is possible for more than one telepresence robot to be in the same physical space). Imagine if avatars become more convincing in their appearance to users. How many people would switch from their current use of 2D media to catch up and chat with friends in this kind of immersive 3D Facebook page? Do you think it would enhance the experience of how they would interact and communicate with others remotely?

How many people would don a VR headset 10 times a day or more to teleport to meet their friends virtually? (The average number of times that someone looks at Facebook on their phones is now 14 times each day.) There is also the perennial problem of motion sickness that 25–40 percent of people say that they have experienced in VR (Mason, 2017). ■

Telepresence robots have also become a regular feature at conferences, including the ACM CHI conference, enabling people to attend who cannot travel. They are typically about 5-feet tall, have a display at the top that shows the remote person's head, and have a base at the bottom holding wheels allowing the robot to move forward, move backward, or turn around. A commercial example is the Beam+ (<https://suitabletech.com/>). To help the robot navigate in its surroundings, two cameras are embedded in the display, one facing outward to provide the remote person with a view of what is in front of them and the other facing downward to provide a view of the floor. The robots also contain a microphone and speakers to enable the remote person to be heard and to hear what is being said locally. Remote users connect via Wi-Fi to the remote site and steer their Beam+ robot using a web interface.

A PhD student from the University College London (UCL) attended her first CHI conference remotely, during which time she gave a demo of her research every day by talking to the attendees using the Beam+ robot (see Figure 5.6). Aside from a time difference of eight hours (meaning that she had to stay up through the night to attend), it was an enriching experience for her. She met lots of new people who not only were interested in her demo but who also learned how she felt about attending the conference remotely. Her colleagues at the conference also dressed up her robot to make her appear more like her, giving the robot a set of foam-cutout arms with waving hands, and they put a university T-shirt on the robot. However, she could not see how she appeared to others at the conference, so local attendees took photos of her Beam+ robot to show her how she looked. She also commented how she could not gauge the volume of her voice, and on one occasion she accidentally set the volume control to be at its highest setting. When speaking to someone, she did not realize how loud



Figure 5.6 Susan Lechelt's Beam+ robot given a human touch with cut-out foam arms and a university logo T-shirt

Source: Used courtesy of Susan Lechelt

she was until another person across the room told her that she was yelling. (The person she was talking to was too polite to tell her to lower her voice.)

Another navigation problem that can occur is when the remote person wants to move from one floor to another in a building. They don't have a way of pressing the elevator button to achieve this. Instead, they have to wait patiently beside the elevator for someone to come along to help them. They also lack awareness of others who are around them. For example, when moving into a room to get a good spot to see a presentation, they may not realize that they have obscured the view of people sitting behind them. It can also be a bit surreal when their image starts breaking up on the robot "face" as the Wi-Fi signal deteriorates. For example, Figure 5.7 shows Johannes Schöning breaking up into a series of pixels that makes him look a bit like David Bowie!

Despite these usability problems, a study of remote users trying a telepresence robot for the first time at a conference found the experience to be positive (Neustaedter et al., 2016). Many felt that it provided them with a real sense of being at the conference—quite different from the experience of watching or listening to talks online—as happens when connecting via a livestream or a webinar. Being able to move around the venue also enabled them to see familiar faces and to bump into people during coffee breaks. For the

conference attendees, the response was also largely positive, enabling them to chat with those who could not make the conference. However, sometimes the robot's physical presence obstructed their view in a room when watching a speaker, and that could be frustrating. It is difficult to know how to tell a telepresence robot discreetly to move out of the way while a talk is in progress and for the remote person to know where to move that is out of the way as they have been told.




Figure 5.7 The image of Johannes Schöning breaking up on the Beam+ robot video display when the Wi-Fi signal deteriorated

Source: Yvonne Rogers

ACTIVITY 5.4

Watch these two videos about Beam and Cisco's telepresence. How does the experience of being at a meeting using a telepresence robot compare with using a telepresence videoconferencing system?

Videos

BeamPro overview of how robotic telepresence works—<https://youtu.be/SQCigphfSvc>

Cisco TelePresence Room EndPoints MX700 and MX800—<https://youtu.be/52lgl0kh0FI>

Comment

The BeamPro allows the remote person to move around a workplace as well as sit in on meetings. They can also have one-on-one conversations with someone at their desk. When moving

around, the remote individual can even bump into other remote workers, in the corridor, for example, who are also using a BeamPro. Hence, it supports a range of informal and formal social interactions. Using a BeamPro also allows someone to feel as if they are at work while still being at home.

In contrast, the Cisco telepresence room has been designed specifically to support meetings between small remote groups to make them feel more natural. When someone is speaking, the camera zooms in on them to have them fill the screen. From the video, it appears effortless and allows the remote groups to focus on their meeting rather than worry about the technology. However, there is limited flexibility—conducting one-on-one meetings, for example. ■

BOX 5.2

Simulating Human Mirroring Through Artificial Smiles

A common phenomenon that occurs during face-to-face conversations is mirroring, where people mimic each other's facial expressions, gestures, or body movements. Have you ever noticed that when you put your hands behind your head, yawn, or rub your face during a conversation with someone that they follow suit? These kinds of mimicking behaviors are assumed to induce empathy and closeness between those conversing (Stel and Vonk, 2010). The more people engage in mimicry, the more they view each other as being similar, which in turn increases the rapport between them (Valdesolo and DeSteno, 2011). Mimicry doesn't always occur during a conversation, however—sometimes it requires a conscious effort, while in other situations it does not occur. Might the use of technology increase its occurrence in conversations?

One way would be to use special video effects. Suppose that an artificial smile could be superimposed on the video face of someone to make them appear to smile. What might happen? Would they both begin to smile and in doing so feel closer to each other? To investigate this possibility of simulating smiling mimicry, Keita Suzuki et al. (2017) developed a technique called FaceShare. The system was developed so that it could deform the image of someone's face to make it appear to smile—even though they were not—whenever their partner's face began smiling. The mimicry method used 3D modeling of key feature points of the face, including the contours, eyes, nose, and mouth to detect where to place the smile. The smile was created by raising the lower eyelids and both ends of the mouth in conjunction with the cheeks. The findings from this research showed that FaceShare was effective at making conversations appear smoother and that the pseudo smiles appearing on someone's video face were judged to be natural. ■

5.5 Co-presence

Together with telepresence, there has been much interest in enhancing *co-presence*, that is, supporting people in their activities when interacting in the same physical space. A number of technologies have been developed to enable more than one person to use them at the same time. The motivation is to enable co-located groups to collaborate more effectively when working, learning, and socializing. Examples of commercial products that support this kind of parallel interaction are Smartboards and Surfaces, which use multitouch, and Kinect, which uses gesture and object recognition. To understand how effective they are, it is important to consider the coordination and awareness mechanisms already in use by people in face-to-face interactions and then to see how these have been adapted or replaced by the technology.

5.5.1 Physical Coordination

When people are working closely together, they talk to each other, issuing commands and letting others know how they are progressing. For example, when two or more people are collaborating, as when moving a piano, they shout instructions to each other, like “Down a bit, left a touch, now go straight forward,” to coordinate their actions. A lot of nonverbal communication is also used, including nods, shakes, winks, glances, and hand-raising in combination with such coordination talk in order to emphasize and sometimes replace it.

For time-critical and routinized collaborative activities, especially where it is difficult to hear others because of the physical conditions, people frequently use gestures (although radio-controlled communication systems may also be used). Various types of hand signals have evolved, with their own set of standardized syntax and semantics. For example, the arm and baton movements of a conductor coordinate the different players in an orchestra, while the arm and orange baton movements of ground personnel at an airport signal to a pilot how to bring the plane into its assigned gate. Universal gestures, such as beckoning, waving, and halting hand movement, are also used by people in their everyday settings.

The use of physical objects, such as wands and batons, can also facilitate coordination. Group members can use them as external thinking props to explain a principle, an idea, or a plan to the others (Brereton and McGarry, 2000). In particular, the act of waving or holding up a physical object in front of others is very effective at commanding attention. The persistence and ability to manipulate physical artifacts may also result in more options being explored in a group setting (Fernaes and Tholander, 2006). They can help collaborators gain a better overview of the group activity and increase awareness of others’ activities.

5.5.2 Awareness

Awareness involves knowing who is around, what is happening, and who is talking with whom (Dourish and Bly, 1992). For example, when attending a party, people move around the physical space, observing what is going on and who is talking to whom, eavesdropping on others’ conversations, and passing on gossip to others. A specific kind of awareness is *peripheral awareness*. This refers to a person’s ability to maintain and constantly update a sense of what is going on in the physical and social context, by keeping an eye on what is happening in the periphery of their vision. This might include noticing whether people are in a good or bad mood by the way they are talking, how fast the drink and food is being consumed, who has entered or left the room, how long someone has been absent, and whether the lonely person in

the corner is finally talking to someone—all while we are having a conversation with someone else. The combination of direct observations and peripheral monitoring keeps people informed and updated on what is happening in the world.

Another form of awareness that has been studied is *situational awareness*. This refers to being aware of what is happening around you in order to understand how information, events, and your own actions will affect ongoing and future events. Having good situational awareness is critical in technology-rich work domains, such as air traffic control or an operating theater, where it is necessary to keep abreast of complex and continuously changing information.

People who work closely together also develop various strategies for coordinating their work, based on an up-to-date awareness of what the others are doing. This is especially so for interdependent tasks, where the outcome of one person's activity is needed for others to be able to carry out their tasks. For example, when putting on a show, the performers will constantly monitor what each other is doing in order to coordinate their performance efficiently. The metaphorical expression *close-knit teams* exemplifies this way of collaborating. People become highly skilled in reading and tracking what others are doing and the information to which they are paying attention.

A classic study of this phenomenon is of two controllers working together in a control room in the London Underground subway system (Heath and Luff, 1992). An overriding observation was that the actions of one controller were tied closely to what the other was doing. One of the controllers (controller A) was responsible for the movement of trains on the line, while the other (controller B) was responsible for providing information to passengers about the current service. In many instances, it was found that controller B overheard what controller A was saying and doing and acted accordingly, even though controller A had not said anything explicitly to him. For example, on overhearing controller A discussing a problem with a train driver over the in-cab intercom system, controller B inferred from the conversation that there was going to be a disruption in the service and so started to announce this to the passengers on the platform before controller A had even finished talking with the train driver. At other times, the two controllers kept a lookout for each other, monitoring the environment for actions and events that they might not have noticed but that could have been important for them to know about so that they could act appropriately.

ACTIVITY 5.5

What do you think happens when one person in a close-knit team does not see or hear something, or misunderstands what has been said, while the others in the group assume that person has seen, heard, or understood what has been said?

Comment

The person who has noticed that someone has not acted in the manner expected may use one of a number of subtle repair mechanisms, say coughing or glancing at something that needs to be attended to. If this doesn't work, they may then resort to stating explicitly aloud what

(Continued)

had previously been signaled implicitly. Conversely, the unaware person may wonder why the event hasn't happened and, likewise, look over at the other team members, cough to get their attention, or explicitly ask them a question. The kind of repair mechanism employed at a given moment will depend on a number of factors, including the relationship among the participants, for instance, whether one is more senior than the others. This determines who can ask what, the perceived fault or responsibility for the breakdown, and the severity of the outcome of not acting there and then on the new information. ■

5.5.3 Shareable Interfaces

A number of technologies have been designed to capitalize on existing forms of coordination and awareness mechanisms. These include whiteboards, large touch screens, and multitouch tables that enable groups of people to collaborate while interacting at the same time with content on the surfaces. Several studies have investigated whether different arrangements of shared technologies can help co-located people work better together (for example, see Müller-Tomfelde, 2010). An assumption is that shareable interfaces provide more opportunities for flexible kinds of collaboration compared with single-user interfaces, through enabling co-located users to interact simultaneously with digital content. The use of fingers or pens as input on a public display is observable by others, increasing opportunities for building situational and peripheral awareness. The sharable surfaces are also considered to be more natural than other technologies, enticing people to touch them without feeling intimidated or embarrassed by the consequences of their actions. For example, small groups found it more comfortable working together around a tabletop compared with sitting in front of a PC or standing in a line in front of a vertical display (Rogers and Lindley, 2004).

BOX 5.3

Playing Together in the Same Place

Augmented reality (AR) sandboxes have been developed for museum visitors to interact with a landscape, consisting of mountains, valleys, and rivers. The sand is real, while the landscape is virtual. Visitors can sculpt the sand into different-shaped contours that change their appearance to look like a river or land, depending on the height of the sand piles. Figure 5.8 shows a AR sandbox that was installed at the V&A museum in London. On observing two young children playing at the sandbox, this author overheard one say to the other while flattening a pile of sand, "Let's turn this land into sea." The other replied "OK, but let's make an island on that." They continued to talk about how and why they should change their landscape. It was a pleasure to watch this dovetailing of explaining and doing.

The physical properties of the sand, together with the real-time changing superimposed landscape, provided a space for children (and adults) to collaborate in creative ways. ■




Figure 5.8 Visitors creating together using an Augmented Reality Sandbox at the V&A Museum in London

Source: Helen Sharp

Often in meetings, some people dominate while others say very little. While this is OK in certain settings, in others it is considered more desirable for everyone to have a say. Is it possible to design shareable technologies so that people can participate around them more equally? Much research has been conducted to investigate whether this is possible. Of primary importance is whether the interface invites people to select, add, manipulate, or remove digital content from the displays and devices. A user study showed that a tabletop that allowed group members to add digital content by using physical tokens resulted in more equitable participation than if only digital input was allowed via touching icons and menus at the tabletop (Rogers et al., 2009). This suggests that it was easier for people who are normally shy in groups to contribute to the task. Moreover, people who spoke the least were found to make the largest contribution to the design task at the tabletop, in terms of selecting, adding, moving, and removing options. This reveals how changing the way people can interact with a surface can affect group participation. It shows that it is possible for more reticent members to contribute without feeling under pressure to speak more.

Experimentation with real-time feedback presented via ambient displays has also been shown to provide a new form of awareness for co-located groups. LEDs glowing in tabletops and abstract visualizations on handheld and wall displays have been designed to represent how different group members are performing, such as turn-taking. The assumption is that this kind of real-time feedback can promote self and group regulation and in so doing modify group members' contributions to make them more equitable. For example, the Reflect Table was designed based on this assumption (Bachour et al., 2008). The table monitors and analyzes ongoing conversations using embedded microphones in front of each person and represents this in the form of increasing numbers of colored LEDs (see Figure 5.9). A study investigated whether students became more aware of how much they were speaking during a group meeting when their relative levels of talk were displayed in this manner and, if so, whether they regulated their levels of participation more effectively. In other words, would the girl in the bottom right reduce her contributions (as she clearly has been talking the most) while the boy in the bottom left increase his (as he has been talking the least)? The findings were mixed: Some participants changed their level to match the levels of others, while others became frustrated and chose simply to ignore the LEDs. Specifically, those who spoke the most changed their behavior the most (that is, reduced their level) while those who spoke the least changed theirs the least (in other words, did not increase their level). Another finding was that participants who believed that it was beneficial to contribute equally to the conversation took more notice of the LEDs and regulated their conversation level accordingly. For example, one participant said that she “refrained from talking to avoid having a lot more lights than the others” (Bachour et al., 2010). Conversely, participants who thought it was not important took less notice. How do you think you would react?



Figure 5.9 The Reflect Table

Source: Used courtesy of Pierre Dillenbourg

An implication from the various user studies on co-located collaboration around tabletops is that designing shareable interfaces to encourage more equitable participation isn't straightforward. Providing explicit real-time feedback on how much someone is speaking in a group may be a good way of showing everyone who is talking too much, but it may be intimidating for those who are talking too little. Allowing discreet and accessible ways for adding and manipulating content to an ongoing collaborative task at a shareable surface may

be more effective at encouraging greater participation from people who normally find it difficult or who are simply unable to contribute verbally to group settings (for example, those on the autistic spectrum, those who stutter, or those who are shy or are non-native speakers).

How best to represent the activity of online social networks in terms of who is taking part has also been the subject of much research. A design principle that has been influential is *social translucence* (Erickson and Kellogg, 2000). This refers to the importance of designing communication systems to enable participants and their activities to be visible to one another. This idea was very much behind the early communication tool, Babble, developed at IBM by David Smith (Erickson et al., 1999), which provided a dynamic visualization of the participants in an ongoing chat room. A large 2D circle was depicted using colored marbles on each user's monitor. Marbles inside the circle conveyed those individuals active in the current conversation. Marbles outside the circle showed users involved in other conversations. The more active a participant was in the conversation, the more the corresponding marble moved toward the center of the circle. Conversely, the less engaged a person was in the ongoing conversation, the more the marble moved toward the periphery of the circle.

Since this early work on visualizing social interactions, there have been a number of virtual spaces developed that provide awareness about what people are doing, where they are, and their availability, with the intention of helping them feel more connected. Working in remote teams can be isolating, especially if they rarely get to see their colleagues face to face. When teams are not co-located, they also miss out on in-person collaboration and valuable informal conversations that build team alignment. This is where the concept of the “online office” comes in. For example, Sococo (<https://www.sococo.com/>) is an online office platform that is bridging the gap between remote and co-located work. It uses the spatial metaphor of a floor plan of an office to show where people are situated, who is in a meeting, and who is chatting with whom. The Sococo map (see Figure 5.10) provides a



Figure 5.10 Sococo floor plan of a virtual office, showing who is where and who is meeting with whom

Source: Used courtesy of Leeann Brumby

bird's-eye view of a team's online office, giving everyone at-a-glance insight into teammates' availability and what's happening organizationally. Sococo also provides the sense of presence and virtual "movement" that you get in a physical office—anyone can pop into a room, turn on their microphone and camera, and meet with another member of their team face to face. Teams can work through projects, get feedback from management, and collaborate ad hoc in their online office regardless of physical location. This allows organizations to take advantage of the benefits of the distributed future of work while still providing a central, online office for their teams.

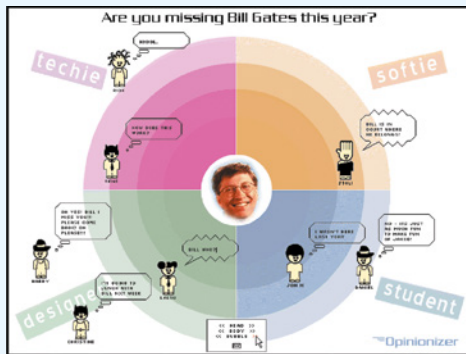
BOX 5.4

Can Technologies Be Designed to Help People Break the Ice and Socialize?

Have you ever found yourself at a party, wedding, conference, or other social gathering, standing awkwardly by yourself, not knowing who to talk to or what to talk about? Social embarrassment and self-consciousness affect most of us at such moments, and such feelings are most acute when one is a newcomer and by oneself, such as a first-time attendee at a conference. How can conversation initiation be made easier and less awkward for people who do not know each other?

A number of mechanisms have been employed by organizers of social events, such as asking old-timers to act as mentors and the holding of various kinds of ice-breaking activities. Badge-wearing, the plying of drink and food, and introductions by others are also common ploys. While many of these methods can help, engaging in ice-breaking activities requires people to act in a way that is different from the way they normally socialize and which they may find equally uncomfortable or painful to do. They often require people to agree to join in a collaborative game, which they may find embarrassing. This can be exacerbated by the fact that once people have agreed to take part, it is difficult for them to drop out because of the perceived consequences that it will have for the others and themselves (such as being seen by the others as a spoilsport or party-pooper). Having had one such embarrassing experience, most people will shy away from any further kinds of ice-breaking activities.

An alternative approach is to design a physical space where people can enter and exit a conversation with a stranger in subtler ways, that is, one where people do not feel threatened or embarrassed and that does not require a high level of commitment. The classic Opinionizer system was designed along these lines, with the goal of encouraging people in an informal gathering to share their opinions visually and anonymously (Brignull and Rogers, 2003). The collective creation of opinions via a public display was intended to provide a talking point for the people standing beside it. Users submitted their opinions by typing them in at a public keyboard. To add color and personality to their opinions, a selection of small cartoon avatars and speech bubbles were available. The screen was also divided into four labeled quadrants representing different backgrounds, such as techie, softie, designer, or student, to provide a factor on which people could comment (see Figure 5.11).



(a)



(b)

Figure 5.11 (a) The Opinionizer interface and (b) a photo of it being used at a book launch party

Source: Helen Sharp

When the Opinionizer was placed in various social gatherings, a honey-pot effect was observed. By this it is meant the creation of a sociable buzz in the immediate vicinity of the Opinionizer as a result of an increase in the number of people moving into the area. Furthermore, by standing in this space and showing an interest, for example visibly facing the screen or reading the text, people gave off a tacit signal to others that they were open to discussion and interested in meeting new people.

There are now a number of commercial ice-breaking phone apps available that use artificial intelligence (AI) matchmaking algorithms to determine which preferences and interests shared among people make them suitable conversational partners. Wearable technology is also being developed as a new form of digital ice-breaker. Limbic Media (<https://limbicmedia.ca/social-wearables/>), for example, has developed a novel pendant device colored with LED lights for this purpose. When two people touch their pendants together, the effect is for them to vibrate. This coming together action can break the ice in a fun and playful way. ■

This video features Limbic Media's novel type of social wearable being used at the 2017 BCT Tech Summit: <https://vimeo.com/216045804>.

A range of other ambient-based displays have been developed and placed in physical work settings with the purpose of encouraging people to socialize and talk more with each other. For example, the Break-Time Barometer was designed to persuade people to come out of their offices for a break to meet others they might not talk with otherwise (Kirkham et al., 2013). An ambient display, based on a clock metaphor, shows how many people are currently in the common room; if there are people present, it also sends an alert that it would be a good time to join them for a break. While the system nudged some people to go for a break in the staff room, it also had the opposite effect on others who used it to determine when breaks weren't happening so that they could take a break without their colleagues being around for company.

5.6 Social Engagement

Social engagement refers to participation in the activities of a social group (Anderson and Binstock, 2012). Often it involves some form of social exchange where people give or receive something from others. Another defining aspect is that it is voluntary and unpaid. Increasingly, different forms of social engagement are mediated by the Internet. For example, there are many websites now that support pro-social behavior by offering activities intended to help others. One of the first websites of this ilk was GoodGym (www.goodgym.org/), which connects runners with isolated older people. While out running, the runners stop for a chat with an older person who has signed up to the service, and the runner helps them with their chores. The motivation is to help others in need while getting fit. There is no obligation, and anyone is welcome to join. Another website that was set up is conservation volunteers (<https://www.tcv.org.uk/>). The website brings together those who want to help out with existing conservation activities. By bringing different people together, social cohesion is also promoted.

Not only has the Internet enabled local people to meet who would not have otherwise, it has proven to be a powerful way of connecting millions of people with a common interest in ways unimaginable before. An example is retweeting a photo that resonates with a large crowd who finds it amusing and wants to pass it on further. For example, in 2014, the most retweeted selfie was one taken by Ellen DeGeneres (an American comedian and television host) at the Oscar Academy Awards of her in front of a star-studded, smiling group of actors and friends. It was retweeted more than 2 million times (more than three-quarters of a million in the first half hour of being tweeted)—far exceeding the one taken by Barack Obama at Nelson Mandela’s funeral the previous year.

There has even been an “epic Twitter battle.” A teenager from Nevada, Carter Wilkerson, asked Wendy’s fast-food restaurant how many retweets were needed for him to receive a whole year’s supply of free chicken nuggets. The restaurant replied “18 million” (see Figure 5.12). From that moment on, his quest became viral with his tweet being retweeted more than 2 million times. Ellen’s record was suddenly put in jeopardy, and she intervened, putting out a series of requests on her show for people to continue to retweet her tweet so her record would be upheld. Carter, however, surpassed her record at the 3.5 million mark. During the Twitter battle, he used his newly found fame to create a website that sold T-shirts promoting his chicken nugget challenge. He then donated all of the proceeds from the sales toward a charity that was close to his heart. The restaurant also gave him a year’s supply of free chicken nuggets—even though he didn’t reach the target of 18 million. Not only that, it also donated \$100,000 to the same charity in honor of Carter achieving a new record. It was a win-win situation (except maybe for Ellen).

Another way that Twitter connects people rapidly and at scale is when unexpected events and disasters happen. Those who have witnessed something unusual may upload an image that they have taken of it or retweet what others have posted to inform others about it. Those who like to reach out in this way are sometimes called *digital volunteers*. For example, while writing this chapter, there was a massive thunderstorm overhead that was very dramatic. I checked out the Twitter hashtag #hove (I was in the United Kingdom) and found that hundreds of people had uploaded photos of the hailstones, flooding, and minute-by-minute updates of how public transport and traffic were being affected. It was easy to get a sense



Figure 5.12 Carter Wilkerson's tweet that went viral

of the scale of the storm before it was picked up by the official media channels, which then used some of the photos and quotes from Twitter in their coverage (see Figure 5.13). Relying on Twitter for breaking news has increasingly become the norm. When word came of a huge explosion in San Bruno, California, the chief of the Federal Emergency Management Agency in the United States logged on to Twitter and searched for the word *explosion*. Based on the tweets coming from that area, he was able to discern that the gas explosion and ensuing fire was a localized event that would not spread to other communities. He noted how he got better situational awareness more quickly from reading Twitter than by hearing about it from official sources.

Clearly, the immediacy and global reach of Twitter provides an effective form of communication, providing first responders and those living in the affected areas with up-to-the-minute information about how a wildfire, storm, or gas plume is spreading. However, the reliability of the tweeted information can sometimes be a problem. For example, some people end up obsessively checking and posting, sometimes without realizing that this can start or fuel rumors by adding news that is old or incorrect. Regulars can go into a frenzy, constantly adding new tweets about an event, as witnessed when an impending flood was announced (Starbird et al., 2010). While such citizen-led dissemination and retweeting of information from disparate sources is well intentioned, it can also flood the Twitter streams, making it difficult to know what is old, actual, or hearsay.



Figure 5.13 A weather warning photo tweeted and retweeted about a severe storm in Hove, United Kingdom

BOX 5.5

Leveraging Citizen Science and Engagement Through Technology

The growth and success of citizen science and citizen engagement has been made possible by the Internet and mobile technology, galvanizing and coordinating the efforts of millions of people throughout the world. Websites, smartphone apps, and social media have been instrumental in leveraging the reach and impact of a diversity of citizen science projects across time and geographical zones (Preece et al., 2018). Citizen science involves local people helping scientists carry out a scientific project at scale. Currently, thousands of such projects have been set up all over the world, whereby volunteers help out in a number of research areas, including biodiversity, air quality, astronomy, and environmental issues. They do so by engaging in scientific activities such as monitoring plants and wildlife, collecting air and water samples, categorizing galaxies, and analyzing DNA sequences. Citizen engagement involves people helping governments, rather than scientists, to improve public services and policies in their communities. Examples include setting up and overseeing a website that offers local services for community disasters and creating an emergency response team when a disaster occurs.

Why would anyone want to volunteer their time for the benefit of science or government? Many people want to learn more about a domain, while others want to be recognized for their contributions (Rotman et al., 2014). Some citizen science apps have developed online mechanisms to support this. For example, iNaturalist (<https://www.inaturalist.org/>) enables volunteers to comment on and help classify others' contributions. ■

DILEMMA

Is It OK to Talk with a Dead Person Using a Chatbot?

Eugenia Kuyda, an AI researcher, lost a close friend in a car accident. He was only in his 20s. She did not want to lose his memory, so she gathered all of the texts he had sent over the course of his life and made a chatbot from them. The chatbot is programmed to respond automatically to text messages so that Eugenia can talk to her friend as if he were still alive. It responds to her questions using his own words.

Do you think this kind of interaction is creepy or comforting to someone who is grieving? Is it disrespectful of the dead, especially if the dead person has not given their consent? What if the friend had agreed to having their texts mashed up in this way in a “pre-death digital agreement”? Would that be more socially acceptable? ■

In-Depth Activity

The goal of this activity is to analyze how collaboration, coordination, and communication are supported in online video games involving multiple players.

The video game *Fortnite* arrived in 2017 to much acclaim. It is an action game designed to encourage teamwork, cooperation, and communication. Download the game from an app store (it is free) and try it. You can also watch an introductory video about it at https://youtu.be/_U2JbFhUPX8.

Answer the following questions.

1. Social issues
 - (a) What is the goal of the game?
 - (b) What kinds of conversations are supported?
 - (c) How is awareness of the others in the game supported?
 - (d) What kinds of social protocols and conventions are used?
 - (e) What types of awareness information are provided?
 - (f) Does the mode of communication and interaction seem natural or awkward?
 - (g) How do players coordinate their actions in the game?
2. Interaction design issues
 - (a) What form of interaction and communication is supported, for instance, text, audio, and/or video?
 - (b) What other visualizations are included? What information do they convey?
 - (c) How do users switch between different modes of interaction, for example, exploring and chatting? Is the switch seamless?
 - (d) Are there any social phenomena that occur specific to the context of the game that wouldn't happen in face-to-face settings?
3. Design issues
 - What other features might you include in the game to improve communication, coordination, and collaboration?

Summary

Human beings are inherently social. People will always need to collaborate, coordinate, and communicate with one another, and the diverse range of applications, web-based services, and technologies that have emerged enable them to do so in more extensive and diverse ways. In this chapter, we looked at some core aspects of sociality, namely, communication and collaboration. We examined the main social mechanisms that people use in different conversational settings when interacting face to face and at a distance. A number of collaborative and telepresence technologies designed to support and extend these mechanisms were discussed, highlighting core interaction design concerns.

Key Points

- Social interaction is central to our everyday lives.
- Social mechanisms have evolved in face-to-face and remote contexts to facilitate conversation, coordination, and awareness.
- Talk and the way it is managed are integral to coordinating social interaction.
- Many kinds of technologies have been developed to enable people to communicate remotely with one another.
- Keeping aware of what others are doing and letting others know what you are doing are important aspects of collaboration and socializing.
- Social media has brought about significant changes in the way people keep in touch and manage their social lives.

Further Reading

boyd, d. (2014) *It's Complicated: The Social Lives of Networked Teens*. Yale. Based on a series of in-depth interviews with a number of teenagers, danah boyd offers new insights into how teenagers across the United States, who have only ever grown up in a world of apps and media, navigate, use, and appropriate them to grow up and develop their identities. A number of topics are covered that are central to what it means to grow up in a networked world, including bullying, addiction, expressiveness, privacy, and inequality. It is insightful and covers much ground.

CRUMLISH, C. and MALONE, E. (2009) *Designing Social Interfaces*. O'Reilly. This is a collection of design patterns, principles, and advice for designing social websites, such as online communities.

GARDNER, H. and DAVIS, K. (2013) *The App Generation: How Today's Youth Navigate Identity, Intimacy, and Imagination in a Digital World*. Yale. This book explores the impact of apps on the young generation, examining how they affect their identity, intimacy, and imagination. It focuses on what it means to be app-dependent versus app-empowered.

ROBINSON, S., MARSDEN, G. and JONES, M. (2015) *There's Not an App for That: Mobile User Experience Design for Life*. Elsevier. This book offers a fresh approach for designers, students, and researchers to dare to think differently by moving away from the default framing of technological design in terms of yet another “looking down” app. It asks the reader instead to look up and around them—to be inspired by how we actually live our lives when “out there” app-less. They also explore what it means to design technologies to be more mindful.

TURKLE, S. (2016) *Reclaiming Conversation: The Power of Talk in a Digital Age*. Penguin. Sherry Turkle has written extensively about the positive and negative effects of digital technology on everyday lives—at work, at home, at school, and in relationships. This book is a very persuasive warning about the negative impacts of perpetual use of smartphones. Her main premise is that as people—both adults and children—become increasingly glued to their phones instead of talking to one another, they lose the skill of empathy. She argues that we need to reclaim conversation to relearn empathy, friendship, and creativity.

Chapter 6

EMOTIONAL INTERACTION

- 6.1 Introduction
- 6.2 Emotions and the User Experience
- 6.3 Expressive Interfaces and Emotional Design
- 6.4 Annoying Interfaces
- 6.5 Affective Computing and Emotional AI
- 6.6 Persuasive Technologies and Behavioral Change
- 6.7 Anthropomorphism

Objectives

The main goals of this chapter are to accomplish the following:

- Explain how our emotions relate to behavior and the user experience.
- Explain what are expressive and annoying interfaces and the effects they can have on people.
- Introduce the area of emotion recognition and how it is used.
- Describe how technologies can be designed to change people's behavior.
- Provide an overview on how anthropomorphism has been applied in interaction design.

6.1 Introduction

When you receive some bad news, how does it affect you? Do you feel upset, sad, angry, or annoyed—or all of these? Does it put you in a bad mood for the rest of the day? How might technology help? Imagine a wearable technology that could detect how you were feeling and provide a certain kind of information and suggestions geared toward helping to improve your mood, especially if it detected that you were having a real downer of a day. Would you find such a device helpful, or would you find it unnerving that a machine was trying to cheer you up? Designing technology to detect and recognize someone's emotions automatically from sensing aspects of their facial expressions, body movements, gestures, and so forth,

is a growing area of research often called *emotional AI* or *affective computing*. There are many potential applications for using automatic emotion sensing, other than those intended to cheer someone up, including health, retail, driving, and education. These can be used to determine if someone is happy, angry, bored, frustrated, and so on, in order to trigger an appropriate technology intervention, such as making a suggestion to them to stop and reflect or recommending a particular activity for them to do.

In addition, *emotional design* is a growing area relating to the design of technology that can engender desired emotional states, for example, apps that enable people to reflect on their emotions, moods, and feelings. The focus is on how to design interactive products to evoke certain kinds of emotional responses in people. It also examines why people become emotionally attached to certain products (for instance, virtual pets), how social robots might help reduce loneliness, and how to change human behavior through the use of emotive feedback.

In this chapter, we include emotional design and affective computing using the broader term, *emotional interaction*, to cover both aspects. We begin by explaining what emotions are and how they shape behavior and everyday experiences. We then consider how and whether an interface's appearance affects usability and the user experience. In particular, we look at how expressive and persuasive interfaces can change people's emotions or behaviors. How technology can detect human emotions using voice and facial recognition is then covered. Finally, the way anthropomorphism has been used in interaction design is discussed.

6.2 Emotions and the User Experience

Consider the different emotions one experiences throughout a common everyday activity—shopping online for a product, such as a new laptop, a sofa, or a vacation. First, there is the realization of needing or wanting one and then the desire and anticipation of purchasing it. This is followed by the joy or frustration of finding out more about what products are available and deciding which to choose from potentially hundreds or even thousands of them by visiting numerous websites, such as comparison sites, reviews, recommendations, and social media sites. This entails matching what is available with what you like or need and whether you can afford it. The thrill of deciding on a purchase may be quickly followed by the shock of how much it costs and the disappointment that it is too expensive. The process of having to revise your decision may be accompanied by annoyance if you discover that nothing is as good as the first choice. It can become frustrating to keep looking and revisiting sites. Finally, when you make your decision, a sense of relief is often experienced. Then there is the process of clicking through the various options (such as color, size, warranty, and so forth) until the online payment form pops up. This can be tedious, and the requirement to fill in the many details raises the possibility of making a mistake. Finally, when the order is complete, you can let out a big sigh. However, doubts can start to creep in—maybe the other one was better after all... .

This rollercoaster set of emotions is what many of us experience when shopping online, especially for big-ticket items where there is a myriad of options from which to choose and where you want to be sure that you make the right choice.

ACTIVITY 6.1

Have you seen one of the terminals shown in Figure 6.1 at an airport after you have gone through security? Were you drawn toward it, and did you respond? If so, which smiley button did you press?



Figure 6.1 A Happyornot terminal located after security at Heathrow Airport

Source: <https://www.rsresearch.com/research/why-metrics-matter>. Used courtesy of Retail Systems Research

Comment

The act of pressing one of the buttons can be very satisfying—providing a moment for you to reflect upon your experience. It can even be pleasurable to express how you feel in this physical manner. Happyornot designed the feedback terminals that are now used in many airports throughout the world. The affordances of the large, colorful, slightly raised buttons laid out in a semicircle, with distinct smileys, makes it easy to know what is being asked of the passerby, enabling them to select among feeling happy, angry, or something in between.

The data collected from the button presses provides statistics for an airport as to when and where people are happiest and angriest after going through security. Happyornot has found that it also makes travelers feel valued. The happiest times to travel, from the data they have collected at various airports, are at 8 a.m. and 9 a.m. The unhappiest times recorded are in the early hours of the morning, presumably because people are tired and grumpier. ■

Emotional interaction is concerned with what makes people feel happy, sad, annoyed, anxious, frustrated, motivated, delirious, and so on, and then using this knowledge to inform the design of different aspects of the user experience. However, it is not straightforward. Should an interface be designed to try to keep a person happy when it detects that they are smiling, or should it try to change them from being in a negative mood to a positive one when it detects that they are scowling? Having detected an emotional state, a decision has to

be made as to what or how to present information to the user. Should it try to “smile” back through using various interface elements, such as emojis, feedback, and icons? How expressive should it be? It depends on whether a given emotional state is viewed as desirable for the user experience or the task at hand. A happy state of mind might be considered optimal for when someone goes to shop online if it is assumed that this will make them more willing to make a purchase.

Advertising agencies have developed a number of techniques to influence people’s emotions. Examples include showing a picture of a cute animal or a child with hungry, big eyes on a website that “pulls at the heartstrings.” The goal is to make people feel sad or upset at what they observe and make them want to do something to help, such as by making a donation. Figure 6.2, for example, shows a web page that has been designed to trigger a strong emotional response in the viewer.



Figure 6.2 A webpage from Crisis (a UK homelessness charity)

Source: <https://www.crisis.org.uk>

Our moods and feelings are also continuously changing, making it more difficult to predict how we feel at different times. Sometimes, an emotion can descend upon us but disappear shortly afterward. For example, we can become startled by a sudden, unexpected loud noise. At other times, an emotion can stay with us for a long time; for example, we can remain annoyed for hours when staying in a hotel room that has a noisy air conditioning unit. An emotion like jealousy can keep simmering for a long period of time, manifesting itself on seeing or hearing something about the person or thing that triggered it.

In a series of short videos, Kia Höök talks about affective computing, explaining how emotion is formed and why it is important to consider when designing user experiences with technology. See www.interaction-design.org/encyclopedia/affective_computing.html.

A good place to start understanding how emotions affect behavior and how behavior affects emotions is to examine how people express themselves and read each other's expressions. This includes understanding the relationship between facial expressions, body language, gestures, and tone of voice. For example, when people are happy, they typically smile, laugh, and relax their body posture. When they are angry, they might shout, gesticulate, tense their hands, and screw up their face. A person's expressions can trigger emotional responses in others. When someone smiles, it can cause others to feel good and smile back.

Emotional skills, especially the ability to express and recognize emotions, are central to human communication. Most people are highly skilled at detecting when someone is angry, happy, sad, or bored by recognizing their facial expressions, way of speaking, and other body signals. They also usually know what emotions to express in a given situation. For example, when someone has just heard they have failed an exam, it is not a good time to smile and be happy for them. Instead, people try to empathize and show that they feel sad, too.

There is an ongoing debate about whether and how emotion causes certain behaviors. For example, does being angry make us concentrate better? Or does being happy make us take more risks, such as spending too much money or vice versa or neither? It could be that we can just feel happy, sad, or angry, and that this does not affect our behavior. Roy Baumeister et al. (2007) argue that the role of emotion is more complicated than a simple cause-and-effect model.

Many theorists, however, argue that emotions cause behavior, for example that fear brings about flight and that anger initiates the fight perspective. A widely accepted explanation, derived from evolutionary psychology, is that when something makes someone frightened or angry, their emotional response is to focus on the problem at hand and try to overcome or resolve the perceived danger. The physiological responses that accompany this state usually include a rush of adrenalin through the body and the tensing of muscles. While the physiological changes prepare people to fight or flee, they also give rise to unpleasant experiences, such as sweating, butterflies in the stomach, quick breathing, heart pounding, and even feelings of nausea.

Nervousness is a state of being that is often accompanied by several emotions, including apprehension and fear. For example, many people get worried and some feel terrified before speaking at a public event or a live performance. There is even a name for this kind of nervousness—*stage fright*. Andreas Komninos (2017) suggests that it is the autonomous system “telling” people to avoid these kinds of potentially humiliating or embarrassing experiences. But performers or professors can't simply run away. They have to cope with the negative emotions associated with having to be in front of an audience. Some are able to turn their nervous state to their advantage, using the increase in adrenalin to help them focus on their performance. Others are only too glad when the performance is over and they can relax again.

As mentioned earlier, emotions can be simple and short-lived or complex and long-lasting. To distinguish between the two types of emotion, researchers have described them in terms of being either automatic or conscious. *Automatic emotions* (also known as *affect*) happen rapidly, typically within a fraction of a second and, likewise, may dissipate just as quickly. *Conscious emotions*, on the other hand, tend to be slow to develop and equally slow to dissipate, and they are often the result of a conscious cognitive behavior, such as weighing the odds, reflection, or contemplation.

BOX 6.1

How Does Emotion Affect Driving Behavior?

Research investigating the influence of emotions on driving behavior has been extensively reviewed (Pêcher et al., 2011; Zhang and Chan, 2016). One major finding is that when drivers are angry, their driving becomes more aggressive, they take more risks such as dangerous overtaking, and they are prone to making more errors. Driving performance has also been found to be negatively affected when drivers are anxious. People who are depressed are also more prone to accidents.

What are the effects of listening to music while driving? A study by Christelle Pêcher et al. (2009) found that people slowed down while driving in a car simulator when they listened to either happy or sad music, as compared to neutral music. This effect is thought to be due to the drivers focusing their attention on the emotions and lyrics of the music. Listening to happy music was also found not only to slow drivers down, but to distract them more by reducing their ability to stay in their lane. This did not happen with the sad music. ■



“it’s a very user-friendly model.”

Source: Jonny Hawkins / Cartoon Stock

Understanding how emotions work provides a way of considering how to design for user experiences that can trigger affect or reflection. For example, Don Norman (2005) suggests that being in a positive state of mind can enable people to be more creative as they are less focused. When someone is in a good mood, it is thought to help them make decisions more quickly. He also suggests that when people are happy, they are more likely to overlook and cope with minor problems that they are experiencing with a device or interface. In contrast, when someone is anxious or angry, they are more likely to be less tolerant. He also suggests that designers pay special attention to the information required to do the task at hand, but especially in the case when designing apps or devices for serious tasks, such as monitoring a process control plant or driving a car. The interface needs to be clearly visible with

unambiguous feedback. The bottom line is “things intended to be used under stressful situations require a lot more care, with much more attention to detail” (Norman, 2005, p. 26).

Don Norman and his colleagues (Ortony et al., 2005) have also developed a model of emotion and behavior. It is couched in terms of different “levels” of the brain. At the lowest level are parts of the brain that are prewired to respond automatically to events happening in the physical world. This is called the *visceral level*. At the next level are the brain processes that control everyday behavior. This is called the *behavioral level*. At the highest level are brain processes involved in contemplating. This is called the *reflective level* (see Figure 6.3). The visceral level responds rapidly, making judgments about what is good or bad, safe or dangerous, pleasurable or abhorrent. It also triggers the emotional responses to stimuli (for instance fear, joy, anger, and sadness) that are expressed through a combination of physiological and behavioral responses. For example, many people will experience fear on seeing a very large hairy spider running across the floor of the bathroom, causing them to scream and run away. The behavioral level is where most human activities occur. Examples include well-learned routine operations such as talking, typing, and swimming. The reflective level entails conscious thought where people generalize across events or step back from their daily routines. An example is switching between thinking about the narrative structure and special effects used in a horror movie and becoming scared at the visceral level when watching the movie.

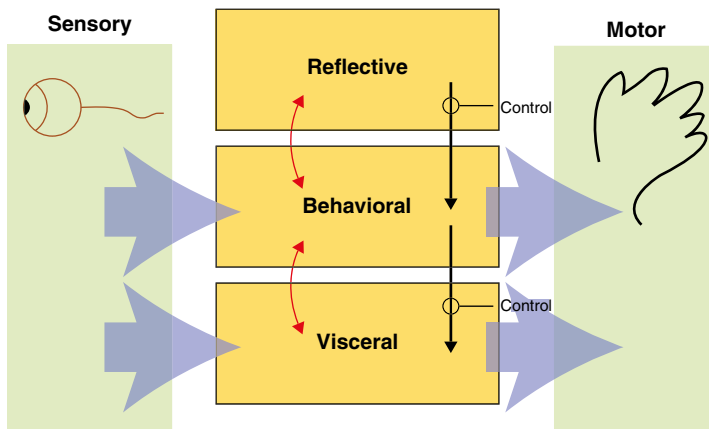


Figure 6.3 Anthony Ortony et al.’s (2005) model of emotional design showing three levels: visceral, behavioral, and reflective

Source: Adapted from Norman (2005), Figure 1.1

One way of using the model is to think about how to design products in terms of the three levels. Visceral design refers to making products look, feel, and sound good. Behavioral design is about use and equates to the traditional values of usability. Reflective design is about considering the meaning and personal value of a product in a particular culture. For example, the design of a Swatch watch (see Figure 6.4) can be viewed in terms of the three levels. The use of cultural images and graphical elements is designed to appeal to users



Figure 6.4 A Swatch watch called Dip in Color

Source: <http://store.swatch.com/suop103-dip-in-color.html>

at the reflective level; its affordances of use at the behavioral level, and the brilliant colors, wild designs, and art attract users' attention at the visceral level. They are combined to create the distinctive Swatch trademark, and they are what draw people to buy and wear their watches.

6.3 Expressive Interfaces and Emotional Design

Designers use a number of features to make an interface expressive. Emojis, sounds, colors, shapes, icons, and virtual agents are used to (1) create an emotional connection or feeling with the user (for instance, warmth or sadness) and/or (2) elicit certain kinds of emotional responses in users, such as feeling at ease, comfort, and happiness. In the early days, emotional icons were used to indicate the current state of a computer or a phone, notably when it was waking up or being rebooted. A classic from the 1980s was the happy Mac icon that appeared on the screen of the Apple computer whenever the machine was booted (see Figure 6.5a). The smiling icon conveyed a sense of friendliness, inviting the user to feel at ease and even smile back. The appearance of the icon on the screen was also meant to be

reassuring, indicating that the computer was working. After being in use for nearly 20 years, the happy and sad Mac icons were laid to rest. Apple now uses more impersonal but aesthetically pleasing forms of feedback to indicate a process for which the user needs to wait, such as “starting up,” “busy,” “not working,” or “downloading.” These include a spinning colorful beach ball (see Figure 6.5b) and a moving clock indicator. Similarly, Android uses a spinning circle to show when a process is loading.




Figure 6.5 (a) Smiling and sad Apple icons depicted on the classic Mac and (b) the spinning beach ball shown when an app freezes

Source: (b) <https://www.macobserver.com/tmo/article/frozen-how-to-force-quit-an-os-x-app-showing-a-spinningbeachball-of-death>

Other ways of conveying expressivity include the following:

- Animated icons (for example, a recycle bin expanding when a file is placed in it and paper disappearing in a puff of smoke when emptied)
- Sonifications indicating actions and events (such as whoosh for a window closing, “schlook” for a file being dragged, or ding for a new email arriving)
- Vibrotactile feedback, such as distinct smartphone buzzes that represent specific messages from friends or family

The style or brand conveyed by an interface, in terms of the shapes, fonts, colors, and graphical elements used, and the way they are combined, also influence its emotional impact. Use of imagery at the interface can result in more engaging and enjoyable experiences (Mullet and Sano, 1995). A designer can also use a number of aesthetic techniques such as clean lines, balance, simplicity, white space, and texture.

The benefits of having aesthetically pleasing interfaces have become more acknowledged in interaction design. Noam Tractinsky (2013) has repeatedly shown how the aesthetics of an interface can have a positive effect on people’s perception of the system’s usability. When the look and feel of an interface is pleasing and pleasurable—for example through beautiful graphics or a nice feel or the way that the elements have been put together—people are likely to be more tolerant and prepared to wait a few more seconds for a website to download. Furthermore, good-looking interfaces are generally more satisfying and pleasurable to use.

BOX 6.2

The Design of the Nest Thermostat Interface

The popular Nest thermostat provides an automatic way of controlling home heating that is personalized to the habits and needs of the occupants. Where possible, it also works out how to save money by reducing energy consumption when not needed. The wall-mounted device does this by learning what temperature the occupants prefer and when to turn the heating on and off in each room by learning their routines.

The Nest thermostat is more than just a smart meter, however. It was also designed to have a minimalist and aesthetically pleasing interface (see Figure 6.6a). It elegantly shows the temperature currently on its round face and to which temperature it has been set. This is very different from earlier generations of automatic thermostats, which were utilitarian box-shaped designs with lots of complicated buttons and a dull screen that provided feedback about the setting and temperature (see Figure 6.6b). It is little wonder that the Nest thermostat has been a success. ■



Figure 6.6 (a) The Nest thermostat and (b) A traditional thermostat

Source: Nest

For more information about the design of other Nest products, see <https://www.wired.com/story/inside-the-second-coming-of-nest/>.

6.4 Annoying Interfaces

In many situations, interfaces may inadvertently elicit negative emotional responses, such as anger. This typically happens when something that should be simple to use or set turns out to be complex. The most common examples are remote controls, printers, digital alarm clocks, and digital TV systems. Getting a printer to work with a new digital camera, trying to switch from watching a DVD to a TV channel, and changing the time on a digital alarm clock in

a hotel can be very trying. Also, complex actions such as attaching the ends of cables between smartphones and laptops, or inserting a SIM card into a smartphone, can be irksome, especially if it is not easy to see which way is correct to insert them.

This does not mean that developers are unaware of such usability problems. Several methods have been devised to help the novice user get set up and become familiarized with a technology. These methods include pop-up help boxes and contextual videos. Another approach to helping users has been to make an interface appear friendlier as a way of reassuring users—especially those who were new to computers or online banking. One technique that was first popularized in the 1990s was the use of cartoon-like companions. The assumption was that novices would feel more at ease with a “friend” appearing on the screen and would be encouraged to try things out after listening, watching, following, and interacting with it. For example, Microsoft pioneered a class of agent-based software, Bob, aimed at new computer users (many of whom were viewed as computer-phobic). The agents were presented as friendly characters, including a pet dog and a cute bunny. An interface metaphor of a warm, cozy living room, replete with fire and furniture, was also provided (see Figure 6.7), again intended to convey a comfortable feeling. However, Bob never became a commercial product. Why do you think that was?



Figure 6.7 “At home with Bob” software developed for Windows 95

Source: Microsoft Corporation

Contrary to the designers’ expectations, many people did not like the idea of Bob, finding the interface too cute and childish. However, Microsoft did not give up on the idea of making its interfaces friendlier and developed other kinds of agents, including the infamous Clippy (a paper clip that had human-like qualities), as part of their Windows 98 operating environment. Clippy typically appeared at the bottom of a user’s screen whenever the system thought the user needed help carrying out a particular task (see Figure 6.8a). It, too, was depicted as a cartoon character, with a warm personality. This time, Clippy was released as a commercial product, but it was not a success. Many Microsoft users found it too intrusive, distracting them from their work.



Figure 6.8 Defunct virtual agents: (a) Microsoft's Clippy and (b) IKEA's Anna

Source: Microsoft Corporation

A number of online stores and travel agencies also began including automated virtual agents in the form of cartoon characters who acted as sales agents on their websites. The agents appeared above or next to a textbox where the user could type in their query. To make them appear as if they were listening to the user, they were animated in a semi human-like way. An example of this was Anna from IKEA (see Figure 6.8b) who occasionally nodded, blinked her eyes, and opened her mouth. These virtual agents, however, have now largely disappeared from our screens, being replaced by virtual assistants who talk in speech bubbles that have no physical appearance, or static images of real agents who the user can talk to via LiveChat.

Interfaces, if designed poorly, can make people sometimes feel insulted, stupid, or threatened. The effect can be to annoy them to the point of losing their temper. There are many situations that cause such negative emotional responses. These include the following:

- When an application doesn't work properly or crashes
- When a system doesn't do what the user wants it to do
- When a user's expectations are not met
- When a system does not provide sufficient information to let the user know what to do
- When error messages pop up that are vague or obtuse
- When the appearance of an interface is too noisy, garish, gimmicky, or patronizing

- When a system requires users to carry out too many steps to perform a task, only to discover a mistake was made somewhere along the line and they need to start all over again
- Websites that are overloaded with text and graphics, making it difficult to locate desired information and resulting in sluggish performance
- Flashing animations, especially flashing banner ads and pop-up ads that cover the user view and which require them to click in order to close them
- The overuse or automatic playing of sound effects and music, especially when selecting options, carrying out actions, running tutorials, or watching website demos
- Featuritis—an excessive number of operations, such as an array of buttons on remote controls
- Poorly laid-out keyboards, touchpads, control panels, and other input devices that cause users to press the wrong keys or buttons persistently

ACTIVITY 6.2

Most people are familiar with the “404 error” message that pops up now and again when a web page does not upload for the link they have clicked or when they have typed or pasted an incorrect URL into a browser. What does it mean and why the number 404? Is there a better way of letting users know when a link to a website is not working? Might it be better for the web browser to say that it was sorry rather than presenting an error message?

Comment

The number 404 comes from the HTML language. The first 4 indicates a client error. The server is telling the user that they have done something wrong, such as misspelling the URL or requesting a page that no longer exists. The middle 0 refers to a general syntax error, such as a spelling mistake. The last 4 indicates the specific nature of the error. For the user, however, it is an arbitrary number. It might even suggest that there are 403 other errors they could make!

Early research by Byron Reeves and Clifford Nass (1996) suggested that computers should be courteous to users in the same way that people are to one another. They found that people are more forgiving and understanding when a computer says that it’s sorry after making a mistake. A number of companies now provide alternative and more humorous “error” landing pages that are intended to make light of the embarrassing situation and to take the blame away from the user (see Figure 6.9). ■

(Continued)



Figure 6.9 An alternative 404 error message

Source: <https://www.creativebloq.com/web-design/best-404-pages-812505>

DILEMMA

Should Voice Assistants Teach Kids Good Manners?

Many families now own a smart speaker, such as an Amazon Echo, with a voice assistant like Alexa running on it. One observation is that young children will often talk to Alexa as if she was their friend, asking her all sorts of personal questions, such as “Are you my friend?” and “What is your favorite music?” and “What is your middle name?” They also learn that is not necessary to say “please” when asking their questions or “thank you” on receiving a response, similar to how they talk to other display-based voice assistants, such as Siri or Cortana. Some parents, however, are worried that this lack of etiquette could develop into a new social norm that could transfer over to how they talk to real human beings. Imagine the scenario where Aunt Emma and Uncle Liam come over to visit their young niece for her 5th birthday, and the first thing that they hear is, “Aunty Emma, get me my drink” or “Uncle Liam, where is my birthday present?” with nary a “please” uttered. How would you feel if you were treated like that?

One would hope that parents would continue to teach their children good manners and the difference between a real human and a voice assistant. However, it is also possible to configure Alexa and other voice assistants to reward children when they are polite to them,

for example, by saying “By the way, thanks for asking so nicely.” Voice assistants could also be programmed to be much more forceful in how they teach good manners, for example, saying, “I won’t answer you unless you say ‘please’ each time you ask me a question.” Would this be taking the role of parenting too far? Mike Elgon (2018) cogently argues why voice assistants should not do this. He questions whether by extending human social norms to voice assistants, we are teaching children that technology can have sensibilities and hence should be thought about in the same way that we consider human feelings. In particular, he wonders whether by being polite to a voice assistant, children might begin to think that they are capable of feeling appreciated or unappreciated and that they have rights just like humans. Do you agree with him, or do you think that there is no harm in developing virtual assistants to teach children good manners and that children will learn? Or, do you believe that children will instinctively know voice assistants don’t have rights or feelings?

6.5 Affective Computing and Emotional AI

Affective computing is concerned with how to use computers to recognize and express emotions in the same way as humans do (Picard, 1998). It involves designing ways for people to communicate their emotional states, through using novel, wearable sensors and creating new techniques to evaluate frustration, stress, and moods by analyzing people’s expressions and conversations. It also explores how affect influences personal health (Jacques et al., 2017). More recently, *emotional AI* has emerged as a research area that seeks to automate the measurement of feelings and behaviors by using AI technologies that can analyze facial expressions and voice in order to infer emotions. A number of sensing technologies can be used to achieve this and, from the data collected, predict aspects of a user’s behavior, for example, forecasting what someone is most likely to buy online when feeling sad, bored, or happy. The main techniques and technologies that have been used to do this are as follows:

- Cameras for measuring facial expressions
- Biosensors placed on fingers or palms to measure galvanic skin response (which is used to infer how anxious or nervous someone is as indicated by an increase in their sweat)
- Affective expression in speech (voice quality, intonation, pitch, loudness, and rhythm)
- Body movement and gestures, as detected by motion capture systems or accelerometer sensors placed on various parts of the body

The use of automated facial coding is gaining popularity in commercial settings, especially in marketing and e-commerce. For example, Affectdex emotion analytics software from Affectiva (www.affectiva.com) employs advanced computer vision and machine learning algorithms to catalog a user’s emotional reactions to digital content, as captured through a webcam, to analyze how engaged the user is with digital online content, such as movies, online shopping sites, and advertisements.

Six fundamental emotions are classified based on the facial expressions that Affectdex collects.

- Anger
- Contempt
- Disgust
- Fear
- Joy
- Sadness

These emotions are indicated as a percentage of what was detected beside the emotion labels above the person's face appearing on a display. For example, Figure 6.10 shows a label of 100 percent happiness and 0 percent for all the other categories above the woman's head on the smartphone display. The white dots overlaying her face are the markers used by the app when modeling a face. They provide the data that determines the type of facial expression being shown, in terms of detecting the presence or absence of the following:

- Smiling
- Eye widening
- Brow raising
- Brow furrowing
- Raising a cheek
- Mouth opening
- Upper-lip raising
- Wrinkling of the nose

If a user screws up their face when an ad pops up, this suggests that they feel disgust, whereas if they start smiling, it suggests that they are feeling happy. The website can then adapt its ad, movie storyline, or content to what it perceives the person needs at that point in their emotional state.



Figure 6.10 Facial coding using Affectiva software

Source: Affectiva, Inc.

Affectiva has also started to analyze drivers' facial expressions when on the road with the goal of improving driver safety. The emotional AI software perceives if a driver is angry and then suggests an intervention. For example, a virtual agent in the car might suggest to the driver to take a deep breath and play soothing music to help relax them. In addition to identifying particular emotions through facial expressions (for example, joy, anger, and surprise), Affectiva uses particular markers to detect drowsiness. These are eye closure, yawning, and blinking rate. Again, upon detecting when a threshold has been reached for these facial expressions, the software might trigger an action, such as getting a virtual agent to suggest to the driver that they pull over where it is safe to do so.

Other indirect methods that are used to reveal the emotional state of someone include eye-tracking, finger pulse, speech, and the words/phrases they use when tweeting, chatting online, or posting to Facebook (van den Broek, 2013). The level of affect expressed by users, the language they use, and the frequency with which they express themselves when using social media can all indicate their mental state, well-being, and aspects of their personality (for instance, whether they are an extrovert or introvert, neurotic or calm, and so on). Some companies may try to use a combination of these measures, such as facial expressions and the language that people use when online, while others may focus on just one aspect, such as the tone of their voice when answering questions over the phone. This type of indirect emotion detection is beginning to be used to help infer or predict someone's behavior, for example, determining their suitability for a job or how they will vote in an election.

Another application of biometric data is being used in streaming video games where spectators watch players, known as *streamers*, play video games. The most popular site is Twitch; millions of viewers visit it each day to watch others compete in games, such as *Fortnite*. The biggest streamers have become a new breed of celebrity, like YouTubers. Some even have millions of dedicated fans. Various tools have been developed to enhance the viewers' experience. One is called All the Feels, which provides an overlay of biometric and webcam-derived data of a streamer onto the screen interface (Robinson et al., 2017). A dashboard provides a visualization of the streamer's heart rate, skin conductance, and emotions. This additional layer of data has been found to enhance the spectator experience and improve the connection between the streamer and spectators. Figure 6.11 shows the emotional state of a streamer using the All the Feels interface.

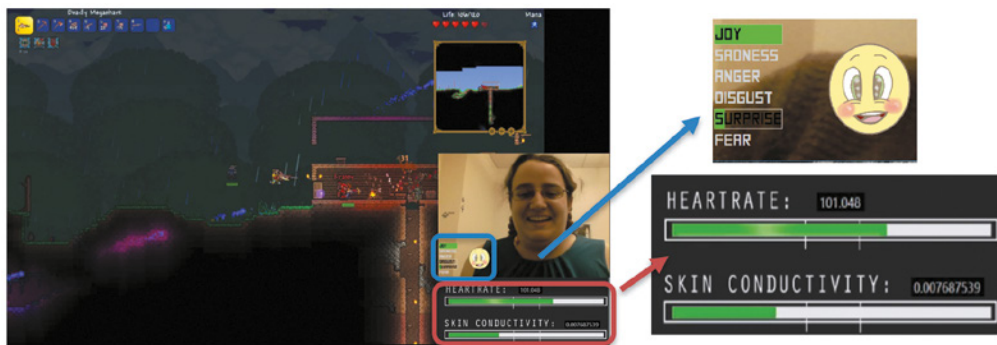


Figure 6.11 All the Feels app showing the biometric data of a streamer playing a videogame

Source: Used courtesy of Katherine Isbister

BOX 6.3

Is It OK for Technology to Work Out How You Are Feeling?

Do you think it is ethical that technology is trying to read your emotions from your facial expressions or from what you write in your tweets and, based on its analysis, filter the online content that you are browsing, such as ads, news, or a movie to match your mood? Might some people think it is an invasion of their privacy?

Human beings will suggest things to each other, often based on what they think the other is feeling. For example, they might suggest a walk in the park to cheer them up. They might also suggest a book to read or a movie to watch. However, some people may not like the idea that an app can do the same, for example, suggesting what you should eat, watch, or do based on how it analyzes your facial expressions. ■

6.6 Persuasive Technologies and Behavioral Change

A diversity of techniques has been used at the interface level to draw people's attention to certain kinds of information in an attempt to change what they do or think. Pop-up ads, warning messages, reminders, prompts, personalized messages, and recommendations are some of the methods that are being deployed on a computer or smartphone interface. Examples include Amazon's one-click mechanism that makes it easy to buy something on its online store and recommender systems that suggest specific books, hotels, restaurants, and so forth, that a reader might want to try based on their previous purchases, choices, and taste. The various techniques that have been developed have been referred to as *persuasive design* (Fogg, 2009). They include enticing, cajoling, or nudging someone into doing something through the use of persuasive technology.

Technology interventions have also been developed to change people's behaviors in other domains besides commerce, including safety, preventative healthcare, fitness, personal relationships, energy consumption, and learning. Here the emphasis is on changing someone's habits or doing something that will improve an individual's well-being through monitoring their behavior. An early example was Nintendo's Pokémon Pikachu device (see Figure 6.12) that was designed to motivate children into being more physically active on a consistent basis. The owner of the digital pet that lives in the device was required to walk, run, or jump each day to keep it alive. The wearer received credits for each step taken—the currency being watts that could be used to buy Pikachu presents. Twenty steps on the pedometer rewarded the player with 1 watt. If the owner did not exercise for a week, the virtual pet became angry and refused to play anymore. This use of positive rewarding and sulking can be a powerful means of persuasion, given that children often become emotionally attached to their virtual pets, especially when they start to care for them.



Figure 6.12 Nintendo's Pokémon Pikachu device

Source: [http://nintendo.wikia.com/wiki/File:Pok%C3%A9mon_Pikachu_2_GS_\(Device\).png](http://nintendo.wikia.com/wiki/File:Pok%C3%A9mon_Pikachu_2_GS_(Device).png)

ACTIVITY 6.3

Watch these two videos:

The Piano Staircase: <http://youtu.be/2lXh2n0aPyw>

The Outdoor Bin: <http://youtu.be/cbEKAwCoCKw>

Do you think that such playful methods are effective at changing people's behavior?

Comment

Volkswagen sponsored an open competition, called *the fun theory*, asking people to transform mundane artifacts into novel enjoyable user experiences in an attempt to change people's behavior for the better. The idea was to encourage a desired behavior by making it more fun. The Piano Staircase and the Outdoor Bin are the most well-known examples; the stairs sounded like piano keys being played as they were climbed, while the bin sounded like a well echoing when something was thrown into it. Research has shown that using these kinds of playful methods is very engaging, and they can help people overcome their social inhibition of taking part in an activity in a public place (Rogers et al., 2010a). ■

HAPIfork is a device that was developed to help someone monitor and track their eating habits (see Figure 6.13). If it detects that they are eating too quickly, it will vibrate (similar to the way a smartphone does when on silent mode), and an ambient light will appear at the end of the fork, providing the eater with real-time feedback intended to slow them down. The assumption is that eating too fast results in poor digestion and poor weight control and that making people aware that they are gobbling their food down can help them think about



Figure 6.13 Someone using the HAPIfork in a restaurant

Source: Helen Sharp

how to eat more slowly at a conscious level. Other data is collected about how long it took them to finish their meal, the number of fork servings per minute, and the time between them. These are turned into a dashboard of graphs and statistics so that the user can see each week whether their fork behavior is improving.

Nowadays, there are many kinds of mobile apps and personal tracking devices available that are intended to help people monitor various behaviors and change them based on the data collected and displayed back to them. These devices include fitness trackers, for example, Fitbit, and weight trackers, such as smart scales. Similar to HAPIfork, these devices are designed to encourage people to change their behavior by displaying dashboards of graphs showing how much exercise they have done or weight they have lost over a day, week, or longer period, compared with what they have done in the previous day, week, or month. These results can also be compared, through online leaderboards and charts, with how well they have done versus their peers and friends. Other techniques employed to encourage people to exercise more or to move when sedentary include goal setting, reminders, and rewards for good behavior. A survey of how people use such devices in their everyday lives revealed that people often bought them simply to try them or were given one as a present, rather than specifically trying to change a particular behavior (Rooksby et al., 2014). How, what, and when they tracked depended on their interests and lifestyles; some used them as a way of showing how fast they could run during a marathon or cycle on a course or how they could change their lifestyle to sleep or eat better.

An alternative approach to collecting quantified data about a behavior automatically is to ask people to write down manually how they are feeling now or to rate their mood and for them to reflect upon how they felt about themselves in the past. A mobile app called Echo, for example, asked people to write a subject line, rate their happiness at that moment, and add a description, photos, and/or videos if they wanted to (Isaacs et al., 2013). Sporadically, the app then asked them to reflect on previous entries. An assumption was that this type

of technology-mediated reflection could increase well-being and happiness. Each reflection was shown as a stacked card with the time and a smiley happiness rating. People who used the Echo app reported on the many positive effects of doing so, including reliving positive experiences and overcoming negative experiences by writing them down. The double act of recording and reflecting enabled them to generalize from the positive experiences and draw positive lessons from them.

The global concern about climate change has also led a number of HCI researchers to design and evaluate various energy-sensing devices that display real-time feedback. One goal is to find ways of helping people reduce their energy consumption, and it is part of a larger research agenda called *sustainable HCI*: see Mankoff et al., 2008; DiSalvo et al., 2010; Hazas et al., 2012. The focus is to persuade people to change their everyday habits with respect to environmental concerns, such as reducing their own carbon footprint, their community's footprint (for example, a school or workplace), or an even larger organization's carbon footprint (such as a street, town, or country).

Extensive research has shown that domestic energy use can be reduced by providing households with feedback on their consumption (Froehlich et al., 2010). The frequency of feedback is considered important; continuous or daily feedback on energy consumption has been found to yield higher savings results than monthly feedback. The type of graphical representation also has an effect. If the image used is too obvious and explicit (for instance, a finger pointing at the user), it may be perceived as too personal, blunt, or “in your face,” resulting in people objecting to it. In contrast, simple images (for example, an infographic or emoticon) that are more anonymous but striking and whose function is to get people's attention may be more effective. They may encourage people to reflect more on their energy use and even promote public debate about what is represented and how it affects them. However, if the image used is too abstract and implicit, other meanings may be attributed to it, such as simply being an art piece (such as an abstract painting with colored stripes that change in response to the amount of energy used), resulting in people ignoring it. The ideal may be somewhere in between. Peer pressure can also be effective, where peers, parents, or children chide or encourage one another to turn lights off, take a shower instead of a bath, and so on.

Another influencing factor is *social norms*. In a classic study by P. Wesley Schultz et al., (2007), households were shown how their energy consumption compared with their neighborhood average. Households above the average tended to decrease their consumption, but those using less electricity than average tended to increase their consumption. The study found that this “boomerang” effect could be counteracted by providing households with an emoticon along with the numerical information about their energy usage: households using less energy than average continued to do so if they received a smiley icon; households using more than average decreased their consumption even more if they were given a sad icon.

In contrast to the Schultz study, where each household's energy consumption was kept private, the Tidy Street project (Bird and Rogers, 2010) that was run in Brighton in the United Kingdom created a large-scale visualization of the street's electricity usage by spraying a stenciled display on the road surface using chalk (see Figure 6.14). The public display was updated each day to represent how the average electricity usage of the street compared to the city of Brighton's average. The goal was to provide real-time feedback that all of the homeowner's and the general public could see change each day over a period of three weeks. The street graph also proved to be very effective in getting people who lived on Tidy Street

to talk to each other about their electricity consumption and habits. It also encouraged them to talk with the many passersby who walked up and down the street. The outcome was to reduce electricity consumption in the street by 15 percent, which was considerably more than other projects in this area have been able to achieve.



Figure 6.14 Aerial view of the Tidy Street public electricity graph

Source: Helen Sharp

BOX 6.4

The Darker Side: Deceptive Technology

Technology is increasingly being used to deceive people into parting with their personal details, which allows Internet fraudsters to access their bank accounts and draw money from them. Authentic-looking letters, appearing to be sent from eBay, PayPal, and various leading banks, are spammed across the world, ending up in people's email in-boxes with messages such as "During our regular verification of accounts, we couldn't confirm your information. Please click here to update and verify your information." Given that many people have an account with one of these corporations, there is a good chance that they will be misled and unwittingly believe what is being asked of them, only to discover a few days later that they are several thousand dollars worse off. Similarly, letters from supposedly super-rich individuals in far-away countries, offering a share of their assets if the email recipient provides them with their bank details, have persistently been spammed worldwide. While many people are becoming increasingly wary of what are known as *phishing scams*, there are still many vulnerable individuals who are gullible to such tactics. ■

The term *phishing* is a play on the term *fishing*, which refers to the sophisticated way of luring users' financial information and passwords. Internet fraudsters are becoming smarter and are constantly changing their tactics. While the art of deception is centuries old, the increasing, pervasive, and often ingenious use of the web to trick people into divulging personal information can have catastrophic effects on society as a whole.

6.7 Anthropomorphism

Anthropomorphism is the propensity people have to attribute human qualities to animals and objects. For example, people sometimes talk to their computers as if they were humans, treat their robot cleaners as if they were their pets, and give all manner of cute names to their mobile devices, routers, and so on. Advertisers are well aware of this phenomenon and often create human-like and animal-like characters out of inanimate objects to promote their products. For example, breakfast cereals, butter, and fruit drinks have all been transmogrified into characters with human qualities (they move, talk, have personalities, and show emotions), enticing the viewer to buy them. Children are especially susceptible to this kind of magic, as witnessed by their love of cartoons where all manner of inanimate objects are brought to life with human-like qualities.

The finding that people, especially children, have a propensity to accept and enjoy objects that have been given human-like qualities has led many designers to capitalize on it, most notably in the design of virtual agents and interactive dolls, robots, and cuddly toys. Early commercial products like ActiMates were designed to encourage children to learn by playing with them. One of the first—Barney (a dinosaur)—attempted to motivate play in children by using human-based speech and movement (Strommen, 1998). The toys were programmed to react to the child and make comments while watching TV or working together on a computer-based task. In particular, Barney was programmed to congratulate the child whenever they produced a right answer and also to react to the content on-screen with appropriate emotions, for instance, cheering at good news and expressing concern at bad news. Interactive dolls have also been designed to talk, sense, and understand the world around them, using sensor-based technologies, speech recognition, and various mechanical servos embedded in their bodies. For example, the interactive doll Luvabella exhibits facial expressions, such as blinking, smiling, and making baby cooing noises in response to how her owner plays and looks after her. The more a child plays with her, the doll will learn to speak, transforming her babble into words and phrases.


A YouTube video (<https://youtu.be/au2VG9xRZZ0>) shows Luvabella in action and asks viewers to decide whether the interactive doll is creepy or cool. What do you think?

Furnishing technologies with personalities and other human-like attributes can make them more enjoyable and fun to interact with. They can also motivate people to carry out various activities, such as learning. Being addressed in the first person (for instance, “Hello, Noah! Nice to see you again. Welcome back. Now what were we doing last time? Oh yes, Exercise 5. Let’s start again.”) is more appealing than being addressed in the impersonal third person (“User 24, commence Exercise 5.”), especially for children. It can make them feel more at ease and reduce their anxiety. Similarly, interacting with screen characters like tutors and wizards can be more engaging than interacting with a dialog box.


ACTIVITY 6.4

A Robot or a Cuddly Pet?

Early robot pets, such as Sony's AIBO, were made of hard materials that made them look shiny and clunky. In contrast, a more recent trend has been to make them look and feel more like real pets by covering them up in fur and making them behave in more cute, pet-like ways. Two contrasting examples are presented in Figure 6.15a and 6.15b. Which do you prefer and why?



(a)



(b)

Figure 6.15 Robot pets: (a) Aibo and (b) The Haptic Creature

Source: (a) Jennifer Preece, (b) Used courtesy of Steve Yohanan. Photo by Martin Dee

Comment

Most people like stroking pets, so they may prefer a soft pet robot that they can also stroke, such as the one shown in Figure 6.15b. A motivation for making robot pets cuddly is to enhance the emotional experience people receive through using their sense of touch. For example, the Haptic Creature on the right is a robot that mimics a pet that might sit in your lap, such as a cat or a rabbit (Yohanan and MacLean, 2008). It is made up of a body, head, and two ears, as well as mechanisms that simulate breathing, a vibrating purr, and the warmth of a living creature. The robot “detects” the way it is touched by means of an array of (roughly 60) touch sensors laid out across its entire body and an accelerometer. When the Haptic Creature is stroked, it responds accordingly, using the ears, breathing, and purring to communicate its emotional state through touch. On the other hand, the sensors are also used by the robot to detect the human’s emotional state through touch. Note how the robot has no eyes, nose, or mouth. Facial expressions are the most common way humans communicate emotional states. Since the Haptic Creature communicates and senses emotional states solely through touch, the face was deliberately left off to prevent people from trying to “read” emotion from it. ■

A number of commercial physical robots have been developed specifically to support care giving for the elderly. Early ones were designed to be about 2 feet tall and were made from white plastic with colored parts that represented clothing or hair. An example was Zora (see Figure 6.16), developed in Belgium, that was marketed as a social robot for healthcare. One was bought by a nursing home in France. Many of the patients developed an emotional attachment to their Zora robot, holding it, cooing, and even giving it kisses on the head. However, some people found this kind of robot care a little demeaning. Certainly, it can never match the human touch and warmth that patients need, but there is no harm in it playing an entertaining and motivating role alongside human caregivers.

This video demonstrates how the Zora robot was used to entertain seniors and to help them get some exercise: <https://youtu.be/jcMNY5EnQNQ>.



Figure 6.16 The Zora robot

Source: <http://zorarobotics.be/>

In-Depth Activity

This in-depth activity requires you to try one of the emotion recognition apps available and to see how well it fares in recognizing different people's facial expressions. Download the AffdexMe app or Age Emotion Detector for Apple or Android. Take a photo of yourself looking natural and see what emotion it suggests.

(Continued)

1. How many emotions does it recognize?
2. Try to make a face for each of the following: sadness, anger, joy, fear, disgust, and surprise. After making a face for each, see how well the app detects the emotion you were expressing.
3. Ask a couple of other people to try it. See whether you can find someone with a beard and ask them to try, too. Does facial hair make it more difficult for the app to recognize an emotion?
4. What other application areas do you think these kinds of apps could be used for besides advertising?
5. What ethical issues does facial recognition raise? Has the app provided sufficient information as to what it does with the photos taken of people's faces?
6. How well would the recognition software work when used in a more natural setting where the user is not making a face for the camera?

Summary

This chapter described the different ways that interactive products can be designed (both deliberately and inadvertently) to make people respond in certain ways. The extent to which users will learn, buy a product online, quit a bad habit, or chat with others depends on the believability of the interface, how comfortable they feel when using a product, and/or how much they can trust it. If the interactive product is frustrating to use, annoying, or patronizing, users will easily become angry and despondent and often they stop using it. If, on the other hand, the product is pleasurable, is enjoyable to use, and makes people feel comfortable and at ease, then they will continue to use it, make a purchase, return to the website, or continue to learn.

This chapter also described various interaction mechanisms that can be used to elicit positive emotional responses in users and ways of avoiding negative ones. Further, it described how new technology has been developed to detect emotional states.

Key Points

- Emotional aspects of interaction design are concerned with how to facilitate certain states (for example, pleasure) or avoid certain reactions (such as frustration) in user experiences.
- Well-designed interfaces can elicit good feelings in people.
- Aesthetically pleasing interfaces can be a pleasure to use.
- Expressive interfaces can provide reassuring feedback to users as well as be informative and fun.
- Badly designed interfaces often make people frustrated, annoyed, or angry.
- Emotional AI and affective computing use AI and sensor technology for detecting people's emotions by analyzing their facial expressions and conversations.
- Emotional technologies can be designed to persuade people to change their behaviors or attitudes.
- Anthropomorphism is the attribution of human qualities to objects.
- Robots are being used in a variety of settings, including households and assisted-living homes.

Further Reading

CALVO, R.A and PETERS, D. (2014) *Positive Computing*. MIT. This book discusses how to design technology for well-being to make a happier and healthier world. As the title suggests, it is positive in its outlook. It covers the psychology of well-being, including empathy, mindfulness, joy, compassion, and altruism. It also describes the opportunities and challenges facing interaction designers who want to develop technology that can improve people's well-being.

HÖÖK, K. (2018) *Designing with the Body*. MIT. This book proposes that interaction design should consider the experiential, felt, and aesthetic stance that encompasses the design and use cycle. The approach suggested by the author is called *soma design*, where body and movements are viewed as very much part of the design process, and where a slow, thoughtful process is promoted that considers fundamental human values. It is argued that adopting this stance can yield better products and create healthier, more sustainable companies.

LEDOUX, J. E. (1998) *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. Simon & Schuster. This book explains what causes us to feel fear, love, hate, anger, and joy, and it explores whether we control our emotions versus them controlling us. The book also covers the origins of human emotions and explains that many evolved to enable us to survive.

McDUFF, D. & CZERWINSKI, M. (2018) Designing Emotionally Sentient Agents. *Communications of the ACM*, Vol. 61 No. 12, pages 74–83. This article provides an accessible overview of the burgeoning area of emotional agents. It presents the challenges, opportunities, dilemmas, concerns, and current applications that are now being developed, including bots, robots, and agents.

NORMAN, D. (2005) *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books. This book is an easy read while at the same time being thought-provoking. We get to see inside Dan Norman's kitchen and learn about the design aesthetics of his collection of teapots. The book also includes essays on the emotional aspects of robots, computer games, and a host of other pleasurable interfaces.

WALTER, A. (2011) *A Book Apart: Designing for Emotion*. Zeldman, Jeffrey. This short book is targeted at web designers who want to understand how to design websites that users will enjoy and want to return to. It covers the classic literature on emotions, and it proposes practical approaches to emotional web design.

Chapter 7

INTERFACES

7.1 Introduction

7.2 Interface Types

7.3 Natural User Interfaces and Beyond

7.4 Which Interface?

Objectives

The main goals of the chapter are to accomplish the following:

- Provide an overview of the many different kinds of interfaces.
- Highlight the main design and research considerations for each of the interfaces.
- Discuss what is meant by a natural user interface (NUI).
- Consider which interface is best for a given application or activity.

7.1 Introduction

When considering how to solve a user problem, the default solution that many developers choose to design is an app that can run on a smartphone. Making this easier still are many easy-to-use app developer tools that can be freely downloaded. It is hardly surprising, therefore, to see just how many apps there are in the world. In December 2018, Apple, for example, had a staggering 2 million apps in its store, many of which were games.

Despite the ubiquity of the smartphone app industry, the web continues to proliferate in offering services, content, resources, and information. A central concern is how to design them to be interoperable across different devices and browsers, which takes into account the varying form factors, size, and shape of smart watches, smartphones, laptops, smart TVs, and computer screens. Besides the app and the web, many other kinds of interfaces have been developed, including voice interfaces, touch interfaces, gesture interfaces, and multimodal interfaces.

The proliferation of technological developments has encouraged different ways of thinking about interaction design and UX. For example, input can be via mice, touchpads, pens, remote controllers, joysticks, RFID readers, gestures, and even brain-computer interaction. Output is equally diverse, appearing in the form of graphical interfaces, speech, mixed realities, augmented realities, tangible interfaces, wearable computing, and more.

The goal of this chapter is to give you an overview of the diversity of interfaces that can be developed for different environments, people, places, and activities. We present a catalog of 20 interface types, starting with command-based and ending with smart ones. For each interface, we present an overview and outline the key research and design considerations. Some are only briefly touched upon, while others, which are more established in interaction design, are described in greater depth.

NOTE

This chapter is not meant to be read from beginning to end; rather, it should be dipped into as needed to find out about a particular type of interface.

7.2 Interface Types

Numerous adjectives have been used to describe the different types of interfaces that have been developed, including *graphical*, *command*, *speech*, *multimodal*, *invisible*, *ambient*, *affective*, *mobile*, *intelligent*, *adaptive*, *smart*, *tangible*, *touchless*, and *natural*. Some of the interface types are primarily concerned with a function (for example, to be intelligent, to be adaptive, to be ambient, or to be smart), while others focus on the interaction style used (such as command, graphical, or multimedia), the input/output device used (for instance, pen-based, speech-based, or gesture-based), or the platform being designed for (for example, tablet, mobile, PC, or wearable). Rather than cover every possible type that has been developed or described, we have chosen to select the main types of interfaces that have emerged over the past 40 years. The interface types are loosely ordered in terms of when they were developed. They are numbered to make it easier to find a particular one. (See the following list for the complete set.) It should be noted, however, that this classification is for convenience of reference. The interface entries are not mutually exclusive since some products can appear in two or more categories. For example, a smartphone can be considered to be mobile, touch, or wearable.

The types of interfaces covered in this chapter include the following:

1. Command
2. Graphical
3. Multimedia
4. Virtual reality
5. Web
6. Mobile
7. Appliance
8. Voice
9. Pen
10. Touch
11. Gesture
12. Haptic
13. Multimodal
14. Shareable

15. Tangible
16. Augmented reality
17. Wearables
18. Robots and drones
19. Brain-computer interaction
20. Smart

Here is a selection of classic HCI videos on the Internet that demonstrate pioneering interfaces:

The Sketchpad: Ivan Sutherland (1963) describes the first interactive graphical interface: https://youtu.be/6orsmFndx_o.

The Mother of All Demos: Douglas Engelbart (1968) describes the first WIMP: <http://youtu.be/yJDv-zdhzMY>.

Put that there (1979): MIT demonstrates the first speech and gesture interface: <https://youtu.be/RyBEUyEtXQo>.

Unveiling the genius of multitouch interface design: Jeff Han gives a TED talk (2007): <http://youtu.be/ac0E6deG4AU>.

Intel's Future Technology Vision (2012): See http://youtu.be/g_cauM3kccl.

7.2.1 Command-Line Interfaces

Early interfaces required the user to type in commands that were typically abbreviations (for example, `ls`) at the prompt symbol appearing on the computer display, to which the system responded (for example, by listing current files). Another way of issuing commands is by pressing certain combinations of keys (such as `Shift+Alt+Ctrl`). Some commands are also a fixed part of the keyboard, such as delete, enter, and undo, while other function keys can be programmed by the user as specific commands (for instance, `F11` commanding print action).

Command-line interfaces were largely superseded by graphical interfaces that incorporated commands such as menus, icons, keyboard shortcuts, and pop-up/predictable text commands as part of an application. Where command-line interfaces continue to have an advantage is when users find them easier and faster to use than equivalent menu-based systems (Raskin, 2000). Users also prefer command-line interfaces for performing certain operations as part of a complex software package, such as for CAD environments (such as Rhino3D and AutoCAD), to allow expert designers to interact rapidly and precisely with the software. They also provide scripting for batch operations, and they are being increasingly used on the web, where the search bar acts as a general-purpose command-line facility, for example, www.yubnub.org.

System administrators, programmers, and power users often find that it is much more efficient and quicker to use command languages such as Microsoft's PowerShell. For example, it is much easier to delete 10,000 files in one go by using one command rather than scrolling through that number of files and highlighting those that need to be deleted. Command languages have also been developed for visually impaired people to allow them to interact in virtual worlds, such as Second Life (see Box 7.1).

BOX 7.1

Command Interfaces for Virtual Worlds

Virtual worlds, such as Second Life, have become popular places for learning and socializing. Unfortunately, people who are visually impaired cannot interact in a visual capacity. A command-based interface, called TextSL, was developed to enable them to participate using a screen reader (Folmer et al., 2009). Commands can be issued to enable the user to move their avatar around, interact with others, and find out about the environment in which they are located. Figure 7.1 shows that the user has issued the command for their avatar to smile and say hello to other avatars who are sitting by a log fire.

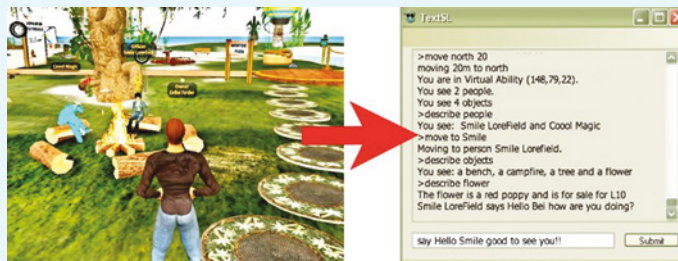


Figure 7.1 Second Life command-based interface for visually impaired users

Source: Used courtesy of Eelke Folmer ■

Watch a video demonstration of TextSL at http://youtu.be/0Ba_w7u44MM.

Research and Design Considerations

In the 1980s, much research investigated ways of optimizing command interfaces. The form of the commands (including use of abbreviations, full names, and familiar names), syntax (such as how best to combine different commands), and organization (for instance, how to structure options) are examples of some of the main areas that have been investigated (Shneiderman, 1998). A further concern was which command names would be the easiest to remember. A number of variables were tested, including how familiar users were with the chosen names. Findings from a number of studies, however, were inconclusive; some found specific names were better remembered than general ones (Barnard et al., 1982), others showed that names selected by users themselves were preferable (see Ledgard et al., 1981; Scapin, 1981), while yet others demonstrated that high-frequency words were better remembered than low-frequency ones (Gunther et al., 1986).

The most relevant design principle is consistency (see Chapter 1, “What Is Interaction Design?”). Therefore, the method used for labeling/naming the commands should be chosen to be as consistent as possible; for example, always use the first letters of the operation when using abbreviations.

7.2.2 Graphical User Interfaces

The Xerox Star interface (described in Chapter 3, “Conceptualizing Interaction”) led to the birth of the graphical user interface (GUI), opening up new possibilities for users to interact with a system and for information to be presented and represented within a graphical interface. Specifically, new ways of visually designing the interface became possible, which included the use of color, typography, and imagery (Mullet and Sano, 1995). The original GUI was called a WIMP (windows, icons, menus, pointer) and consisted of the following:

- *Windows*: Sections of the screen that can be scrolled, stretched, overlapped, opened, closed, and moved using a mouse
- *Icons*: Pictograms that represent applications, objects, commands, and tools that are opened or activated when clicked on
- *Menus*: Lists of options that can be scrolled through and selected in the way a menu is used in a restaurant
- *Pointing device*: A mouse controlling the cursor as a point of entry to the windows, menus, and icons on the screen

The first generation of WIMP interfaces were primarily boxy in design; user interaction took place through a combination of windows, scroll bars, checkboxes, panels, palettes, and dialog boxes that appeared on the screen in various forms (see Figure 7.2). Developers were largely constrained by the set of widgets available to them, of which the dialog box was most prominent. (A widget is a standardized display representation of a control, like a button or scroll bar, that can be manipulated by the user.) Nowadays, GUIs have been adapted for mobile and touchscreens. Instead of using a mouse and keyboard as input, the default action for most users is to swipe and touch using a single finger when browsing and interacting with digital content. (For more on this subject, see the sections on touch and mobile interfaces.)



Figure 7.2 The boxy look of the first generation of GUIs

The basic building blocks of the WIMP are still part of the modern GUI used as part of a display, but they have evolved into a number of different forms and types. For example, there are now many different types of icons and menus, including audio icons and audio menus, 3D animated icons, and even tiny icon-based menus that can fit onto a smartwatch screen (see Figure 7.3). Windows have also greatly expanded in terms of how they are used and what they are used for; for example, a variety of dialog boxes, interactive forms, and feedback/error message boxes have become pervasive. In addition, a number of graphical elements that were not part of the WIMP interface have been incorporated into the GUI. These include toolbars and docks (a row or column of available applications and icons of other objects such as open files) and rollovers (where text labels appear next to an icon or part of the screen as the cursor is rolled over it). Here, we give an overview of the design considerations concerning the basic building blocks of the WIMP/GUI: windows, menus, and icons.



Figure 7.3 Simple smartwatch menus with one, two, or three options

Source: <https://developer.apple.com/design/human-interface-guidelines/watchos/interface-elements/menus/>

Window Design

Windows were invented to overcome the physical constraints of a computer display, enabling more information to be viewed and tasks to be performed on the same screen. Multiple windows can be opened at any one time, for example, web browsers, word processing documents, photos, and slideshows, enabling the user to switch between them when needing to look at or work on different documents, files, and apps. They can also enable multiple instances of one app to be opened, such as when using a web browser.

Scrolling bars within windows also enable more information to be viewed than is possible on one screen. Scroll bars can be placed vertically and horizontally in windows to enable upward, downward, and sideway movements through a document and can be controlled using a touchpad, mouse, or arrow keys. Touch interfaces enable users to scroll content simply by swiping the screen to the left or right or up or down.

One of the problems of having multiple windows open is that it can be difficult to find specific ones. Various techniques have been developed to help users locate a particular window, a common one being to provide a list as part of an app menu. macOS also provides a function that shrinks all windows that are open for a given application so that they can be seen side by side on one screen. The user needs only to press one function key and then move the cursor over each one to see what they are called in addition to a visual preview. This technique enables users to see at a glance what they have in their workspace, and it also allows them easily to select one to bring forward. Another option is to display all of the windows open for a particular application, for example, Microsoft Word. Web browsers, like Firefox, also show thumbnails of the top sites visited and a selection of sites that you have saved or visited, which are called *highlights* (see Figure 7.4).




Figure 7.4 Part of the home page for the Firefox browser showing thumbnails of top sites visited and suggested highlight pages (bottom rows)

A particular kind of window that is commonly used is the *dialog box*. Confirmations, error messages, checklists, and forms are presented through dialog boxes. Information in the dialog boxes is often designed to guide user interaction, with the user following the sequence of options provided. Examples include a sequenced series of forms (such as Wizards) presenting the necessary and optional choices that need to be filled in when choosing a PowerPoint presentation or an Excel spreadsheet. The downside of this style of interaction is that there is a tendency to cram too much information or data entry fields into one box, making the interface confusing, crowded, and difficult to read (Mullet and Sano, 1995).

BOX 7.2

The Joys of Filling In Forms on the Web

For many of us, shopping on the Internet is generally an enjoyable experience. For example, choosing a book on Amazon or flowers from Interflora can be done at our leisure and convenience. The part that we don't enjoy, however, is filling in the online form to give the company the necessary details to pay for the selected items. This can often be a frustrating and time-consuming experience, especially as there is much variability between sites. Sometimes, it requires users to create an account and a new password. At other times, guest checkout is enabled. However, if the site has a record of your email address in its database, it won't allow you to use the guest option. If you have forgotten your password, you need to reset it, and this requires switching from the form to your email account. Once past this hurdle, different kinds of interactive forms pop up for you to enter your mailing address and credit card details. The form may provide the option of finding your address by allowing you to enter a postal or ZIP code. It may also have asterisks that denote fields that must be filled in.

Having so much inconsistency can frustrate the user, as they are unable to use the same mental model for filling in checkout forms. It is easy to overlook or miss a box that needs to be filled in, and after submitting the page, an error message may come back from the system saying it is incomplete. This may require the user to have to enter sensitive information again, as it will have been removed in the data processing stage (for example, the user's credit card number and the three or four-digit security code on the back or front of the card, respectively).

To add to the frustration, many online forms often accept only fixed data formats, meaning that, for some people whose information does not fit within its constraints, they are unable to complete the form. For example, one kind of form will accept only a certain type of mailing address format. The boxes are provided for: address line 1 and address line 2, providing no extra lines for addresses that have more than two lines; a line for the town/city; and a line for the ZIP code (if the site is based in the United States) or other postal code (if based in another country). The format for the codes is different, making it difficult for non-U.S. residents (and U.S. residents for other country sites) to fill in this part.

Another gripe about online registration forms is the country of residence box that opens up as a never-ending menu, listing all of the countries in the world in alphabetical order. Instead of typing in the country in which they reside, users are required to select the one they are from, which is fine if you happen to live in Australia or Austria but not if you live in Venezuela or Zambia (see Figure 7.5).

This is an example of where the design principle of recognition over recall (see Chapter 4, "Cognitive Aspects") does not apply and where the converse is true. A better design is to have a predictive text option, where users need only to type in the first one or two letters of their country to cause a narrowed-down list of choices to appear from which they can select within the interface. Or, one smart option is for the form to preselect the user's country of origin by using information shared from the user's computer or stored in the cloud. Automating the filling in of online forms, through providing prestored information about a user (for example, their address and credit card details), can obviously help reduce usability problems—provided they are OK with this.

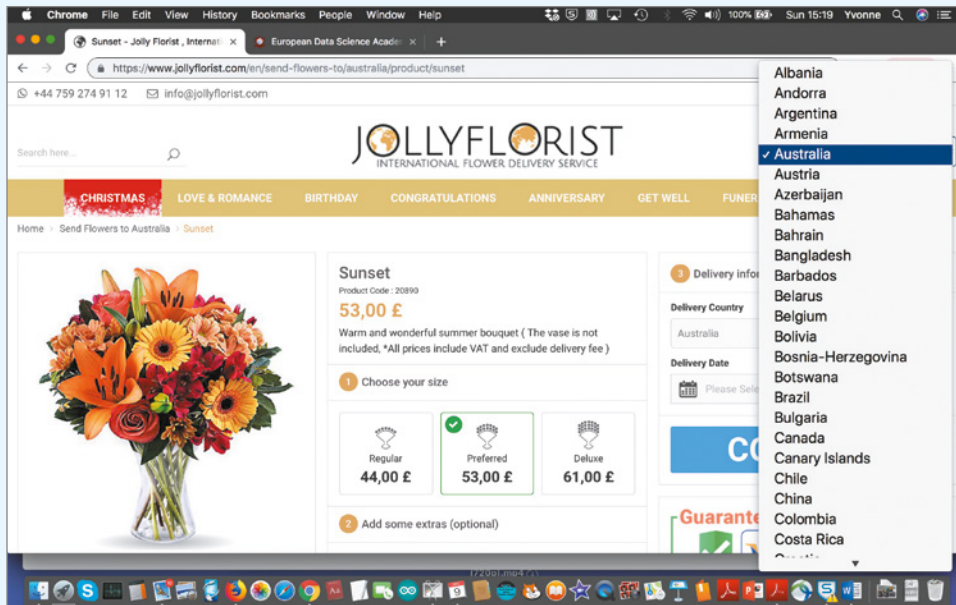


Figure 7.5 A scrolling menu of country names

Source: <https://www.jollyflorist.com> ■

ACTIVITY 7.1

Go to the Interflora site in the United Kingdom, click the international delivery option, and then click “select a country.” How are the countries ordered? Is it an improvement over the scrolling pop-up menu?

Comment

Earlier versions of the full list of countries to which flowers could be sent by interflora.co.uk listed eight countries at the top, starting with the United Kingdom and then the United States, France, Germany, Italy, Switzerland, Austria, and Spain. This was followed by the remaining set of countries listed in alphabetical order. The reason for having this particular ordering is likely to have been because the top eight are the countries that have most customers, with the U.K. residents using the service the most. The website has changed now to show top countries by national flag followed by a table format, grouping all of the countries in alphabetical order using four columns across the page (see Figure 7.6). Do you think this is an improvement over the use of a single scrolling list of country names shown in Figure 7.5? The use of letter headings and shading makes searching quicker.

| F | G | H | I | J |
|------------------|------------|-----------|-------------|---------|
| Fiji | Gabon | Haiti | Iceland | Jamaica |
| Finland | Germany | Holland | India | Japan |
| France | Gibraltar | Honduras | Indonesia | Jordan |
| French Guyana | Greece | Hong Kong | Iran | |
| French Polynesia | Greenland | Hungary | Ireland | |
| | Guadeloupe | | Israel | |
| | Guam | | Italy | |
| | Guatemala | | Ivory Coast | |

Figure 7.6 An excerpt of the listing of countries in alphabetical order from interflora.co.uk

Source: <https://www.interflora.co.uk> ■

Research and Design Considerations

A key research concern is *window management*—finding ways of enabling users to move fluidly between different windows (and displays) and to be able to switch their attention rapidly between windows to find the information they need or to work on the document/task within each window without getting distracted. Studies of how people use windows and multiple displays have shown that *window activation time* (that is, the time a window is open and with which the user interacts with it) is relatively short—an average of 20 seconds—suggesting that people switch frequently between different documents and applications (Hutchings et al., 2004). Widgets like the taskbar are often used for switching between windows.

Another technique is the use of tabs that appear at the top of the web browser that show the name and logo of the web pages that have been visited. This mechanism enables users to rapidly scan and switch among the web pages they have visited. However, the tabs can quickly multiply if a user visits a number of sites. To accommodate new ones, the web browser reduces the size of the tabs by shortening the information that appears on each. The downside of doing this, however, is it can make it more difficult to read and recognize web pages when looking at the smaller tabs. It is possible to reverse this shrinking by removing unwanted tabs by clicking the delete icon for each one. This has the effect of making more space available for the remaining tabs.

There are multiple ways that an online form can be designed to obtain details from someone. It is not surprising, therefore, that there are so many different types that are in use. Design guidelines are available to help decide which format and widgets are best to use. For example, see <https://www.smashingmagazine.com/printed-books/form-design-patterns/>. Another option is to automate form completion by asking the user to store their personal details on their machine or in a company's database, requiring them only to enter security information. However, many people are becoming leery of storing their personal data in this way—fearful because of the number of data breaches that are often reported in the news.

Menu Design

Interface menus are typically ordered across the top row or down the side of a screen using category headers as part of a menu bar. The contents of the menus are also for the large part invisible, only dropping down when the header is selected or rolled over with a mouse.

The various options under each menu are typically ordered from top to bottom in terms of most frequently used options and grouped in terms of their similarity with one another; for example, all formatting commands are placed together.

There are numerous menu interface styles, including flat lists, drop-down, pop-up, contextual, collapsible, mega, and expanding ones, such as cascading menus. *Flat menus* are good at displaying a small number of options at the same time or where the size of the display is small, for example on smartphones, cameras, and smartwatches. However, they often have to nest the lists of options within each, requiring several steps to be taken by a user to get to the list with the desired option. Once deep down in a nested menu, the user then has to take the same number of steps to get back to the top of the menu. Moving through previous screens can be tedious.

Expanding menus enable more options to be shown on a single screen than is possible with a single flat menu list. This makes navigation more flexible, allowing for the selection of options to be done in the same window. An example is the *cascading menu*, which provides secondary and even tertiary menus to appear alongside the primary active drop-down menu, enabling further related options to be selected, such as when selecting track changes from the tools menu leads to a secondary menu of three options by which to track changes in a Word document. The downside of using expanding menus, however, is that they require precise control. Users can often end up making errors, namely, overshooting or selecting the wrong options. In particular, cascading menus require users to move their cursor over the menu item, while holding the mouse or touchpad down, and then to move their cursor over to the next menu list when the cascading menu appears and select the next desired option. This can result in the user under or overshooting a menu option, or sometimes accidentally closing the entire menu. Another example of an expandable menu is a *mega menu*, in which many options can be displayed using a 2D drop-down layout (see Figure 7.7). This type of menu is popular with online shopping sites, where lots of items can be viewed at a glance on the same screen without the need to scroll. Hovering, tapping, or clicking is used to reveal more details for a selected item.

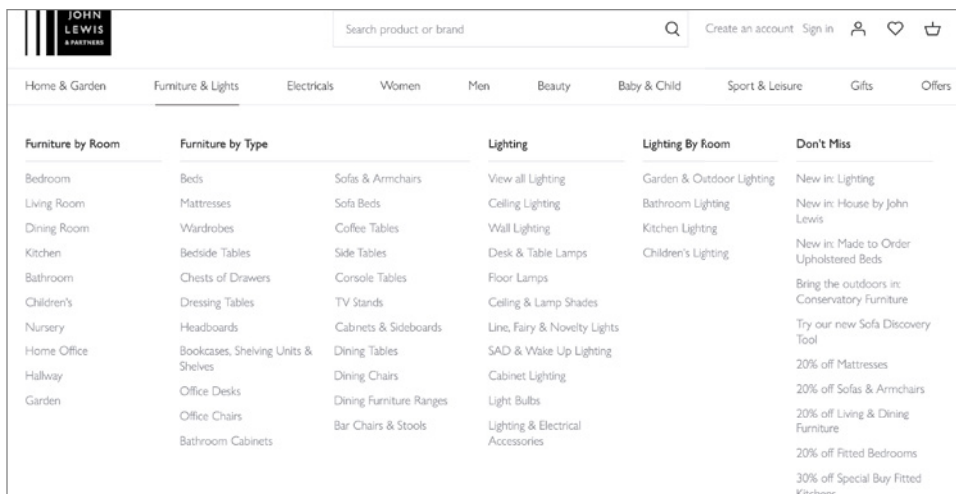


Figure 7.7 A mega menu

Source: <https://www.johnlewis.com>

Collapsible menus provide an alternative approach to expanding menus in that they allow further options to be made visible by selecting a header. The headings appear adjacent to each other, providing the user with an overview of the content available (see Figure 7.8). This reduces the amount of scrolling needed. *Contextual menus* provide access to often-used commands associated with a particular item, for example, an icon. They provide appropriate commands that make sense in the context of a current task. They appear when the user presses the Control key while clicking an interface element. For example, clicking a photo on a website together with holding down the Ctrl key results in a small set of relevant menu options appearing in an overlapping window, such as open it in a new window, save it, or copy it. The advantage of contextual menus is that they provide a limited number of options associated with an interface element, overcoming some of the navigation problems associated with cascading and expanding menus.




Figure 7.8 A template for a collapsible menu

Source: <https://inclusive-components.design/collapsible-sections/>. Reproduced with permission of Smashing Magazine

ACTIVITY 7.2

Open an application that you use frequently (for instance, a word processor, email client, or web browser) on a PC/laptop or tablet and look at the menu header names (but do not open them just yet). For each menu header—File, Edit, Tools, and so on—write down what options you think are listed under each. Then look at the contents under each header. How many options were you able to remember, and how many did you put in the wrong category? Now try to select the correct menu header for the following options (assuming that they are included in the application): Replace, Save, Spelling, and Sort. Did you select the correct header each time, or did you have to browse through a number of them?

Comment

Popular everyday applications, like word processors, have grown enormously in terms of the functions they offer. The current version (2019) of Microsoft Word, for example, has 8 menu headers and numerous toolbars. Under each menu header there are on average 15 options,

some of which are hidden under subheadings and appear only when they are rolled over with the mouse. Likewise, for each toolbar, there is a set of tools available, be it for Drawing, Formatting, Web, Table, or Borders. Remembering the location of frequently used commands like Spelling and Replace is often achieved by remembering their spatial location. For infrequently used commands, like sorting a list of references into alphabetical order, users can spend time flicking through the menus to find the command Sort. It is difficult to remember that the command Sort should be under the Table heading, since what it is doing is not a table operation, but a tool to organize a section of a document. It would be more intuitive if the command was under the Tool header along with similar tools like Spelling. What this example illustrates is just how difficult it can be to group menu options into clearly defined and obvious categories. Some fit into several categories, while it can be difficult to group others. The placement of options in menus can also change between different versions of an application as more functions are added. ■

Research and Design Considerations

An important design consideration is to decide which terms to use for menu options. Short phrases like “bring all to front” can be more informative than single words like “front.” However, the space for listing menu items is often restricted, such that menu names need to be short. They also need to be distinguishable, that is, not easily confused with one another so that the user does not choose the wrong one by mistake. Operations such as Quit and Save should also be clearly separated to avoid the accidental loss of work.

The choice of which type of menu to use will often be determined by the application and type of device for which is being designed. Which is best will also depend on the number of menu options and the size of the display available in which to present them. Flat menus are best for displaying a small number of options at one time, while expanding and collapsible menus are good for showing a large number of options, such as those available in file and document creation/editing applications. Usability testing comparing drop-down menus with mega menus has shown the latter to be more effective and easier to navigate. The main reason is that megamenus enable users to readily scan many items at a glance on the same page, and in doing so find what they are looking for (Nielsen and Li, 2017).

Icon Design

The appearance of icons in an interface came about following the Xerox Star project. They were used to represent objects as part of the desktop metaphor, namely, folders, documents, trashcans, inboxes, and outboxes. The assumption behind using icons instead of text labels is that they are easier to learn and remember, especially for non-expert computer users. They can also be designed to be compact and variably positioned on a screen.

Icons have become a pervasive feature of the interface. They now populate every app and operating system and are used for all manner of functions besides representing desktop objects. These include depicting tools (for example, Paint 3D), status (such as, Wi-Fi strength), categories of apps (for instance, health or personal finance), and a diversity of abstract operations (including cut, paste, next, accept, and change). They have also gone through many changes in their look and feel—black and white, color, shadowing, photorealistic images, 3D rendering, and animation have all been used.

Whereas early icon designers were constrained by the graphical display technology of the day, current interface developers have much more flexibility. For example, the use of anti-aliasing techniques enables curves and non-rectilinear lines to be drawn, enabling more photo-illustrative styles to be developed (*anti-aliasing* means adding pixels around a jagged border of an object to smooth its outline visually). App icons are often designed to be both visually attractive and informative. The goal is to make them inviting, emotionally appealing, memorable, and distinctive.

Different graphical genres have been used to group and identify different categories of icons. Figure 7.9 shows how colorful photorealistic images were used in the original Apple Aqua set, each slanting slightly to the left, for the category of *user* applications (such as email) whereas monochrome straight on and simple images were used for the class of *utility* applications (for instance, printer setup). The former has a fun feel to them, whereas the latter has a more serious look about them. While a number of other styles have since been developed, the use of slanting versus straight facing icons to signify different icon categories is still in use.



Figure 7.9 Two styles of Apple icons used to represent different kinds of functions

Icons can be designed to represent objects and operations in the interface using concrete objects and/or abstract symbols. The mapping between the icon and underlying object or operation to which it refers can be similar (such as a picture of a file to represent the object file), analogical (for instance, a picture of a pair of scissors to represent cut), or arbitrary (for example, the use of an X to represent delete). The most effective icons are generally those that are isomorphic since they have a direct mapping between what is being represented and how it is represented. Many operations in an interface, however, are of actions to be performed on objects, making it more difficult to represent them using direct mapping. Instead, an effective technique is to use a combination of objects and symbols that capture the salient part of an action by using analogy, association, or convention (Rogers, 1989). For example, using a picture of a pair of scissors to represent cut in a word-processing application provides a sufficient clue as long as the user understands the convention of cut for deleting text.

Another approach that many smartphone designers use is flat 2D icons. These are simple and use strong colors and pictograms or symbols. The effect is to make them easily recognizable and distinctive. Examples shown in Figure 7.10a include the white ghost on a yellow background (Snapchat), a white line bubble with a solid white phone handset in a speech bubble on a lime-green background (WhatsApp), and the sun next to a cloud (weather).



Figure 7.10 2D icons designed for (a) a smartphone and (b) a smartwatch

Source: (a) Helen Sharp (b) <https://support.apple.com/en-ca/HT205550>

Icons that appear on toolbars or palettes as part of an application or presented on small device displays (such as digital cameras or smartwatches) have much less screen real estate available. Because of this, they have been designed to be simple, emphasizing the outline form of an object or symbol and using only grayscale or one or two colors (see Figure 7.10b). They tend to convey the status, tool, or action using a concrete object (for example, the airplane symbol signaling whether the airplane mode is on or off) and abstract symbols (such as three waves that light up from none to all to convey the strength/power of the area's Wi-Fi).

ACTIVITY 7.3

Sketch simple icons to represent the following operations to appear on a digital camera screen:

- Turn image 90-degrees sideways.
- Crop the image.
- Auto-enhance the image.
- More options.

Show them to someone else, tell them that they are icons for a new digital camera intended to be really simple to use, and see whether they can understand what each represents.

Comment

Figure 7.11 shows the basic Edit Photo icons on an iPhone that appear at the bottom of the screen when a user selects the edit function. The box with extended lines and two arrows is the icon for cropping an image; the three overlapping translucent circles represents “different lenses” that can be used, the wand in the top-right corner means “auto-enhance,” and the circle with three dots in it means more functions.




Figure 7.11 The basic Edit Photo icons that appear at the top and bottom of an iPhone display ■

Research and Design Considerations

There are many icon libraries available that developers can download for free (for instance, <https://thenounproject.com/> or <https://fontawesome.com/>). Various online tutorials and books on how to design icons are also available (see Hicks, 2012) together with sets of proprietary guidelines and style guides. For example, Apple provides its developers with style guides, explaining why certain designs are preferable to others and how to design icon sets. Style

guides are also covered in more depth in Chapter 13, “Interaction Design in Practice.” On its developers’ website (developer.apple.com), advice is given on how and why certain graphical elements should be used when developing different types of icon. Among the various guidelines, it suggests that different categories of application (for example, Business, Utilities, Entertainment, and so on) should be represented by a different genre, and it recommends displaying a tool to communicate the nature of a task, such as a magnifying glass for searching or a camera for a photo editing tool. Android and Microsoft also provide extensive guidance and step-by-step procedures on how to design icons for its applications on its website.

To help disambiguate the meaning of icons, text labels can be used under, above, or to the side of their icons. This method is effective for toolbars that have small icon sets, such as those appearing as part of a web browser, but it is not as good for applications that have large icon sets, for example, photo editing or word processing, since the screen can get cluttered making it sometimes harder and longer to find an icon. To prevent text/icon clutter on the interface, a hover function can be used where a text label appears adjacent to or above an icon after the user holds the cursor over it for a second and for as long as the user keeps the cursor on it. This method allows identifying information to be temporarily displayed when needed.

7.2.3 Multimedia

Multimedia, as the name implies, combines different media within a single interface, namely, graphics, text, video, sound, and animation, and links them together with various forms of interactivity. Users can click links in an image or text that triggers another media such as an animation or a video. From there they can return to where they were previously or jump to another media source. The assumption is that a combination of media and interactivity can provide better ways of presenting information than can a single media, for example, just text or video alone. The added value of multimedia is that it can be easier for learning, better for understanding, more engaging, and more pleasant (Scaife and Rogers, 1996).

Another distinctive feature of multimedia is its ability to facilitate rapid access to multiple representations of information. Many multimedia encyclopedias and digital libraries have been designed based on this multiplicity principle, providing an assortment of audio and visual materials on a given topic. For example, when looking to find information about the heart, a typical multimedia-based encyclopedia will provide the following:

- One or two video clips of a real live heart pumping and possibly a heart transplant operation
- Audio recordings of the heart beating and perhaps an eminent physician talking about the cause of heart disease
- Static diagrams and animations of the circulatory system, sometimes with narration
- Several columns of hypertext, describing the structure and function of the heart

Hands-on interactive simulations have also been incorporated as part of multimedia learning environments. An early example was the *Cardiac Tutor*, developed to teach students about cardiac resuscitation. It required students to save patients by selecting the correct set of procedures in the correct order from various options displayed on the computer screen (Eliot and Woolf, 1994). Other kinds of multimedia narratives and

games have also been developed to support discovery learning by encouraging children to explore different parts of the display by noticing a hotspot or other kind of link. For example, <https://KidsDiscover.com/apps/> has many tablet apps that use a combination of animations, photos, interactive 3D models, and audio to teach kids about science and social studies topics. Using swiping and touching, kids can reveal, scroll through, select audio narration, and watch video tours. Figure 7.12, for example, has a “slide” mechanism as part of a tablet interface that enables the child to do a side-by-side comparison of what Roman ruins looks like now and in ancient Roman times.



Figure 7.12 An example of a multimedia learning app designed for tablets

Source: KidsDiscover app “Roman Empire for iPad”

Another example of a learning app with an interesting UI can be seen at <https://www.abcmouse.com/apps>.

Multimedia has largely been developed for training, educational, and entertainment purposes. But to what extent is the assumption that learning (such as reading and scientific inquiry skills) and playing can be enhanced through interacting with engaging multimedia interfaces true? What actually happens when users are given unlimited, easy access to multiple media and simulations? Do they systematically switch between the various media and “read” all of the multiple representations on a particular subject, or are they more selective in what they look at and listen to?

ACTIVITY 7.4

Watch this video of Don Norman appearing in his first multimedia CD-ROM book (1994), where he pops up every now and again in boxes or at the side of the page to illustrate the points being discussed on that page: <http://vimeo.com/18687931>.

How do you think students used this kind of interactive e-textbook?

Comment

Anyone who has interacted with educational multimedia knows just how tempting it is to play the video clips and animations while skimming through accompanying text or static diagrams. The former is dynamic, easy, and enjoyable to watch, while the latter is viewed as static and difficult to read from the screen. In an evaluation of the original Voyager's "First Person: Donald Norman, Defending Human Attributes in the Age of the Machine," students consistently admitted to ignoring the text on the interface in search of clickable icons of the author, which when selected would present an animated video of him explaining some aspect of design (Rogers and Aldrich, 1996). Given the choice to explore multimedia material in numerous ways, ironically, users tend to be highly selective as to what they actually pay attention to, adopting a channel-hopping mode of interaction. While enabling the users to select the information they want to view or features to explore for themselves, there is the danger that multimedia environments may in fact promote fragmented interactions where only part of the media is ever viewed. In a review of research comparing reading from screens versus paper, Lauren Singer and Patricia Alexandra (2017) found that despite students saying they preferred reading from screens, their actual performance was worse than when using paper-based textbooks.

Hence, online multimedia material may be good for supporting certain kinds of activities, such as browsing, but less optimal for others, for instance reading at length about a topic. One way to encourage more systematic and extensive interactions (when it is considered important for the activity at hand) is to require certain activities to be completed that entail the reading of accompanying text, before the user is allowed to move on to the next level or task. ■

Research and Design Considerations

A core research question is how to encourage users to interact with all aspects of a multimedia app, especially given the tendency to select videos to watch rather than text to read. One technique is to provide a diversity of hands-on interactivities and simulations that require the user to complete a task, solve a problem, or explore different aspects of a topic that involves reading some accompanying text. Specific examples include electronic notebooks that are integrated as part of the interface, where users can type in their own material; multiple-choice quizzes that provide feedback about how well they have done; interactive puzzles where they have to select and position different pieces in the right combination; and simulation-type games

(Continued)

where they have to follow a set of procedures to achieve some goal for a given scenario. Another approach is to employ *dynalinking*, where information depicted in one window explicitly changes in relation to what happens in another. This can help users keep track of multiple representations and see the relationship between them (Scaife and Rogers, 1996).

Specific guidelines are available that recommend how best to combine multiple media in relation to different kinds of task, for example, when to use audio with graphics, sound with animations, and so on, for different learning tasks. As a rule of thumb, audio is good for stimulating the imagination, movies for depicting action, text for conveying details, and diagrams for conveying ideas. From such generalizations, it is possible to devise a presentation strategy for online learning. This could be along the lines of the following:

1. Stimulate the imagination through playing an audio clip.
2. Present an idea in diagrammatic form.
3. Display further details about the concept through hypertext.

7.2.4 Virtual Reality

Virtual reality (VR) has been around since the 1970s when researchers first began developing computer-generated graphical simulations to create “the illusion of participation in a synthetic environment rather than external observation of such an environment” (Gigante, 1993, p. 3). The goal was to create user experiences that feel virtually real when interacting with an artificial environment. Images are displayed stereoscopically to the users—most commonly through VR headsets—and objects within the field of vision can be interacted with via an input device like a joystick.

The 3D graphics can be projected onto Cave Automatic Virtual Environment (CAVE) floor and wall surfaces, desktops, 3D TV, headsets, or large shared displays, for instance, IMAX screens. One of the main attractions of VR is that it can provide opportunities for new kinds of immersive experiences, enabling users to interact with objects and navigate in 3D space in ways not possible in the physical world or a 2D graphical interface. Besides looking at and navigating through a 360-degree visual landscape, auditory and haptic feedback can be added to make the experience feel even more like the real world. The resulting user experience can be highly engaging; it can feel as if one really is flying around a virtual world. People can become completely absorbed by the experience. The sense of presence can make the virtual setting seem convincing. *Presence*, in this case, means “a state of consciousness, the (psychological) sense of being in the virtual environment” (Slater and Wilbur, 1997, p. 605), where someone behaves in a similar way to how they would if at an equivalent real event.

VR simulations of the world can be constructed to have a higher level of fidelity with the objects they represent compared to other forms of graphical interfaces, for example, multimedia. The illusion afforded by the technology can make virtual objects appear to be very life-like and behave according to the laws of physics. For example, landing and take-off terrains developed for flight simulators can appear to be very realistic. Moreover, it is assumed that learning and training applications can be improved through having a greater fidelity to the represented world.

Another distinguishing feature of VR is the different viewpoints it can offer. Players can have a first-person perspective, where their view of the game or environment is through their own eyes, or a third-person perspective, where they see the world through an avatar visually represented on the screen. An example of a first-person perspective is that experienced in first-person shooter games such as DOOM, where the player moves through the environment without seeing a representation of themselves. It requires the user to imagine what they might look like and decide how best to move around. An example of a third-person perspective is that experienced in Tomb Raider, where the player sees the virtual world above and behind the avatar of Lara Croft. The user controls Lara's interactions with the environment by controlling her movements, for example, making her jump, run, or crouch. Avatars can be represented from behind or from the front, depending on how the user controls its movements. First-person perspectives are typically used for flying/driving simulations and games, for instance, car racing, where it is important to have direct and immediate control to steer the virtual vehicle. Third-person perspectives are more commonly used in games, learning environments, and simulations, where it is important to see a representation of self with respect to the environment and others in it. In some virtual environments, it is possible to switch between the two perspectives, enabling the user to experience different viewpoints on the same game or training environment.

In the beginning, head-mounted displays were used to present VR experiences. However, the visuals were often clunky, the headset uncomfortable to wear, and the immersive experience sometimes resulting in motion sickness and disorientation. Since then, VR technology has come of age and improved greatly. There are now many off-the-shelf VR headsets (for example Oculus Go, HTC Vive, and Samsung Gear VR) that are affordable and comfortable. They also have more accurate head tracking that allow developers to create more compelling games, movies, and virtual environments.

“Out of Home Entertainment” and VR arcades have also become popular worldwide and provide a range of social VR experiences, targeted at the general public. For example, Hyper-Reality has developed a number of spooky games, for 1–4 players, such as Japanese Adventures, Escape the Lost Pyramid, and the Void. Each game lasts for about 40 minutes, where players have to carry out a set of tasks, such as finding a lost friend in a realm. The immersive entertainment is full of surprises at every turn. One moment a player might be on solid ground and the next in complete darkness. The pleasure is often in not knowing what is going to happen next and being able to recount the experiences afterward with friends and family.

Another application area is how VR can enrich the experience of reporting and witnessing current affairs and news, especially feelings of empathy and compassion to real-life experiences (Aronson-Rath et al., 2016). For example, the BBC together with Aardman Interactive and University College London researchers developed a VR experience called “We Wait,” where they put the viewer in a place that few foreign reporters have been, namely, on a boat with a group of refugees crossing the Mediterranean Sea (Steed et al., 2018). The goal was to let news reporters and other participants experience how it felt to be there on the boat with the refugees. They used a particular artistic polygon style rather than realism to create the characters sitting on the boat (see Figure 7.13). The characters had expressive eyes intended to convey human emotion in response to gaze interaction. The avatars were found to generate an empathic response from participants.




Figure 7.13 Snapshot of polygon graphics used to represent avatars for the “We Wait” VR experience
Source: Steed, Pan, Watson and Slater, <https://www.frontiersin.org/articles/10.3389/frobt.2018.00112/full>.
Licensed Under CC-BY 4.0

VR is also starting to be used by airlines and travel companies to enrich someone’s planning experience of their travel destinations. For example, the airline KLM has developed a platform called iFly VR (<https://360.iflymagazine.com/>) that provides an immersive experience intended to inspire people to discover more about the world. A potential danger of this approach, however, is that if the VR experience is too lifelike it might make people feel they have ‘been there, done that’ and hence don’t need to visit the actual place. However, KLM’s rationale is quite the opposite; if you make the virtual experience so compelling people will want to go there even more. Their first foray into this adventure follows the famous “Fearless Chef” Kiran Jethwa into a jungle in Thailand to look for the world’s most remarkable coffee beans.

MagicLeap has pushed the envelope even further into new realms of virtual reality; combining cameras, sensors, and speakers in a headset that provides quite a different experience—one where the user can create their own worlds using various virtual tools, for example, painting a forest or building a castle—that then come alive in the actual physical space in which they reside. In this sense, it is not strictly VR, as it allows the wearer to see the virtual world and virtual objects they have created, or curated, blend with the physical objects in their living room or other space in which they are located. It is as if by magic the two are in the same world. In some ways it is a form of augmented reality (AR) - described in section 7.2.16.

Watch this video of MagicLeap's *Create World* where the virtual world meets the physical world in magical ways: <https://youtu.be/K5246156rcQ>.

Peter Rubin's (2018) guide to VR published in *Wired* magazine provides a summary and speculation about its future: <https://www.wired.com/story/wired-guide-to-virtual-reality/>.

Research and Design Considerations

VR has been developed to support learning and training for numerous skills. Researchers have designed apps to help people learn to drive a vehicle, fly a plane, and perform delicate surgical operations—where it is very expensive and potentially dangerous to start learning with the real thing. Others have investigated whether people can learn to find their way around a real building/place before visiting it by first navigating a virtual representation of it, see Gabrielli et al. (2000).

An early example of VR was the Virtual Zoo project. Allison et al. (1997) found that people were highly engaged and very much enjoyed the experience of adopting the role of a gorilla, navigating the environment, and watching other gorillas respond to their movements and presence.

Virtual environments (VE) have also been designed to help people practice social and speaking skills and confront their social phobias, see Cobb et al. (2002) and Slater et al. (1999). An underlying assumption is that the environment can be designed as a safe place to help people gently overcome their fears (for example, spiders, talking in public, and so forth) by confronting them through different levels of closeness and unpleasantness (such as by seeing a small virtual spider move far away, seeing a medium one sitting nearby, and then finally touching a large one). Studies have shown that people can readily suspend their disbelief, imagining a virtual spider to be a real one or a virtual audience to be a real audience. For example, Slater et al. (1999) found that people rated themselves as being less anxious after speaking to a virtual audience that was programmed to respond to them in a positive fashion than after speaking to virtual audiences programmed to respond to them negatively.

Core design considerations include the importance of having a virtual self-body as part of a VR experience to enhance the feeling of presence; how to prevent users from experiencing simulator sickness through experimenting with galvanic stimulation; determining the most effective ways of enabling users to navigate through them, for instance, first person versus third person; how to control their interactions and movements, for example, use of head and body movements; how best to enable users to interact with information in VR, for example, use of keypads, pointing, joystick buttons; and how to enable users to collaborate and communicate with others in the virtual environment.

(Continued)

A central concern is the level of realism to target. Is it necessary to design avatars and the environments that they inhabit to be life-like, using rich graphics, or can simpler and more abstract forms be used, but which nonetheless are equally capable of engendering a sense of presence? Do you need to provide a visual representation of the arm and hands for holding objects for a self-avatar, or is it enough to have continuous movement of the object? Research has shown that it is possible for objects to appear to be moving with invisible hands as if they were present. This has been coined as the “tomato presence,” that is, where presence is maintained using a stand-in object in VR (for instance, a tomato). (See <https://owlchemylabs.com/tomatopresence/>.)

3D software toolkits are also available, making it much easier for developers and researchers to create virtual environments. The most popular is Unity. 3D worlds can be created using their APIs, toolkits, and physics engines to run on multiple platforms, for example, mobile, desktop, console, TV, VR, AR, and the Web.

7.2.5 Website Design

Early websites were largely text-based, providing hyperlinks to different places or pages of text. Much of the design effort was concerned with the information architecture, that is, how best to structure information at the interface level to enable users to navigate and access it easily and quickly. For example, Jakob Nielsen (2000) adapted his and Rolf Molich’s usability guidelines (Nielsen and Molich, 1990) to make them applicable to website design, focusing on simplicity, feedback, speed, legibility, and ease of use. He also stressed how critical download time was to the success of a website. Simply, users who have to wait too long for a page to appear are likely to move on somewhere else.

Since then, the goal of web design has been to develop sites that are not only usable but also aesthetically pleasing. Getting the graphical design right, therefore, is critical. The use of graphical elements (such as background images, color, bold text, and icons) can make a website look distinctive, striking, and pleasurable for the user when they first view it and also to make it readily recognizable on their return. However, there is the danger that designers can get carried away with the appearance at the expense of making it difficult to find something and navigate through it.

Steve Krug (2014) discusses this usability versus attractiveness dilemma in terms of the difference between how designers create websites and how users actually view them. He argues that many web designers create sites as if the user was going to pore over each page, reading the finely crafted text word for word; looking at the use of images, color, icons, and so forth; examining how the various items have been organized on the site; and then contemplating their options before they finally select a link. Users, however, often behave quite differently. They will glance at a new page, scan part of it, and click the first link that catches their interest or looks like it might lead them to what they want.

Much of the content on a web page is not read. In Krug’s words, web designers are “thinking great literature” (or at least “product brochure”), while the user’s reality is much closer to a “billboard going by at 60 miles an hour” (Krug, 2014, p. 21). While somewhat of a caricature of web designers and users, his depiction highlights the discrepancy between the

meticulous ways that designers create their websites and the rapid and less than systematic approach that users take to view them. To help navigate their way through the many choices that web developers have to make, Jason Beaird and James George (2014) have come up with a number of guidelines intended to help web developers achieve a balance between using color, layout and composition, texture, typography, and imagery. They also cover mobile and responsive web design. Other website guidelines are mentioned in Chapter 16.

Web designers now have a number of languages available to design websites, such as Ruby and Python. HTML5 and web development tools, such as JavaScript and CSS, are also used. Libraries, such as React, and open source toolkits, such as Bootstrap, enable developers to get started quickly when prototyping their ideas for a website. WordPress also provides users with an easy-to-use interface and hundreds of free templates to use as a basis when creating their own website. In addition, built-in optimization and responsive, mobile-ready themes are available. Customized web pages are available for smartphone browsers that provide scrolling lists of articles, games, tunes, and so on, rather than hyperlinked pages.

Another interface element that has become an integral part of any website is *breadcrumb navigation*. Breadcrumbs are category labels that appear on a web page that enable users to peruse other pages without losing track of where they have come from (see Figure 7.14). The term comes from the way-finding technique that Hansel used in the Brothers Grimm fairy tale *Hansel and Gretel*. The metaphor conjures up the idea of leaving a path to follow back. Breadcrumbs are also used by search engine optimization tools that match up a user's search terms with relevant web pages using the breadcrumbs. Breadcrumbs also extol usability in a number of ways, including helping users know where they are relative to the rest of the website, enabling one-click access to higher site levels, attracting first time visitors to continue to browse a website after having viewed the landing page (Mifsud, 2011). Therefore, using them is good practice for other web applications besides websites.




Figure 7.14 A breadcrumb trail on the BestBuy website showing three choices made by the user to get to Smart Lights

Source: <https://www.bestbuy.ca>

With the arrival of tablets and smartphones, web designers needed to rethink how to design web browsers and websites for them, as they realized the touchscreen affords a different interaction style than PC/laptops. The standard desktop interface was found to not work as well on a tablet or smartphone. In particular, the typical fonts, buttons, and menu tabs were too small and awkward to select when using a finger. Instead of double-clicking interface elements, as users do with a mouse or trackpad, tablet and smartphone screens enable finger tapping. The main methods of navigation are by swiping and pinching. A new style of website emerged that mapped better to this kind of interaction style but also one that the

user could interact with easily when using a mouse and trackpad. Responsive websites were developed that change their layout, graphic design, font, and appearance depending on the screen size (smartphone, tablet, or PC) on which it was being displayed.

If you look at the design of many websites, you will see that the front page presents a banner at the top, a short promotional video about the company/product/service, arrows to the left or right to indicate where to flick to move through pages, and further details appearing beneath the home page that the user can scroll through. Navigation is largely done by swiping pages horizontally or scrolling up and down.

Tips on designing websites for tablets versus mobile phones can be found here: <https://css-tricks.com/a-couple-of-best-practices-for-tablet-friendly-design/>

BOX 7.3

In-Your-Face Web Ads

Web advertising has become pervasive and invasive. Advertisers realized how effective flashing and animated ads were for promoting their products, taking inspiration from the animated neon light advertisements used in city centers, such as London's Piccadilly Circus. But since banner ads emerged in the 1990s, advertisers have become even more cunning in their tactics. In addition to designing even flashier banner ads, more intrusive kinds of web ads have begun to appear on our screens. Short movies and garish cartoon animations, often with audio, now pop up in floating windows that zoom into view or are tagged on at the front end of an online newspaper or video clip. Moreover, this new breed of in-your-face, often personalized web ads frequently requires the user either to wait until they end or to find a check box to close the window down. Sites that provide free services, such as Facebook, YouTube, and Gmail, are also populated with web ads. The problem for users is that advertisers pay significant revenues to online companies to have their advertisements placed on their websites, entitling them to say where, what, and how they should appear. One way users can avoid them is to set up ad blockers when browsing the web. ■

Research and Design Considerations

There are numerous classic books on web design and usability (for example, Krug, 2014; Cooper et al., 2014). In addition, there are many good online sites offering guidelines and tips. For example, the BBC provides online guidance specifically for how to design responsive websites that includes topics such as context, accessibility, and modular design. See <https://www.bbc.co.uk/gel/guidelines/how-to-design-for-the-web>. Key design considerations for all websites are captured well by three core questions proposed by Keith Instone (quoted in Veen, 2001): Where am I? What's here? Where can I go?

ACTIVITY 7.5

Look at a fashion brand's website, such as Nike, and describe the kind of interface used. How does it contravene the design principles outlined by Jeffrey Veen? Does it matter? For what type of user experience is it providing? What was your experience in engaging with it?

Comment

Fashion companies' sites, like Nike, are often designed to be more like a cinematic experience and use rich multimedia elements, including videos, sounds, music, animations, and interactivity. Branding is central. In this sense, it contravenes what are considered core usability guidelines. Specifically, the site has been designed to entice the visitor to enter the virtual store and watch high-quality and innovative movies that show cool dudes wearing their products. Often, multimedia interactivities are embedded into the sites to help the viewer move to other parts of the site, for example by clicking on parts of an image or a video playing. Screen widgets are also provided, such as menus, skip over, and next buttons. It is easy to become immersed in the experience and forget that it is a commercial store. It is also easy to get lost and not to know—Where am I? What's here? Where can I go? But this is precisely what companies such as Nike want their visitors to do and to enjoy: the experience. ■

7.2.6 Mobile Devices

Mobile devices have become pervasive, with people increasingly using them in all aspects of their everyday and working lives—including phones, fitness trackers, and watches. Customized mobile devices are also used by people in a diversity of work settings where they need access to real-time data or information while walking around. For example, they are now commonly used in restaurants to take orders, at car rental agencies to check in car returns, in supermarkets for checking stock, and on the streets for multiplayer gaming.

Larger-sized tablets are also used in mobile settings. For example, many airlines provide their flight attendants with one so that they can use their customized flight apps while airborne and at airports; sales and marketing professionals also use them to demonstrate their products or to collect public opinions. Tablets and smartphones are also commonly used in classrooms that can be stored in special “tabcabbies” provided by schools for safe keeping and recharging.

Smartphones and smartwatches have an assortment of sensors embedded in them, such as an accelerometer to detect movement, a thermometer to measure temperature, and galvanic skin response to measure changes in sweat level on one's skin. Other apps may be designed for fun. An example of an early app developed by magician Steve Sheraton simply for a moment of pleasure is iBeer (see Figure 7.15). Part of its success was due to the ingenious use of the accelerometer inside the phone. It detects the tilting of the iPhone and uses this information to mimic a glass of beer being consumed. The graphics and sounds are also very enticing; the color of the beer together with frothy bubbles and accompanying sound effects gives the illusion of virtual beer being swished around a virtual glass. The beer can be drained if the phone is tilted enough, followed by a belch sound when it has been finished.



Figure 7.15 The iBeer smartphone app

Source: Hottrix

Smartphones can also be used to download contextual information by scanning barcodes in the physical world. Consumers can instantly download product information by scanning barcodes using their iPhone when walking around a supermarket, including allergens, such as nuts, gluten, and dairy. For example, the GoodGuide app enables shoppers to scan products in a store by taking a photo of their barcode to see how they rate for healthiness and impact on the environment. Others include concert tickets and location-based notifications.

Another method that provides quick access to relevant information is the use of quick response (QR) codes that store URLs and look like black-and-white checkered squares (see Figure 7.16). They work by people taking a picture using their camera phone that then takes them to a particular website. However, despite their universal appeal to companies as a way of providing additional information or special offers, not many people actually use them in practice. One of the reasons is that they can be slow, tricky, and cumbersome to use *in situ*. People have to download a QR reader app first, open it, and then try to hold it over the QR code to take a photo, which can take time to open up a webpage.



Figure 7.16 QR code appearing on a magazine page

ACTIVITY 7.6

Smartwatches, such those made by Google, Apple, and Samsung, provide a multitude of functions including fitness tracking, streaming music, texts, email, and the latest tweets. They are also context and location aware. For example, on detecting the user's presence, promotional offers may be pinged to them from nearby stores, tempting them in to buy. How do you feel about this? Do you think it is the same or worse compared to the way advertisements appear on a user's smartphone? Is this kind of context-based advertising ethical?

Comment

Smartwatches are similar to smartphones in that they, too, get pinged with promotions and ads for nearby restaurants and stores. However, the main difference is that when worn on a wrist, smartwatches are ever-present; the user only needs to glance down at it to notice a new notification, whereas they have to take their phones out of their pockets and purses to see what new item has been pinged (although some people hold their smartphone permanently in their hands). This means that their attention is always being given to the device, which could make them susceptible to responding to notifications and spending more money. While some people might like to get 10 percent off on coffee if they walk into the cafe that has just sent them a digital voucher, for others such notifications may be seen as very annoying as they are constantly bombarded with promotions. Worse still, it could tempt children and vulnerable people who are wearing such a watch to spend money when perhaps they shouldn't or to nag their parents or caretakers to buy it for them. However, smartwatch companies are aware of this potential problem, and they provide settings that the user can change in terms of the level and type of notifications they want to receive. ■

Research and Design Considerations

Mobile interfaces typically have a small screen and limited control space. Designers have to think carefully about what type of dedicated hardware controls to include, where to place them on the device, and then how to map them to the software. Apps designed for mobile interfaces need to take into account the ability to navigate through content when using a mobile display is constrained, whether using touch, pen, or keypad input. The use of vertical and horizontal scrolling provides a rapid way of scanning through images, menus, and lists. A number of mobile browsers have also been developed that allow users to view and navigate the Internet, magazines, or other media in a more streamlined way. For example, Microsoft's Edge browser was one of the first mobile browsers that was designed to make it easier to find, view, and manage content on the go. It provides a customized reading view that enables the user to re-organize the content of a web page to make it easier for them to focus on what they want to read. The trade-off, however, is that it makes it less obvious how to perform other functions that are no longer visible on the screen.

(Continued)

Another key concern for mobile display design is the size of the area on the display that the user touches to make something happen, such as a key, icon, button, or app. The space needs to be big enough for “all fingers” to press accurately. If the space is too small, the user may accidentally press the wrong key, which can be annoying. The average fingertip is between one and two centimeters wide, so target areas should be at least 7 mm to 10 mm so that they can be accurately tapped with a fingertip. Fitts’ law (see Chapter 16) is often used to help with evaluating hit area. In their developer design guidelines, Apple also suggests providing ample touch targets for interactive elements, with a minimum tappable area of 44 pts. × 44 pts. for all controls.

A number of other guidelines exist providing advice on how to design interfaces for mobile devices (for instance, see Babich, 2018). An example is avoiding clutter by prioritizing one primary action per screen.

7.2.7 Appliances

Appliances include machines for everyday use in the home (for example, washing machines, microwave ovens, refrigerators, toasters, bread makers, and smoothie makers). What they have in common is that most people using them will be trying to get something specific done in a short period of time, such as starting a wash, watching a program, buying a ticket, or making a drink. They are unlikely to be interested in spending time exploring the interface or looking through a manual to see how to use the appliance. Many of them now have LED displays that provide multiple functions and feedback about a process (such as temperature, minutes remaining, and so on). Some have begun to be connected to the Internet with companion devices, enabling them to be controlled by remote apps. An example is a coffee maker that can be controlled to come on at a certain time from an app running on a smartphone or controlled by voice.

Research and Design Considerations

Alan Cooper et al. (2014) suggest that appliance interfaces require the designer to view them as transient interfaces, where the interaction is short. All too often, however, designers provide full-screen control panels or an unnecessary array of physical buttons that serve to frustrate and confuse the user where only a few, presented in a structured way, would be much better. Here the two fundamental design principles of simplicity and visibility are paramount. Status information, such as what the photocopier is doing, what the ticket machine is doing, and how much longer the wash is going to take should be provided in a simple form and at a prominent place on the interface. A key design question is: as soft displays increasingly become part of an appliance interface, for example, LCD and touchscreens, what are the trade-offs with replacing the traditional physical controls, such as dials, buttons, and knobs, with these soft display controls?

ACTIVITY 7.7

Look at the controls on your toaster (or the one in Figure 7.17 if you don't have one nearby) and describe what each does. Consider how these might be replaced with an LCD screen. What would be gained and lost from changing the interface in this way?



Figure 7.17 A typical toaster with basic physical controls

Source: <https://uk.russellhobbs.com/product/brushed-stainless-steel-toaster-2-slice>

Comment


Standard toasters have two main controls, the lever to press down to start the toasting and a knob to set the amount of time for the toasting. Many come with a small eject button that can be pressed if the toast starts to burn. Some also come with a range of settings for different ways of toasting (such as one side, frozen, and so forth), selected by moving a dial or pressing buttons.

Designing the controls to appear on an LCD screen would enable more information and options to be provided, for example, only toast one slice, keep the toast warm, or automatically pop up when the toast is burning. It would also allow precise timing of the toasting in minutes and seconds. However, it is likely to increase the complexity of what previously was a set of logical and very simple actions. This has happened in the evolution of microwaves, washing machines, and tea kettles that have digital interfaces. They also offer many more options for warming food up, washing clothes, or the temperature to heat water. The downside of increasing the number of choices, especially when the interface is not designed well to support this, is that it can make for a more difficult user experience for mundane tasks. ■

7.2.8 Voice User Interfaces

A *voice user interface* (VUI) involves a person talking with a spoken language app, such as a search engine, a train timetable, a travel planner, or a phone service. It is commonly used for inquiring about specific information (for instance, flight times or the weather) or issuing a command to a machine (such as asking a smart TV to select an Action movie or asking a smart speaker to play some upbeat music). Hence, VUIs use an interaction type of command or conversation (see Chapter 3), where users speak and listen to an interface rather than click on, touch, or point to it. Sometimes, the interaction style can involve the user responding where the system is proactive and initiates the conversation, for example, asking the user if they would like to stop watching a movie or listen to the latest breaking news.

The first generation of speech systems earned a reputation for *mishearing* all too often what a person said (see cartoon). However, they are now much more sophisticated and have higher levels of recognition accuracy. Machine learning algorithms have been developed that are continuing to improve their ability to recognize what someone is saying. For speech output, actors are often used to record answers, messages, and prompts, which are much friendlier, more convincing, and more pleasant than the artificially-sounding synthesized speech that was typically used in the early systems.



Source: Reproduced with permission of King Features Syndicate

VUIs have become popular for a range of apps. Speech-to-text systems, such as Dragon, enable people to dictate rather than have to type, whether it is entering data into a spreadsheet, using a search engine, or writing a document. The words spoken appear on the screen. For some people, this mode of interaction is more efficient, especially when they are on the move. Dragon claims on their website that it is three times faster than typing and it is 99 percent accurate. Speech technology is also used by people with visual impairments, including speech recognition word processors, page scanners, web readers, and VUIs for operating home control systems, including lights, TV, stereo, and other home appliances.

One of the most popular applications of speech technology is call routing, where companies use an automated speech system to enable users to reach one of their services during a phone call. Callers voice their needs in their own words, for example, “I’m having problems with my Wi-Fi router,” and in response are automatically forwarded to the appropriate service (Cohen et al., 2004). This is useful for companies, as it can reduce operating costs. It can also increase revenue by reducing the number of lost calls. The callers may be happier, as their call can be routed to an available agent (real or virtual) rather than being lost or sent to voicemail.

In human conversations, people often interrupt each other, especially if they know what they want, rather than waiting for someone to go through a series of options. For example, they may stop the waitress at a restaurant in midflow when describing the specials if they know what they want, rather than let her go through the entire list. Similarly, speech technology has been designed with a feature called *barge-in* that allows callers to interrupt a system message and provide their request or response before the message has finished playing. This can be useful if the system has numerous options from which the caller may choose, and the chooser knows already what they want.

There are several ways that a VUI dialog can be structured. The most common is a directed dialogue where the system is in control of the conversation, asking specific questions and requiring specific responses, similar to filling in a form (Cohen et al., 2004):

System: Which city do you want to fly to?

Caller: London

System: Which airport: Gatwick, Heathrow, Luton, Stansted, or City?

Caller: Gatwick

System: What day do you want to depart?

Caller: Monday next week.

System: Is that Monday, May 5?

Caller: Yes

Other systems are more flexible, allowing the user to take more initiative and specify more information in one sentence (for example, “I’d like to go to Paris next Monday for two weeks”). The problem with this approach is that there is more chance for error, since the caller might assume that the system can follow all of their needs in one pass as a real travel agent would (for example, “I’d like to go to Paris next Monday for two weeks, and would like the cheapest possible flight, preferably leaving from Gatwick airport and definitely with no stop-overs ...”). The list is simply too long and would overwhelm the system’s parser. Carefully guided prompts can be used to get callers back on track and help them speak appropriately (for instance, “Sorry, I did not get all that. Did you say you wanted to fly next Monday?”).

A number of speech-based phone apps exist that enable people to use them while mobile, making them more convenient to use than text-based entry. For example, people can voice queries into their phone using Google Voice or Apple Siri rather than entering text manually. Mobile translators allow people to communicate in real time with others who speak a different language by letting a software app on their phone do the talking (for example, Google Translate). People speak in their own language using their phone while the software translates what each person is saying into the language of the other one. Potentially, this means people from all over the world (there are more than 6,000 languages) can talk to one another without having to learn another language.

Voice assistants, like Amazon’s Alexa and Google Home, can be instructed by users to entertain in the home by telling jokes, playing music, keeping track of time, and enabling users to play games. Alexa also offers a range of “skills,” which are voice-driven capabilities intended to provide a more personalized experience. For example, “Open the Magic Door” is an interactive story skill that allows users to choose their path in a story by selecting different options through the narrative. Another one, “Kids court,” allows families to settle arguments in an Alexa-run court while learning about the law. Many of the skills are designed

to support multiple users taking part at the same time, offering the potential for families to play together. Social interaction is encouraged by the smart speaker that houses Alexa or Home. Smart speakers sit in a common space for all to use (similar to a toaster or refrigerator). In contrast, handheld devices, such as a smartphone or tablet, support only single use and ownership.

Despite advances in speech recognition, conversational interaction is limited mainly to answering questions and responding to requests. It can be difficult for VUIs to recognize children's speech, which is not as articulate as adults. For example, Druga et al. (2017) found that young children (3–4 years old) experienced difficulty interacting with conversational and chat agents, resulting in them becoming frustrated. Also, voice assistants don't always recognize who is talking in a group, such as a family, and always need to be called by their name each time someone wants to interact with it. There is still a way to go before voice assistant interaction resembles human conversation.

Research and Design Considerations

Key research questions are what conversational mechanisms to use to structure the voice user interface and how human-like they should be. Some researchers focus on how to make it appear natural (that is, like human conversation), while others are concerned more with how to help people navigate efficiently through a menu system by enabling them to recover easily from errors (their own or the system's), to be able to escape and go back to the main menu (similar to the undo button of a GUI), and to guide those who are vague or ambiguous in their requests for information or services using prompts. The type of voice actor, male, female, neutral, or dialect, and form of pronunciation are also topics of research. Do people prefer to listen to and are more patient with a female or male voice? What about one that is jolly versus one that is serious?

Michael Cohen et al. (2004) discuss the pros and cons of using different techniques for structuring the dialogue and managing the flow of voice interactions, the different ways of expressing errors, and the use of conversational etiquette—all still relevant for today's VUIs. A number of commercial guidelines are available for voice interfaces. For example, Cathy Pearl (2016) has written a practical book that provides a number of VUI design principles and topics, including which speech recognition engine to use, how to measure the performance of VUIs, and how to design VUIs for different interfaces, for example, a mobile app, toy, or voice assistant.

7.2.9 Pen-Based Devices

Pen-based devices enable people to write, draw, select, and move objects on an interface using light pens or styluses that capitalize on the well-honed drawing and writing skills that are developed from childhood. They have been used to interact with tablets and large displays, instead of mouse, touch, or keyboard input, for selecting items and supporting freehand sketching. Digital ink, such as Anoto, uses a combination of an ordinary ink pen with a digital camera that




Figure 7.18 The Anoto pen being used to fill in a paper form and a schematic showing its internal components

Source: www.grafichewanda.it/anoto.php?language=EN

digitally records everything written with the pen on special paper (see Figure 7.18). The pen works by recognizing a special nonrepeating dot pattern that is printed on the paper. The nonrepeating nature of the pattern means that the pen is able to determine which page is being written on and where on the page the pen is pointing. When writing on digital paper with a digital pen, infrared light from the pen illuminates the dot pattern, which is then picked up by a tiny sensor. The pen decodes the dot pattern as the pen moves across the paper and stores the data temporarily in the pen. The digital pen can transfer data that has been stored in the pen via Bluetooth or a USB port to a computer. Handwritten notes can also be converted and saved as standard typeface text. This can be useful for applications that require people to fill in paper-based forms and also for taking notes during meetings.

Another advantage of digital pens is that they allow users to annotate existing documents, such as spreadsheets, presentations, and diagrams quickly and easily in a similar way to how they would do this when using paper-based versions. This is useful for a team who is working together and communicating to each other from different locations. One problem with using pen-based interactions on small screens, however, is that sometimes it can be difficult to see options on the screen because a user's hand can obscure part of it when writing.

BOX 7.4

Electronic Ink

Digital ink is not to be confused with the term *electronic ink* (or *e-ink*). Electronic ink is a display technology designed to mimic the appearance of ordinary ink on paper used in e-readers, such as the Kindle. The display used reflects light like ordinary paper. ■

7.2.10 Touchscreens

Single touchscreens, used in walk-up kiosks such as ticket machines or museum guides, ATMs, and cash registers (for instance, restaurants), have been around for a while. They work by detecting the presence and location of a person's touch on the display; options are selected by tapping on the screen. *Multitouch surfaces*, on the other hand, support a much wider range of more dynamic fingertip actions, such as swiping, flicking, pinching, pushing, and tapping. They do this by registering touches at multiple locations using a grid (see Figure 7.19). This multitouch method enables devices, such as smartphones and tabletops, to recognize and respond to more than one touch at the same time. This enables users to use multiple digits to perform a variety of actions, such as zooming in and out of maps, moving photos, selecting letters from a virtual keyboard when writing, and scrolling through lists. Two hands can also be used together to stretch and move objects on a tabletop surface, similar to how both hands are used to stretch an elastic band or scoop together a set of objects.




Figure 7.19 A multitouch interface

Source: www.sky-technology.eu/en/blog/article/item/multi-touch-technology-how-it-works.html

The flexibility of interacting with digital content afforded by finger gestures has resulted in many ways of experiencing digital content. This includes reading, scanning, zooming, and searching interactive content on tablets, as well as creating new digital content.

Research and Design Considerations

Touchscreens have become pervasive, increasingly becoming the main interface that many people use on a daily basis. However, they are different from GUIs, and a central design concern is what types of interaction techniques to use to best support different activities. For example, what is the optimal way to enable users to choose from menu options, find files, save documents, and so forth, when using a touch interface? These operations are well mapped to interaction styles available in a GUI, but it is not as obvious how to support them on a touch interface. Alternative conceptual models have been developed for the user to carry out these actions on the interface, such as the use of cards, carousels, and stacks (see Chapter 3). The use of these objects enables users to swipe and move through digital content quickly. However, it is also easy to swipe too far when using a carousel. Typing on a virtual keyboard with two thumbs or one fingertip is also not as fast and efficient as using both hands when using a conventional keyboard, although many people have learned to be very adept at pecking at virtual keys on a smartphone. Predictive text can also be used to help people type faster.

Both hands may be used on multitouch tabletops to enable users to make digital objects larger and smaller or to rotate them. Dwelling touches (pressing and holding a finger down) can also be used to enable a user to perform dragging actions and to bring up pop-up menus. One or more fingers can also be used together with a dwell action to provide a wider range of gestures. However, these can be quite arbitrary, requiring users to learn them rather than being intuitive. Another limitation of touchscreens is that they do not provide tactile feedback in the same way that keys or mice do when pressed. To compensate, visual, audio, and haptic feedback can be used. See also the section on shareable interfaces for more background on multitouch design considerations.

7.2.11 Gesture-Based Systems

Gestures involve moving arms and hands to communicate (for instance, waving to say goodbye or raising an arm to speak in class) or to provide information to someone (for example, holding two hands apart to show the size of something). There has been much interest in how technology can be used to capture and recognize a user's gestures for input by tracking them using cameras and then analyzing them using machine learning algorithms.

David Rose (2018) created a video that depicts many sources of inspiration for where gesture is used in a variety of contexts, including those made by cricket umpires, live concert signers for the deaf, rappers, Charlie Chaplin, mime artists, and Italians. His team at IDEO developed a gesture system to recognize a small set of gestures and used these to control a Philips HUE light set and a Spotify station. They found that gestures need to be sequential to be understood in the way a sentence is composed of a noun, then verb, and object plus operation. For example, for “speaker, on,” they used a gesture on one hand to designate the

noun, and another on the other hand to designate the verb. So, to change the volume, the user needs to point to a speaker with their left hand while raising their right hand to signal turn the volume up.

Watch David Rose's inspirations for gesture video at <https://vimeo.com/224522900>.

One area where gesture interaction has been developed is in the operating room. Surgeons need to keep their hands sterile during operations but also need to be able to look at X-rays and scans during an operation. However, after being scrubbed and gloved, they need to avoid touching any keyboards, phones, and other nonsterile surfaces. A far from ideal workaround is to pull their surgical gown over their hands and manipulate a mouse through the gown. As an alternative, Kenton O'Hara et al. (2013) developed a touchless gesture-based system, using Microsoft's Kinect technology, which recognized a range of gestures that surgeons could use to interact with and manipulate MRI or CT images, including single-handed gestures for moving forward or backward through images, and two-handed gestures for zooming and panning (see Figure 7.20).



Figure 7.20 Touchless gesturing in the operating theater

Source: Used courtesy of Kenton O'Hara

Research and Design Considerations

A key design concern for using gestural input is to consider how a computer system recognizes and delineates the user's gestures. In particular, how does it determine the start and end point of a hand or arm movement, and how does it know the difference between a deictic gesture (a deliberate pointing movement) and hand waving (an unconscious gesticulation) that is used to emphasize what is being said verbally?

In addition to being used as a form of input, gestures can be represented as output to show real-time avatar movement or someone's own arm movements. Smartphones, laptops, and some smart speakers (for example, Facebook's Portal) have cameras that can perceive three dimensions and record a depth for every pixel. This can be used to create a representation of someone in a scene, for example, how they are posing and moving, and also to respond to their gestures. One design question that this raises is how realistic must the mirrored graphical representation of the user be in order for them to be believable and for the user to connect their gestures with what they are seeing on the screen.

7.2.12 Haptic Interfaces

Haptic interfaces provide tactile feedback, by applying vibration and forces to the person, using actuators that are embedded in their clothing or a device that they are carrying, such as a smartphone or smartwatch. Gaming consoles have also employed vibration to enrich the experience. For example, car steering wheels that are used with driving simulators can vibrate in various ways to provide the feel of the road. As the driver makes a turn, the steering wheel can be programmed to feel like it is resisting—in the way that a real steering wheel does.

Vibrotactile feedback can also be used to simulate the sense of touch between remote people who want to communicate. Actuators embedded in clothing can be designed to re-create the sensation of a hug or a squeeze by being buzzed on various parts of the body. Another use of haptics is to provide real-time feedback to guide people when learning a musical instrument, such as a violin or drums. For example, the MusicJacket (van der Linden et al., 2011) was developed to help novice violin players learn how to hold their instrument correctly and develop good bowing action. Vibrotactile feedback was provided via the jacket to give nudges at key places on the arm and torso to inform the student when they were either holding their violin incorrectly or their bowing trajectory had deviated from a desired path (see Figure 7.21). A user study with novice players showed that they were able to react to the vibrotactile feedback and adjust their bowing or their posture in response.

Another form of feedback is called *ultrahaptics*, which creates the illusion of touch in midair. It does this by using ultrasound to make three-dimensional shapes and textures that can be felt but not seen by the user (www.ultrahaptics.com). This technique can be used to create the illusion of having buttons and sliders that appear in midair. One potential use is in the automotive industry to replace existing physical buttons and knobs or touchscreens. The ultra-haptic buttons and knobs can be designed to appear next to the driver when needed, for example, when detecting the driver wants to turn down the volume or change the radio station.



Figure 7.21 The MusicJacket with embedded actuators that nudge the player to move their arm up to be in the correct position

Source: Helen Sharp

Haptics are also being embedded into clothing, sometimes called *exoskeletons*. Inspired by the “right trousers” in the Wallace and Gromit animated short movie, Jonathan Rossiter and his team (2018) developed a new kind of exoskeleton that can help people stand up and move around using artificial muscles that consist of air bubbles which are activated using tiny electric motors (see Figure 7.22). These are stiffened or relaxed using graphene parts to make the trousers move. One application area is to help people who have walking difficulties and those who need to exercise but find it difficult to do so.

7.2.13 Multimodal Interfaces

Multimodal interfaces are intended to enrich user experiences by multiplying the way information is experienced and controlled at the interface through using different modalities, such as touch, sight, sound, and speech (Bouchet and Nigay, 2004). Interface techniques that have been combined for this purpose include speech and gesture, eye-gaze and gesture, haptic and audio output, and pen input and speech (Dumas et al., 2009). The assumption is that multimodal interfaces can support more flexible, efficient, and expressive means of human–computer interaction that are more akin to the multimodal experiences that humans encounter in the physical world (Oviatt, 2017). Different input/outputs may be used at the same time, for example, using voice commands and gestures simultaneously to move through a virtual environment, or alternately using speech commands followed by gesturing. The most common combination of technologies used for multimodal interfaces is speech and vision processing (Deng and Huang, 2004). Multimodal interfaces can also be combined with multisensor input to enable other aspects of the human body to be tracked. For example, eye gaze, facial expressions, and lip movements can also be tracked to provide data about a user’s

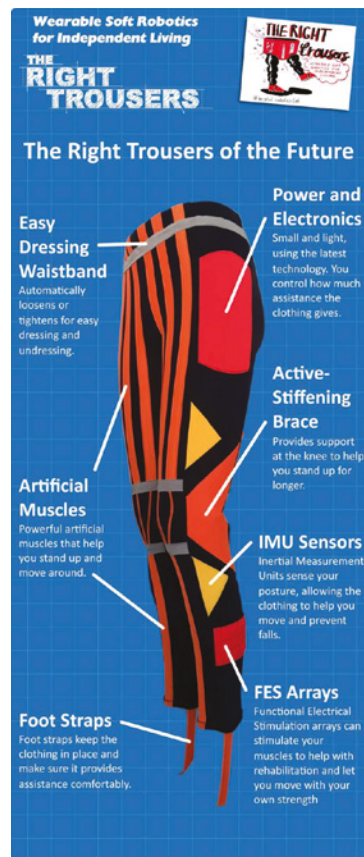


Figure 7.22 Trousers with artificial muscles that use a new kind of bubble haptic feedback
 Source: Used courtesy of The Right Trousers Project: Wearable Soft Robotics for Independent Living

Research and Design Considerations

Haptics are now commonly used in gaming consoles, smartphones, and controllers to alert or heighten a user experience. Haptic feedback is also being developed in clothing and other wearables as a way of simulating being touched, stroked, prodded, or buzzed. A promising application area is sensory-motor skills, such as in sports training and learning to play a musical instrument. For example, patterns of vibrations have been placed across snowboarders' bodies to indicate which moves to take while snowboarding. A study reported faster reaction times than when the same instructions were given verbally (Spelmezan et al., 2009). Other uses are posture trainers that buzz when a user slouches and fitness trackers that also buzz when they detect that their users have not taken enough steps in the past hour.

(Continued)

A key design question is where best to place the actuators on the body, whether to use a single or a sequence of touches, when to activate, and at what intensity and how often to use them to make the feeling of being touched convincing (e.g., Jones and Sarter, 2008). Providing continuous haptic feedback would be simply too annoying. People would also habituate too quickly to the feedback. Intermittent buzzes can be effective at key moments when a person needs to attend to something but not necessarily tell them what to do. For example, a study by Johnson et al. (2010) of a commercially available haptic device, intended to improve posture by giving people a vibrotactile buzz when they slouched, found that while the buzzing did not show them how to improve their posture, it did improve their body awareness.

Different kinds of buzzes can also be used to indicate different tactile experiences that map to events; for example, a smartphone could transmit feelings of slow tapping to feel like water dropping, which is meant to indicate it is about to rain and transmit the sensation of heavy tapping to indicate a thunderstorm is looming.

attention or other behavior. This kind of sensing can provide input for customizing user interfaces and experiences to the perceived need, desire, or level of interest.

A person's body movement can also be tracked so that it can be represented back to them on a screen in the form of an avatar that appears to move just like them. For example, the Kinect was developed as a gesture and body movement gaming input system for the Xbox. Although now defunct in the gaming industry, it proved effective at detecting multimodal input in real time. It consisted of an RGB camera for facial and gesture recognition, a depth sensor (an infrared projector paired with a monochrome camera) for movement tracking, and downward-facing mics for voice recognition (see Figure 7.23). The Kinect looked for



Figure 7.23 Microsoft's Xbox Kinect

Source: Stephen Brashear / Invision for Microsoft / AP Images

someone's body. On finding it, it locked onto it and measured the three-dimensional positioning of the key joints in their body. This information was converted into a graphical avatar of the user that could be programmed to move just like them. Many people readily saw themselves as the avatar and learnt how to play games in this manner.

Research and Design Considerations

Multimodal systems rely on recognizing aspects of a user's behavior, including handwriting, speech, gestures, eye movements, or other body movements. In many ways, this is much harder to accomplish and calibrate than single modality systems that are programmed to recognize one aspect of a user's behavior. The most researched modes of interaction are speech, gesture, and eye-gaze tracking. A key research question is what is actually gained from combining different input and outputs and whether talking and gesturing as humans do with other humans is a natural way of interacting with a computer (see Chapter 4). Guidelines for multimodal design can be found in Reeves et al. (2004) and Oviatt et al. (2017).

7.2.14 Shareable Interfaces

Shareable interfaces are designed for more than one person to use. Unlike PCs, laptops, and mobile devices, which are aimed at single users, shareable interfaces typically provide multiple inputs and sometimes allow simultaneous input by collocated groups. These include large wall displays, for example SmartBoards (see Figure 7.24a), where people use their own pens or gestures, and interactive tabletops, where small groups can interact with information being displayed on the surface using their fingertips. Examples of interactive tabletops include Smart's SmartTable and Circle Twelve's DiamondTouch (Dietz and Leigh, 2001; see Figure 7.24b). The DiamondTouch tabletop is unique in that it can distinguish between different users touching the surface concurrently. An array of antennae is embedded in the touch surface and each one transmits a unique signal. Each user has their own receiver embedded in a mat on which they're standing or a chair in which they're sitting. When a user touches the tabletop, very small signals are sent through the user's body to their receiver that identifies which antenna has been touched and sends this to the computer. Multiple users can interact simultaneously with digital content using their fingertips.

Watch this video of Circle Twelve's demonstration of the Diamond Touch tabletop:
<http://youtu.be/S9QRdXITndU>.

An advantage of shareable interfaces is that they provide a large interactional space that can support flexible group working, enabling groups to create content together at the same time. Compared with a co-located group trying to work around a single-user PC or laptop, where typically one person takes control, making it more difficult for others to take part,



(a)



(b)

Figure 7.24 (a) A SmartBoard in use during a meeting and (b) Mitsubishi's interactive tabletop interface

Source: (a) Used courtesy of SMART Technologies Inc. (b) Mitsubishi Electric Research Labs

multiple users can interact with large display. Users can point to and touch the information being displayed, while simultaneously viewing the interactions and having the same shared point of reference (Rogers et al., 2009). There are now a number of tabletop apps that have been developed for museums and galleries which enable visitors to learn about various aspects of the environment (see Clegg et al., 2019).

Another type of shareable interface is software platforms that enable groups of people to work together simultaneously even when geographically apart. Early examples included shared editing tools developed in the 1980s (for example, ShRedit). Various commercial products now exist that enable multiple remote people to work on the same document at the same time (such as Google Docs and Microsoft Excel). Some enable up to 50 people to edit the same document at the same time with more watching on. These software programs provide various functions, such as synchronous editing, tracking changes, annotating, and commenting. Another collaborative tool is the Balsamiq Wireframes editor, which provides a range of shared functions, including collaborative editing, threaded comments with callouts, and project history.

Research and Design Considerations

Early research on shareable interfaces focused largely on interactional issues, such as how to support electronically based handwriting and drawing, and the selecting and moving of objects around the display (Elrod et al., 1992). The PARCTAB system (Schilit et al., 1993) investigated how information could be communicated between palm-sized, A4-sized, and whiteboard-sized displays using shared software tools, such as Tivoli (Rønby-Pedersen et al., 1993). Another concern was how to develop fluid and direct styles of interaction with large displays, both wall-based and tabletop, involving freehand and pen-based gestures (see Shen et al., 2003). Current research is concerned with how to support ecologies of devices so that groups can share and create content across multiple devices, such as tabletops and wall displays (see Brudy et al., 2016).

A key research issue is whether shareable surfaces can facilitate new and enhanced forms of collaborative interaction compared with what is possible when groups work together using their own devices, like laptops and PCs (see Chapter 5, “Social Interaction”). One benefit is easier sharing and more equitable participation. For example, tabletops have been designed to support more effective joint browsing, sharing, and manipulation of images during decision-making and design activities (Shen et al., 2002; Yuill and Rogers, 2012). Core design concerns include whether size, orientation, and shape of the display have an effect on collaboration. User studies have shown that horizontal surfaces compared with vertical ones support more turn-taking and collaborative working in co-located groups (Rogers and Lindley, 2004), while providing larger-sized tabletops does not necessarily improve group working but can encourage a greater division of labor (Ryall et al., 2004).

The need for both personal and shared spaces has been investigated to see how best to enable users to move between working on their own and together as a group. Several researchers have designed cross-device systems, where a variety of devices, such as tablets, smartphones, and digital pens can be used in conjunction with a shareable surface. For example, SurfaceConstellations was developed for linking mobile devices to create novel cross-device workspace environments (Marquardt et al., 2018). Design guidelines and summaries of empirical research on tabletops and multitouch devices can be found in Müller-Tomfelde (2010).

7.2.15 Tangible Interfaces

Tangible interfaces use sensor-based interaction, where physical objects, such as bricks, balls, and cubes, are coupled with digital representations (Ishii and Ullmer, 1997). When a person manipulates the physical object(s), it is detected by a computer system via the sensing mechanism embedded in the physical object, causing a digital effect to occur, such as a sound, animation, or vibration (Fishkin, 2004). The digital effects can take place in a number of media and places, or they can be embedded in the physical object itself. For example, Oren Zuckerman and Mitchel Resnick's (2005) early Flow Blocks prototype depicted changing numbers and lights that were embedded in the blocks, depending on how they were connected. The flow blocks were designed to simulate real-life dynamic behavior and react when arranged in certain sequences.

Another type of tangible interface is where a physical model, for example, a puck, a piece of clay, or a model, is superimposed on a digital desktop. Moving one of the physical pieces around the tabletop causes digital events to take place on the tabletop. One of the earliest tangible interfaces, Urp, was built to facilitate urban planning; miniature physical models of buildings could be moved around on the tabletop and used in combination with tokens for wind and shadow-generating tools, causing digital shadows surrounding them to change over time and visualizations of airflow to vary. Tangible interfaces differ from other approaches, such as mobile, insofar as the representations are artifacts in their own right that the user can directly act upon, lift up, rearrange, sort, and manipulate.

The technologies that have been used to create tangibles include RFID tags and sensors embedded in physical objects and digital tabletops that sense the movements of objects and subsequently provide visualizations surrounding the physical objects. Many tangible systems have been built with the goal of encouraging learning, design activities, playfulness, and collaboration. These include planning tools for landscape and urban planning (see Hornecker, 2005; Underkoffler and Ishii, 1998). Another example is Tinkersheets, which combine tangible models of shelving with paper forms for exploring and solving warehouse logistics problems (Zufferey, et al., 2009). The underlying simulation allows students to set parameters by placing small magnets on the form.

Tangible computing has been described as having no single locus of control or interaction (Dourish, 2001). Instead of just one input device, such as a mouse, there is a coordinated interplay of different devices and objects. There is also no enforced sequencing of actions and no modal interaction. Moreover, the design of the interface objects exploits their affordances to guide the user in how to interact with them. A benefit of tangibility is that physical objects and digital representations can be positioned, combined, and explored in creative ways, enabling dynamic information to be presented in different ways. Physical objects can also be held in both hands and combined and manipulated in ways not possible using other interfaces. This allows for more than one person to explore the interface together and for objects to be placed on top of each other, beside each other, and inside each other; the different configurations encourage different ways of representing and exploring a problem space. In so doing, people are able to see and understand situations differently, which can lead to greater insight, learning, and problem-solving than with other kinds of interfaces (Marshall et al., 2003).

A number of toolkits have been developed to encourage children to learn coding, electronics, and STEM subjects. These include littleBits (<https://littlebits.com/>), MicroBit (<https://microbit.org/>), and MagicCubes (<https://uclmagiccube.weebly.com/>). The toolkits

provide children with opportunities to connect physical electronic components and sensors to make digital events occur. For example, the MagicCubes can be programmed to change color depending on the speed at which they are shaken; slow is blue and very fast is multicolor. Research has shown that the tangible toolkits provide many opportunities for discovery learning, exploration, and collaboration (Lechelt et al., 2018). The Cubes have been found to encourage a diversity of children, between the ages of 6 and 16, and those with cognitive disabilities, to learn through collaborating, frequently showing and telling each other and their instructors about their discoveries. These moments are facilitated by the cube's form factor, making it easy to show off to others, for example, by waving a cube in the air (see Figure 7.25).



Figure 7.25 Learning to code with the MagicCubes; sharing, showing, and telling
Source: Elpida Makriyannis

Tangible toolkits have also been developed for the visually impaired. For example, Torino (renamed by Microsoft to Code Jumper) was developed as a programming language for teaching programming concepts to children age 7–11, regardless of level of vision (Morrison et al., 2018). It consists of a set of beads that can be connected and manipulated to create physical strings of code that play stories or music.

BOX 7.5

VoxBox—A Tangible Questionnaire Machine

Traditional methods for gathering public opinions, such as surveys, involve approaching people *in situ*, but it can disrupt the positive experience they are having. VoxBox (see Figure 7.26) is a tangible system designed to gather opinions on a range of topics *in situ* at an event through playful and engaging interaction (Golsteijn et al., 2015). It is intended to encourage wider participation by grouping similar questions, encouraging completion, gathering answers to open and closed questions, and connecting answers and results. It was designed as a large physical system that provides a range of tangible input mechanisms through which people give their opinions, instead of using, for example, text messages or social media input. The various input mechanisms include sliders, buttons, knobs, and spinners about which people are all familiar. In addition, the system has a transparent tube at the side that drops a ball step by step as sets of questions are completed to act as an incentive for completion and as a progress indicator. The results of the selections are aggregated and presented as simple digital visualizations on the other side (for example, 95 percent are engaged; 5 percent are bored). VoxBox has been used at a number of events drawing in the crowds, who become completely absorbed in answering questions in this tangible format.



Figure 7.26 VoxBox—front and back of the tangible machine questionnaire

Source: Yvonne Rogers ■

Research and Design Considerations

Researchers have developed conceptual frameworks that identify the novel and specific features of a tangible interface (see Fishkin, 2004; Ullmar et al., 2005; Shaer and Hornecker, 2010). A key design concern is what kind of coupling to use between the physical action and digital effect. This includes determining where the digital feedback is provided in relation to the physical artifact that has been manipulated. For example, should it appear on top of the object, beside it, or in some other place? The type and placement of the digital media will depend to a large extent on the purpose of using a tangible interface. If it is to support learning, then an explicit mapping between action and effect is critical. In contrast, if it is for entertainment purposes, for example, playing music or storytelling, then it may be better to design them to be more implicit and unexpected. Another key design question is what kind of physical artifact to use to enable the user to carry out an activity in a natural way. Bricks, cubes, and other component sets are most commonly used because of their flexibility and simplicity, enabling people to hold them in both hands and to construct new structures that can be easily added to or changed. Sticky notes and cardboard tokens can also be used for placing material onto a surface that is transformed or attached to digital content (Klemmer et al. 2001; Rogers et al., 2006).

Another research question is with what types of digital outputs should tangible interfaces be combined? Overlaying physical objects with graphical feedback that changes in response to how the object is manipulated has been the main approach. In addition, audio and haptic feedback has also been used. Tangibles can also be designed to be an integral part of a multi-modal interface.

7.2.16 Augmented Reality

Augmented reality (AR) became an overnight success with the arrival of Pokémon Go in 2016. The smartphone app became an instant hit worldwide. Using a player's smartphone camera and GPS signal, the AR game makes it seem as if virtual Pokémon characters are appearing in the real world—popping up all over the place, such as on buildings, on streets, and in parks. As players walk around a given place, they may be greeted with rustling bits of grass that signal a Pokémon nearby. If they walk closer, a Pokémon may pop up on their smartphone screen, as if by magic, and look as if they are actually in front of them. For example, one might be spotted sitting on a branch of a tree or a garden fence.

AR works by superimposing digital elements, like Pokémon, onto physical devices and objects. Closely related to AR is the concept of *mixed reality*, where views of the real world are combined with views of a virtual environment (Drascic and Milgram, 1996). To begin, augmented reality was mostly a subject of experimentation within medicine, where virtual objects, for example X-rays and scans, were overlaid on part of a patient's body to aid the physician's understanding of what was being examined or operated on.

AR was then used to aid controllers and operators in rapid decision-making. One example is air traffic control, where controllers are provided with dynamic information about the aircraft in their section that is overlaid on a video screen showing real planes landing, taking off, and taxiing. The additional information enables the controllers to identify planes easily,

which were difficult to make out—something especially useful in poor weather conditions. Similarly, *head-up displays* (HUDs) are used in military and civil planes to aid pilots when landing during poor weather conditions. A HUD provides electronic directional markers on a fold-down display that appears directly in the field of view of the pilot. A number of high-end cars now provide AR windshield technology, where navigation directions can literally look like they are painted on the road ahead of the driver (see Chapter 2, “The Process of Interaction Design”).

Instructions for building or repairing complex equipment, such as photocopiers and car engines, have also been designed to replace paper-based manuals, where drawings are superimposed upon the machinery itself, telling the mechanic what to do and where to do it. There are also many AR apps available now for a range of contexts, from education to car navigation, where digital content is overlaid on geographic locations and objects. To reveal the digital information, users open the AR app on a smartphone or tablet and the content appears superimposed on what is viewed through the screen.

Other AR apps have been developed to aid people walking in a city or town. Directions (in the form of a pointing hand or arrow) and local information (for instance, the nearest bakery) are overlaid on the image of the street ahead that appears on someone’s smartphone screen. These change as the person walks up the street. Virtual objects and information are also being combined to make more complex augmented realities. Figure 7.27 shows a weather alert with animated virtual lightning effects alongside information about a nearby café and the price of properties for sale or rent on a street. Holograms of people and other objects are also being introduced into AR environments that can appear to move



Figure 7.27 Augmented reality overlay used on a car windshield

Source: <https://wayray.com>

and/or talk. For example, virtual tour guides are beginning to appear in museums, cities, and theme parks, which can appear to by moving, talking, or gesturing to visitors who are using an AR app.

The availability of mapping platforms, such as those provided by Niantics and Google, together with Apple's ARKit, SparkAR Studio, and Google's ARCore, has made it easier for developers and students alike to develop new kinds of AR games and AR apps. Another popular AR game that has emerged since Pokémon Go is Jurassic World Alive, where players walk around in the real world to find as many virtual dinosaurs as they can. It is similar to Pokémon Go but with different gaming mechanisms. For example, players have to study the dinosaurs they come across by collecting their DNA and then re-creating it. Microsoft's HoloLens toolkit has also enabled new mixed reality user experiences to be created, allowing users to create or interact with virtual elements in their surroundings.

Most AR apps use the backward-facing camera on a smartphone or tablet to overlay the virtual content onto the real world. Another approach is to use the forward-facing camera (used for selfies) to superimpose digital content onto the user's face or body. The most popular app that has used this technique is SnapChat, which provides numerous filters with which people can experiment plus the opportunity to create their own filters. Adding accessories such as ears, hair, moving lips, and headgear enables people to transform their physical appearance in all sorts of fun ways.

These kinds of virtual try-ons work by analyzing the user's facial features and building a 2-D or 3-D model in real time. So, when they move their head, the make-up or accessories appear to move with them as if they are really on their face. Several AR mirrors now exist in retail that allow shoppers to try on sunglasses, jewelry, and make-up. The goal is to let them try on as many different products as they like to see how they look with them on. Clearly, there are advantages to virtual try-ons: it can be more convenient, engaging, and easier compared to trying on the real thing. There are disadvantages too, however, in that they only give an impression of what you look like. For example, the user cannot feel the weight of a virtual accessory on their head or the texture of virtual make-up on their face.

The same technology can be used to enable people to step into historical, famous, film, or stage characters (for instance, David Bowie or Queen Victoria). For example, a virtual try-on app that was developed as part of a cultural experience was the Magic-Face (Javornik, et al., 2017). The goal was to enable audiences to experience firsthand what it was like to try on the make-up of a character from an opera. The opera chosen was Philip Glass's *Akhnaten*, set in ancient Egyptian time (see Figure 7.28a). The virtual make-up developed were for a Pharaoh and his wife. The app was developed by University College London researchers alongside the English National Opera and AR company, Holition. To provide a real-world context, the app was designed to run on a tablet display that was disguised as a real mirror and placed in an actor's dressing room (see Figure 7.28b). On encountering the mirror *in situ*, visiting school children were fascinated by the way the virtual make-up made them look like Akhnaten and his wife, Nefertiti. The singers and make-up artists who were in the production also tried it out and saw great potential for using the app to enhance their existing repertoire of rehearsal and make-up tools.



Figure 7.28 (a) A principal singer trying on the virtual look of Akhnaten and (b) a framed AR mirror in the ENO dressing room

Source: Used courtesy of Ana Javornik

Research and Design Considerations

A key research concern when designing augmented reality is what form the digital augmentation should take and when and where it should appear in the physical environment (Rogers et al., 2005). The information (such as navigation cues) needs to stand out but not distract the person from their ongoing activity in the physical world. It also needs to be simple and align with the real-world objects taking into account that the user will be moving. Another concern is how much digital content to overlay on the physical world and how to attract the user's attention to it. There is the danger that the physical world becomes overloaded with digital ads and information “polluting” it to the extent that people will turn the AR app off.

One of the limitations of current AR technology is that sometimes the modeling can be slightly off so that the overlaying of the digital information appears in the wrong place or is out of sync with what is being overlaid. This may not be critical for fun applications, but it may be disconcerting if eye shadow appears on someone's ear. It may also break the magic of the AR experience. Ambiguity and uncertainty may be exploited to good effect in mixed reality games, but it could be disastrous in a more serious context, such as the military or medical setting.

7.2.17 Wearables

Wearables are a broad category of devices that are worn on the body. These include smartwatches, fitness trackers, fashion tech, and smart glasses. Since the early experimental days of wearable computing, where Steve Mann (1997) donned head and eye cameras to enable him to record what he saw while also accessing digital information on the move, there have been many innovations and inventions, including Google Glass.

New flexible display technologies, e-textiles, and physical computing (for example, Arduino) provide opportunities to design wearables that people will actually want to wear. Jewelry, caps, glasses, shoes, and jackets have all been the subject of experimentation designed to provide the user with a means of interacting with digital information while on the move in the physical world. Early wearables focused on convenience, enabling people to carry out a task (for example, selecting music) without having to take out and control a handheld device. Examples included a ski jacket with integrated music player controls that enabled the wearer to simply touch a button on their arm with their glove to change a music track. More recent applications have focused on how to combine textiles, electronics, and haptic technologies to promote new forms of communication. For example, CuteCircuit developed the KineticDress, which was embedded with sensors that followed the body of the wearer to capture their movements and interaction with others. These were then displayed through electroluminescent embroidery that covered the external skirt section of the dress. Depending on the amount and speed of the wearer's movement, it changed patterns, displaying the wearer's mood to the audience and creating a magic halo around her.

Exoskeleton clothing (see section 7.2.12) is also an area where fashion meets technology in order to augment and assist people who have problems with walking by literally walking or exercising the person wearing them. In this way, it combines haptics with a wearable. Within the construction industry, exoskeleton suits have also been developed to provide additional power to workers—a bit like Superman—where metal frameworks are fitted with motorized muscles to multiply the wearer's strength. It can make lifting objects feel lighter and in doing so protect the worker from physical injuries.

DILEMMA

Google Glass: Seeing Too Much?

Google Glass was a wearable that went on sale in 2014 in various fashion styles (see Figure 7.29). It was designed to look like a pair of glasses, but with one lens of the glass being an interactive display with an embedded camera that could be controlled with speech input. It allowed the wearer to take photos and videos on the move and look at digital content, such as email, texts, and maps. The wearer could also search the web using voice commands, and the results would appear on the screen. A number of applications were developed beyond those for everyday use, including WatchMeTalk, which provided live captions to help the

(Continued)



Figure 7.29 Google Glass

Source: Google Inc.

hearing-impaired in their day-to-day conversations and Preview for Glass that enabled the wearer to watch a movie trailer the moment they looked at a movie poster.

However, being in the company of someone wearing a Google Glass was felt by many to be unnerving, as the wearer looked up and to the right to view what was on the glass screen rather than looking at you and into your eyes. One of the criticisms of wearers of Google Glass was that it made them appear to be staring into the distance. Others were worried that those wearing Google Glass were recording everything that was happening in front of them. As a reaction, a few bars and restaurants in the United States implemented a “no Glass” policy to prevent customers from recording other patrons.

The original Google Glass was retired after a couple of years. Since then, other types of smart glasses have come onto the market that synch a user’s smartphone with the display and camera on the glasses via Bluetooth. These include Vuzic Blade, which has a camera onboard and voice control that is connected to Amazon Echo devices, along with the provision of turn-by-turn navigation and location-based alerts; and Snap’s Spectacles, which simply allows the wearer to share photos and video with their friends on Snapchat they take when wearing these glasses. ■

Watch the interesting video of London through Google Glass at <http://youtu.be/Z3AIdnzZUsE> and the Talking Shoe concept at <http://youtu.be/VcaSwxbRkcE>.

Research and Design Considerations

A core design concern specific to wearable interfaces is comfort. Users need to feel comfortable wearing clothing that is embedded with technology. It needs to be light, small, not get in the way, fashionable, and (with the exception of the displays) preferably hidden in the clothing. Another related issue is hygiene. Is it possible to wash or clean the clothing once worn? How easy is it to remove the electronic gadgetry and replace it? Where are the batteries going to be placed, and how long is their lifetime? A key usability concern is how does the user control the devices that are embedded in their clothing. Are touch, speech, or more conventional buttons and dials preferable?

A number of technologies can be developed and combined to create wearables including LEDs, sensors, actuators, tangibles, and AR. There is abundant scope for thinking creatively about when and whether to make something wearable as opposed to mobile. In Chapter 1, “What Is Interaction Design?” we mentioned how assistive technology can be designed to be fashionable in order to overcome stigmas of having to wear a monitoring device (for instance, for glucose levels), substitution (for example, a prosthetic) or amplifying device (for example, hearing aids).

7.2.18 Robots and Drones

Robots have been around for some time, most notably as characters in science-fiction movies, but they also play an important role as part of manufacturing assembly lines, as remote investigators of hazardous locations (for example, nuclear power stations and bomb disposal), and as search and rescue helpers in disasters (for instance in forest fires) or faraway places (like Mars). *Console interfaces* have been developed to enable humans to control and navigate robots in remote terrains, using a combination of joysticks and keyboard controls together with cameras and sensor-based interactions (Baker et al., 2004). The focus has been on designing interfaces that enable users to steer and move a remote robot effectively with the aid of live video and dynamic maps.

Domestic robots that help with the cleaning and gardening have become popular. Robots are also being developed to help the elderly and disabled with certain activities, such as picking up objects and cooking meals. Pet robots, in the guise of human companions, have been commercialized. Several research teams have taken the “cute and cuddly” approach to designing robots, signaling to humans that the robots are more pet-like than human-like. For example, Mitsubishi developed Mel the penguin (Sidner and Lee, 2005) whose role was to host events, while the Japanese inventor Takanori Shibata developed Paro in 2004, a baby harp seal that looks like a cute furry cartoon animal, and whose role was as a companion (see Figure 7.30). Sensors were embedded in the pet robots, enabling them to detect certain human behaviors and respond accordingly. For example, they can open, close, and move their eyes, giggle, and raise their flippers. The robots encourage being cuddled or spoken to,

as if they were real pets or animals. The appeal of pet robots is thought to be partially due to their therapeutic qualities, being able to reduce stress and loneliness among the elderly and infirm (see Chapter 6, “Emotional Interaction,” for more on cuddly robot pets). Paro has since been used to help patients with dementia to make them feel more at ease and comforted (Griffiths, 2014). Specifically, it has been used to encourage social behavior among patients who often anthropomorphize it. For example, they might say as a joke “it’s farted on me!” which makes them and others around them laugh, leading to further laughter and joking. This form of encouraging of social interaction is thought to be therapeutic.



Figure 7.30 (a) Mel, the penguin robot, designed to host activities; (b) Japan’s Paro, an interactive seal, designed as a companion, primarily for the elderly and sick children

Source: (a) Mitsubishi Electric Research Labs (b) Parorobots.com

Watch the video of Robot Pets of the Future at <http://youtu.be/wBFws1lhuv0>.

Drones are a form of unmanned aircraft that are controlled remotely. They were first used by hobbyists and then by the military. Since then, they have become more affordable, accessible, and easier to fly. As a result, they have begun to be used in a wider range of contexts. These include entertainment, such as carrying drinks and food to people at festivals and parties; agricultural applications, such as flying them over vineyards and fields to collect data that is useful to farmers (see Figure 7.31); and helping to track poachers in wildlife parks in Africa (Preece, 2016). Compared with other forms of data collection, they can fly low and stream photos to a ground station where the images can be stitched together into maps and then used to determine the health of a crop or when is the best time to harvest the crop.

Watch the video of Rakuten delivering beer via drone to golfers on a golf course at <https://youtu.be/ZameOVS2Skw>.



Figure 7.31 A drone being used to survey the state of a vineyard

Source: Drone inspecting vineyard / Shutterstock

Research and Design Considerations

An ethical concern is whether it is acceptable to create robots that exhibit behaviors that humans will consider to be human- or animal-like. While this form of attribution also occurs for agent interfaces (see Chapter 3), having a physical embodiment—as robots do—can make people suspend their disbelief even more, viewing the robots as pets or humans.


This raises the moral question as to whether such anthropomorphism should be encouraged. Should robots be designed to be as human-like as possible, looking like us with human features, such as eyes and a mouth, behaving like us, communicating like us, and emotionally responding like us? Or, should they be designed to look like robots and behave like robots, for instance, vacuum cleaner robots that serve a clearly defined purpose? Likewise, should the interaction be designed to enable people to interact with the robot as if it were another human being, for example, by talking to it, gesturing at it, holding its hand, and smiling at it? Or, should the interaction be designed to be more like human–computer interaction, in other words, by pressing buttons, knobs, and dials to issue commands?

For many people, the cute pet approach to robotic interfaces seems preferable to one that seeks to design them to be more like fully fledged human beings. Humans know where they stand with pets and are less likely to be unnerved by them and, paradoxically, are more likely to suspend their disbelief in the companionship they provide.

Another ethical concern is whether it is acceptable to use unmanned drones to take a series of images or videos of fields, towns, and private property without permission or people knowing what is happening. They are banned from certain areas such as airports, where they present a real danger. Another potential problem is the noise they make when flying. Having a drone constantly buzzing past your house or delivering drinks to golf players or festival goers nearby can be very annoying.

7.2.19 Brain–Computer Interfaces

Brain–computer interfaces (BCI) provide a communication pathway between a person’s brain waves and an external device, such as a cursor on a screen or a tangible puck that moves via airflow. The person is trained to concentrate on the task (for example, moving the cursor or the puck). Several research projects have investigated how this technique can be used to assist and augment human cognitive or sensory-motor functions. The way BCIs work is by detecting changes in the neural functioning of the brain. Our brains are filled with neurons that comprise individual nerve cells connected to one another by dendrites and axons. Every time we think, move, feel, or remember something, these neurons become active. Small electric signals rapidly move from neuron to neuron, which to a certain extent can be detected by electrodes that are placed on a person’s scalp. The electrodes are embedded in specialized headsets, hairnets, or caps.



"Frankly, I'm not sure this whole idea-sharing thing is working."

Source: Tim Cordell / Cartoon Stock

Brain–computer interfaces have also been developed to control various games. For example, Brainball was developed as a game to be controlled by players’ brain waves in which they compete to control a ball’s movement across a table by becoming more relaxed and focused. Other possibilities include controlling a robot and being able to fly a virtual plane. Pioneering medical research, conducted by the BrainGate research group at Brown University, has started using brain-computer interfaces to enable people who are paralyzed to control robots (see Figure 7.32). For example, a robotic arm controlled by a tethered BCI has enabled patients who are paralyzed to feed themselves (see video mentioned next). Another startup company, NextMind, is developing a noninvasive brain-sensing device intended for the mass market to enable users to play games and control electronic and mobile devices in real time using just their thoughts. It is researching how to combine brain-sensing technology with innovative machine-learning algorithms that can translate brain waves into digital commands.

Watch a video of a woman who is paralyzed moving a robot with her mind at <http://youtu.be/ogBX18maUiM>.




Figure 7.32 A brain-computer interface being used by a woman who is paralyzed to select letters on a screen (developed by BrownGate)

Source: Brown University

7.2.20 Smart Interfaces

The motivation for many new technologies is to make them smart, whether it is a smartphone, smartwatch, smart building, smart home, or smart appliance (for example smart lighting, smart speakers, or virtual assistants). The adjective is often used to suggest that the device has some intelligence and it is connected to the Internet. More generally, smart devices are designed to interact with users and other devices connected to a network, many of which are automated, not requiring users to interact with them directly (Silverio-Fernández et al., 2018). The goal is to make them context-aware, that is, to understand what is happening around them and execute appropriate actions. To achieve this, some have been programmed with AI so that they can learn the context and a user's behavior. Using this intelligence, they then change settings or switch things on according to the user's assumed preferences. An example is the smart Nest thermostat that is designed to learn from a householder's behavior. Rather than make the interface invisible, the designers chose to turn it into an aesthetically pleasing one that could be easily viewed (see Box 6.2).

Smart buildings have been designed to be more energy efficient, efficient, and cost effective. Architects are motivated to use state-of-the-art sensor technology to control building systems, such as ventilation, lighting, security, and heating. Often, the inhabitants of such buildings are considered to be the ones at fault for wasting energy, as they may leave the lights and heating on overnight when not needed, or they forget to lock a door or window. One benefit of having automated systems take control of building services is to reduce these kinds of human errors—a phrase often used by engineers is to take the human “out of the loop.” While some smart buildings and homes have improved how they are managed and cut costs, they can also be frustrating to the user, who sometimes would like to be able to open a window to let fresh air in or raise a blind to let in natural lighting. Taking the human out of

the loop means that these operations are no longer available. Windows are locked or sealed, and heating is controlled centrally.

Instead of simply introducing ever more automation that takes the human out of the loop further, another approach is to consider the needs of the inhabitants in conjunction with introducing smart technology. For example, a new approach that focuses on inhabitants is called *human–building interaction* (HBI). It is concerned with understanding and shaping people’s experiences with, and within, built environments (Alavi et al., 2019). The focus is on human values, needs, and priorities in addressing people’s interactions with “smart” environments.

7.3 Natural User Interfaces and Beyond

As we have seen, there are many kinds of interfaces that can be used to design for user experiences. The staple for many years was the GUI, then the mobile device interface, followed by touch, and now wearables and smart interfaces. Without question, they have been able to support all manner of user activities. What comes next? Will other kinds of interfaces that are projected to be more natural become more mainstream?

A *natural user interface* (NUI) is designed to allow people to interact with a computer in the same way that they interact with the physical world—using their voice, hands, and bodies. Instead of using a keyboard, mouse, or touchpad (as is the case with GUIs), NUIs enable users to speak to machines, stroke their surfaces, gesture at them in the air, dance on mats that detect feet movements, smile at them to get a reaction, and so on. The naturalness refers to the use of everyday skills humans have developed and learned, such as talking, writing, gesturing, walking, and picking up objects. In theory, they should be easier to learn and map more readily onto how people interact with the world than compared with learning to use a GUI.

Instead of having to remember which function keys to press to open a file, a NUI means a person only has to raise their arm or say “open.” But how natural are NUIs? Is it more natural to say “open” than to flick a switch when you want to open a door? And is it more natural to raise both arms to change a channel on the TV than to press a button on a remote device or tell it what to do by speaking to it? Whether a NUI is natural depends on a number of factors, including how much learning is required, the complexity of the app or device’s interface, and whether accuracy and speed are needed (Norman, 2010). Sometimes a gesture is worth a thousand words. Other times, a word is worth a thousand gestures. It depends on how many functions the system supports.

Consider the sensor-based faucets that were described in Chapter 1. The gesture-based interface works mostly (with the exception of people wearing black clothing that cannot be detected) because there are only two functions: (1) turning on the water by waving one’s hands under the tap, and (2) turning off the water by removing them from the sink. Now think about other functions that faucets usually provide, such as controlling water temperature and flow. What kind of a gesture would be most appropriate for changing the temperature and then the flow? Would one decide on the temperature first by raising one’s left arm and the flow by raising one’s right arm? How would someone know when to stop raising their arm to get the right temperature? Would they need to put a hand under the tap

to check? But if they put their right hand under the tap, that might have the effect of decreasing the flow? And when does the system know that the desired temperature and flow has been reached? Would it require having both arms suspended in midair for a few seconds to register that was the desired state? It is a difficult problem on how to provide these choices, and it is probably why sensor-based faucets in public bathrooms all have their temperature and flow set to a default.

Our overview of different interface types in this chapter has highlighted how gestural, voice, and other kinds of NUIs have made controlling input and interacting with digital content easier and more enjoyable, even though sometimes they can be less than perfect. For example, using gestures and whole-body movements have proven to be highly enjoyable as a form of input for computer games and physical exercises. Furthermore, new kinds of gesture, voice, and touch interfaces have made the web and online tools more accessible to those who are visually impaired. For example, the iPhone's VoiceOver control features have empowered visually impaired individuals to be able to easily send email, use the web, play music, and so on, without having to buy an expensive customized phone or screen reader. Moreover, being able to purchase a regular phone means not being singled out for special treatment. And while some gestures may feel cumbersome for sighted people to learn and use, they may not be so for blind or visually impaired people. The iPhone VoiceOver press and guess feature that reads out what you tap on the screen (for example, "messages," "calendar," "mail: 5 new items") can open up new ways of exploring an application while a three-finger tap can become a natural way to turn the screen off.

An emerging class of human-computer interfaces are those that rely largely on subtle, gradual, and continuous changes triggered by information obtained implicitly from the user together with the use of AI algorithms that are coded to learn about the user's behavior and preferences. These are connected with lightweight, ambient, context-aware, affective, and augmented cognition interfaces (Solovey et al., 2014). Using brain, body, behavioral, and environmental sensors, it is now possible to capture subtle changes in people's cognitive and emotional states in real time. This opens up new doors in human-computer interaction. In particular, it allows for information to be used as both continuous and discrete input, potentially enabling new outputs to match and be updated with what people might want and need at any given time. Adding AI to the mix will also enable a new type of interface to emerge that goes beyond simply being natural and smart—one that allows people to develop new superpowers that will enable them to work synergistically with technology to solve ever-more complex problems and undertake unimaginable feats.

7.4 Which Interface?

This chapter presented an overview of the diversity of interfaces that is now available or currently being researched. There are many opportunities to design for user experiences that are a far cry from those originally developed using the command-based interfaces of the 1980s. An obvious question this raises is which one and how do you design it? In many contexts, the requirements for the user experience that have been identified will determine what kind of interface might be appropriate and what features to include. For example, if a healthcare app is being developed to enable patients to monitor their dietary intake, then a

mobile device that has the ability to scan barcodes and/or take pictures of food items that can be compared with a database would be a good interface to use, enabling mobility, effective object recognition, and ease of use. If the goal is to design a work environment to support collocated group decision-making activities, then combining shareable technologies and personal devices that enable people to move fluidly among them would be good to consider using.

But how to decide which interface is preferable for a given task or activity? For example, is multimedia better than tangible interfaces for learning? Is voice effective as a command-based interface? Is a multimodal interface more effective than a single media interface? Are wearable interfaces better than mobile interfaces for helping people find information in foreign cities? How does VR differ from AR, and which is the ultimate interface for playing games? In what way are tangible environments more challenging and captivating than virtual worlds? Will shareable interfaces, such as interactive furniture, be better at supporting communication and collaboration compared with using networked desktop technologies? And so forth. These questions are currently being researched. In practice, which interface is most appropriate, most useful, most efficient, most engaging, most supportive, and so on will depend on the interplay of a number of factors, including reliability, social acceptability, privacy, ethical, and location concerns.

In-Depth Activity

Choose a game that you or someone you know plays a lot on a smartphone (for example, Candy Crush Saga, Fortnite, or Minecraft). Consider how the game could be played using different interfaces other than the smartphone's. Select three different interfaces (for instance, tangible, wearable, and shareable) and describe how the game could be redesigned for each of these, taking into account the user group being targeted. For example, the tangible game could be designed for children, the wearable interface for young adults, and the shareable interface for older people.

1. Go through the research and design considerations for each interface and consider whether they are relevant for the game setting and what considerations they raise.
2. Describe a hypothetical scenario of how the game would be played for each of the three interfaces.
3. Consider specific design issues that will need to be addressed. For example, for the shareable surface would it be best to have a tabletop or a wall-based surface? How will the users interact with the game elements for each of the different interfaces—by using a pen, fingertips, voice, or other input device? How do you turn a single-player game into a multiple player one? What rules would you need to add?
4. Compare the pros and cons of designing the game using the three different interfaces with respect to how it is played on the smartphone.

Summary

This chapter provided an overview of the diversity of interfaces that can be designed for user experiences, identifying key design issues and research questions that need to be addressed. It has highlighted the opportunities and challenges that lie ahead for designers and researchers who are experimenting with and developing innovative interfaces. It also explained some of the assumptions behind the benefits of different interfaces—some that are currently supported and others that are still unsubstantiated. The chapter presented a number of interaction techniques that are particularly suited (or not) for a given interface type. It also discussed the dilemmas facing designers when using a particular kind of interface, for example, abstract versus realism, menu selection versus free-form text input, and human-like versus non-human-like. Finally, it presented pointers to specific design guidelines and exemplary systems that have been designed using a given interface.

Key Points

- Many interfaces have emerged post the WIMP/GUI era, including voice, wearable, mobile, tangible, brain-computer, smart, robots, and drones.
- A range of design and research questions need to be considered when deciding which interface to use and what features to include.
- Natural user interfaces may not be as natural as graphical user interfaces—it depends on the task, user, and context.
- An important concern that underlies the design of any kind of interface is how information is represented to the user (be it speech, multimedia, virtual reality, augmented reality), so that they can make sense of it with respect to their ongoing activity, for example, playing a game, shopping online, or interacting with a pet robot.
- Increasingly, new interfaces that are context-aware or monitor people raise ethical issues concerned with what data is being collected and for what it is being used.

Further Reading

Many practical books have been published on interface design. Some have been revised into second editions. Publishers like New Riders and O'Reilly frequently offer up-to-date books for a specific interface area (for example web or voice). Some are updated on a regular basis while others are published when a new area emerges. There are also a number of excellent online resources, sets of guidelines, and thoughtful blogs and articles.

DASGUPTA, R. (2019) *Voice User Interface Design: Moving from GUI to Mixed Modal Interaction*. Apress. This is a guide that covers the challenges of moving from GUI design to mixed-modal interactions. It describes how our interactions with devices are rapidly changing, illustrating this through a number of case studies and design principles of VUI design.

ROWLAND, C., GOODMAN, E., CHARLIER, M., LIGHT, A. and LUI, A. (2015) *Designing Connected Products*. O'Reilly. This collection of chapters covers the challenges of designing connected products that go beyond the traditional scope of interaction design and software development. It provides a road map and covers a range of aspects, including pairing devices, new business models, and flow of data in products.

GOOGLE *Material Design* <https://material.io/design/> This online resource provides a living online document that visually illustrates essential interface design principles. It is beautifully laid out and very informative to click through all of the interactive examples that it provides. It shows how to add some physical properties to the digital world to make it feel more intuitive to use across platforms.

KRISHNA, G. (2015) *The Best Interfaces Are No Interfaces*. New Riders. This polemical and funny book challenges the reader to think beyond the screen when designing new interfaces.

KRUG, S. (2014) *Don't Make Me Think!* (3rd edn). New Riders Press. The third edition of this very accessible classic guide on web design presents up-to-date principles and examples on web design with a focus on mobile usability. It is highly entertaining with lots of great illustrations.

NORMAN, D. (2010) Natural interfaces are not natural, *interactions*, May/June, 6–10. This is a thought-provoking essay by Don Norman about what is natural may not appear to be natural, which is still very relevant today.



INTERVIEW with Leah Buechley

Leah Buechley is an independent designer, engineer, and educator. She has a PhD in computer science and a degree in physics. She began her studies as a dance major and has also been deeply engaged in theater, art, and design over the years. She was the founder and director of the high-low tech group at the MIT media lab from 2009 to 2014. She has always blended the sciences and the arts in her education and her career, as witnessed by her current work, which consists of computer science, industrial design, interaction design, art, and electrical engineering.

What is the focus of your work?

I'm most interested in changing the culture of technology and engineering to make it more diverse and inclusive. To achieve that goal, I blend computation and electronics with a range of different materials and employ techniques drawn from art, craft, and design. This approach leads to technologies and learning experiences that appeal to a diverse group of people.

Can you give me some examples of how you mesh the digital with physical materials?

My creative focus for the last several years has been computational design—a process in which objects are designed via an algorithm and then constructed with a combination of fabrication and hand building. I'm especially excited about computational ceramics and have been developing a set of

tools and techniques that enable people to integrate programming and hand building with clay.

I've also been working on a project called LilyPad Arduino (or LilyPad) for over 10 years. LilyPad is a construction kit that enables people to embed computers and electronics into fabric. It's a set of sewable electronic pieces, including microcontrollers, sensors, and LEDs, that are stitched together with conductive thread. People can use the kit to make singing pillows, glow-in-the-dark handbags, and interactive ball gowns.

Another example is the work my former students and I have done in paper-based computing. My former student Jie Qi developed a kit called Chibitronics circuit stickers that lets you build interactive paper-based projects. Based on her years of research in my group at MIT, the kit is a set of flexible peel-and-stick electronic stickers. You can connect ultra-thin LEDs, microcontrollers, and sensors with conductive ink, tape, or thread to quickly make beautiful electronic sketches.

The LilyPad and Chibitronics kits are now used by people around the world to learn computing and electronics. It's been fascinating and exciting to see this research have a tangible impact.

Why would anyone want to wear a computer in their clothing?

Computers open up new creative possibilities for designers. Computers are simply

(Continued)

a new tool, albeit an especially powerful one, in a designer's toolbox. They allow clothing designers to make garments that are dynamic and interactive. Clothing that can, for example, change color in response to pollution levels, sparkle when a loved one calls you on the phone, or notify you when you blood pressure increases.

How do you involve people in your research?

I engage with people in a few different ways. First, I design hardware and software tools to help people build new and different kinds of technology. The LilyPad is a good example of this kind of work. I hone these designs by teaching workshops to different groups of people. And once a tool is stable, I work hard to disseminate it to users in the real world. The LilyPad has been commercially available since 2007, and it has been fascinating and exciting to see how a group of real-world designers—who are predominantly female—is using it to build things like smart sportswear, plush video game controllers, soft robots, and interactive embroideries.

I also strive to be as open as possible with my own design and engineering explorations. I document and publish as much information as I can about the materials, tools, and processes I use. I apply an open source approach not only to the software and hardware I create but, as much as I can, to the entire creative process. I develop and share tutorials, classroom and workshop curricula, materials references, and engineering techniques.

What excites you most about your work?

I am infatuated with materials. There is nothing more inspiring than a sheet of

heavy paper, a length of wool felt, a slab of clay, or a box of old motors. My thinking about design and technology is largely driven by explorations of materials and their affordances. So, materials are always delightful. For example, the shape and surface pattern of the cup in Figure 7.33 were computationally designed. A template of the design was then laser cut and pressed into a flat sheet or “slab” of clay. Finally, the clay was folded into shape and then fired and glazed using traditional ceramic techniques. But the real-world adoption of tools I've designed and the prospect this presents for changing technology culture is perhaps what's most exciting. My most dearly held goal is to expand and diversify technology culture, and it's tremendously rewarding to see evidence that my work is doing that.



Figure 7.33 An example of a computational cup

Source: Used courtesy of Leah Buechley

Chapter 8

DATA GATHERING

8.1 Introduction

8.2 Five Key Issues

8.3 Data Recording

8.4 Interviews

8.5 Questionnaires

8.6 Observation

8.7 Choosing and Combining Techniques

Objectives

The main goals of the chapter are to accomplish the following:

- Discuss how to plan and run a successful data gathering program.
- Enable you to plan and run an interview.
- Empower you to design a simple questionnaire.
- Enable you to plan and carry out an observation.

8.1 Introduction

Data is everywhere. Indeed, it is common to hear people say that we are drowning in data because there is so much of it. So, what is data? Data can be numbers, words, measurements, descriptions, comments, photos, sketches, films, videos, or almost anything that is useful for understanding a particular design, user needs, and user behavior. Data can be quantitative or qualitative. For example, the time it takes a user to find information on a web page and the number of clicks to get to the information are forms of quantitative data. What the user says about the web page is a form of qualitative data. But what does it mean to collect these and other kinds of data? What techniques can be used, and how useful and reliable is the data that is collected?

This chapter presents some techniques for data gathering that are commonly used in interaction design activities. In particular, data gathering is a central part of discovering requirements and evaluation. Within the requirements activity, data gathering is conducted

to collect sufficient, accurate, and relevant data so that design can proceed. Within evaluation, data gathering captures user reactions and their performance with a system or prototype. All of the techniques that we will discuss can be done with little to no programming or technical skills. Recently, techniques for scraping large volumes of data from online activities, such as Twitter posts, have become available. These and other techniques for managing huge amounts of data, and the implications of their use, are discussed in Chapter 10, “Data at Scale.”

Three main techniques for gathering data are introduced in this chapter: interviews, questionnaires, and observation. The next chapter discusses how to analyze and interpret the data collected. Interviews involve an interviewer asking one or more interviewees a set of questions, which may be highly structured or unstructured; interviews are usually synchronous and are often face-to-face, but they don’t have to be. Increasingly, interviews are conducted remotely using one of the many teleconferencing systems, such as Skype or Zoom, or on the phone. Questionnaires are a series of questions designed to be answered asynchronously, that is, without the presence of the investigator. These questionnaires may be paper-based or available online. Observation may be direct or indirect. Direct observation involves spending time with individuals observing their activities as they happen. Indirect observation involves making a record of the user’s activity as it happens, to be studied at a later date. All three techniques may be used to collect qualitative or quantitative data.

Although this is a small set of basic techniques, they are flexible and can be combined and extended in many ways. Indeed, it is important not to focus on just one data gathering technique, if possible, but to use them in combination so as to avoid biases that are inherent in any one approach.

8.2 Five Key Issues

Five key issues require attention for any data gathering session to be successful: goal setting, identifying participants, the relationship between the data collector and the data provider, triangulation, and pilot studies.

8.2.1 Setting Goals

The main reason for gathering data is to glean information about users, their behavior, or their reaction to technology. Examples include understanding how technology fits into family life, identifying which of two icons representing “send message” is easier to use, and finding out whether the planned redesign for a handheld meter reader is headed in the right direction. There are many different reasons for gathering data, and before beginning, it is important to set specific goals for the study. These goals will influence the nature of data gathering sessions, the data gathering techniques to be used, and the analysis to be performed (Robson and McCartan, 2016).

The goals may be expressed more or less formally, for instance, using some structured or even mathematical format or using a simple description such as the ones in the previous paragraph. Whatever the format, however, they should be clear and concise. In interaction design, it is more common to express goals for data gathering informally.

8.2.2 Identifying Participants

The goals developed for the data gathering session will indicate the types of people from whom data is to be gathered. Those people who fit this profile are called the *population* or *study population*. In some cases, the people from whom to gather data may be clearly identifiable—maybe because there is a small group of users and access to each one is easy. However, it is more likely that the participants to be included in data gathering need to be chosen, and this is called *sampling*. The situation where all members of the target population are accessible is called *saturation sampling*, but this is quite rare. Assuming that only a portion of the population will be involved in data gathering, then there are two options: probability sampling or nonprobability sampling. In the former case, the most commonly used approaches are simple random sampling or stratified sampling; in the latter case, the most common approaches are convenience sampling or volunteer panels.

Random sampling can be achieved by using a random number generator or by choosing every n th person in a list. Stratified sampling relies on being able to divide the population into groups (for example, classes in a secondary school) and then applying random sampling. Both convenience sampling and volunteer panels rely less on choosing the participants and more on the participants being prepared to take part. The term *convenience sampling* is used to describe a situation where the sample includes those who were available rather than those specifically selected. Another form of convenience sampling is *snowball sampling*, in which a current participant finds another participant and that participant finds another, and so on. Much like a snowball adds more snow as it gets bigger, the population is gathered up as the study progresses.

The crucial difference between probability and nonprobability methods is that in the former you can apply statistical tests and generalize to the whole population, while in the latter such generalizations are not robust. Using statistics also requires a sufficient number of participants. Vera Toepoel (2016) provides a more detailed treatment of sampling, particularly in relation to survey data.

BOX 8.1

How Many Participants Are Needed?

A common question is, how many participants are needed for a study? In general, having more participants is better because interpretations of statistical test results can be stated with higher confidence. What this means is that any differences found among conditions are more likely to be caused by a genuine effect rather than being due to chance.

More formally, there are many ways to determine how many participants are needed. Four of these are saturation, cost and feasibility analysis, guidelines, and prospective power analysis (Caine, 2016).

- Saturation relies on data being collected until no new relevant information emerges, and so it is not possible to know the number in advance of the saturation point being reached.
- Choosing the number of participants based on cost and feasibility constraints is a practical approach and is justifiable; this kind of pragmatic decision is common in industrial projects but rarely reported in academic research.

(Continued)

- Guidelines may come from experts or from “local standards,” for instance, from an accepted norm in the field.
- Prospective power analysis is a rigorous method used in statistics that relies on existing quantitative data about the topic; in interaction design, this data is often unavailable, making this approach infeasible, such as when a new technology is being developed.

Kelly Caine (2016) investigated the sample size (number of participants) for papers published at the international Computer-Human Interaction (CHI) conference in 2014. She found that several factors affected the sample size, including the method being used and whether the data was collected in person or remotely. In this set of papers, the sample size varied from 1 to 916,000, with the most common size being 12. So, a “local standard” for interaction design would therefore suggest 12 as a rule of thumb. ■

8.2.3 Relationship with Participants

One significant aspect of any data gathering is the relationship between the person (people) doing the gathering and the person (people) providing the data. Making sure that this relationship is clear and professional will help to clarify the nature of the study. How this is achieved varies in different countries and different settings. In the United States and United Kingdom, for example, it is achieved by asking participants to sign an informed consent form, while in Scandinavia such a form is not required. The details of this form will vary, but it usually asks the participants to confirm that the purpose of the data gathering and how the data will be used has been explained to them and that they are willing to continue. It usually explains that their data will be private and kept in a secure place. It also often includes a statement that participants may withdraw at any time and that in this case none of their data will be used in the study.

The informed consent form is intended to protect the interests of both the data gatherer and the data provider. The gatherer wants to know that the data they collect can be used in their analysis, presented to interested parties, and published in reports. The data provider wants reassurance that the information they give will not be used for other purposes or in any context that would be detrimental to them. For example, they want to be sure that personal contact information and other personal details are not made public. This is especially true when people with disabilities or children are being interviewed. In the case of children, using an informed consent form reassures parents that their children will not be asked threatening, inappropriate, or embarrassing questions, or be asked to look at disturbing or violent images. In these cases, parents are asked to sign the form. Figure 8.1 shows an example of a typical informed consent form.

This kind of consent is also not generally required when gathering requirements data for a commercial company where a contract usually exists between the data collector and the data provider. An example is where a consultant is hired to gather data from company staff during the course of discovering requirements for a new interactive system to support timesheet entry. The employees of this company would be the users of the system, and the consultant would therefore expect to have access to the employees to gather data about the timesheet activity. In addition, the company would expect its employees to cooperate in this exercise. In this case, there is already a contract in place that covers the data

| Crowdsourcing Design for Citizen Science Organizations | | | | | | | |
|--|--|------------------------|--|-----------|--|------|--|
| SHORT VERSION OF CONSENT FORM for participants at the University of Maryland – 18 YEARS AND OLDER | | | | | | | |
| <p>You are invited to participate in a research project being conducted by the researchers listed on the bottom of the page. In order for us to be allowed to use any data you wish to provide, we must have your consent.</p> <p>In the simplest terms, we hope you will use the mobile phone, tabletop, and project website at the University of Maryland to</p> <ul style="list-style-type: none"> • Take pictures • Share observations about the sights you see on campus • Share ideas that you have to improve the design of the phone or tabletop application or website • Comment on pictures, observations, and design ideas of others <p>The researchers and others using CampusNet will be able to look at your comments and pictures on the tabletop and/or website, and we may ask if you are willing to answer a few more questions (either on paper, by phone, or face-to-face) about your whole experience. You may stop participating at any time.</p> <p>A long version of this consent form is available for your review and signature, or you may opt to sign this shorter one by <i>checking off all the boxes that reflect your wishes and signing and dating the form below.</i></p> <p><input type="checkbox"/> I agree that any photos I take using the CampusNet application may be uploaded to the tabletop at the University of Maryland and/or a website now under development.</p> <p><input type="checkbox"/> I agree to allow any comments, observations, and profile information that I choose to share with others via the online application to be visible to others who use the application at the same time or after me.</p> <p><input type="checkbox"/> I agree to be videotaped/audiotaped during my participation in this study.</p> <p><input type="checkbox"/> I agree to complete a short questionnaire during or after my participation in this study.</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 10px;"> <tbody> <tr> <td style="width: 30%; padding: 5px;">NAME [Please print]</td> <td style="width: 70%;"></td> </tr> <tr> <td style="padding: 5px;">SIGNATURE</td> <td></td> </tr> <tr> <td style="padding: 5px;">DATE</td> <td></td> </tr> </tbody> </table> <p>[Contact information of Senior Researcher responsible for the project]</p> | | NAME [Please print] | | SIGNATURE | | DATE | |
| NAME [Please print] | | | | | | | |
| SIGNATURE | | | | | | | |
| DATE | | | | | | | |

Figure 8.1 Example of an informed consent form

gathering activity, and therefore an informed consent form is less likely to be required. As with most ethical issues, the important thing is to consider the situation and make a judgment based on the specific circumstances. Increasingly, projects and organizations that collect personal data from people need to demonstrate that it is protected from unauthorized

access. For example, the European Union's General Data Protection Regulation (GDPR) came into force in May 2018. It applies to all EU organizations and offers the individual unprecedented control over their personal data.

For more information about GDPR and data protection law in Europe and the United Kingdom, see:
<https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/>

Incentives to take part in data gathering sessions may also be needed. For example, if there is no clear advantage to the respondents, incentives may persuade them to take part; in other circumstances, respondents may see it as part of their job or as a course requirement to take part. For example, if support sales executives are asked to complete a questionnaire about a new mobile sales application, then they are likely to agree if the new device will impact their day-to-day lives. In this case, the motivation for providing the required information is clear. However, when collecting data to understand how appealing a new interactive app is for school children, different incentives would be appropriate. Here, the advantage for individuals to take part is not so obvious.

8.2.4 Triangulation

Triangulation is a term used to refer to the investigation of a phenomenon from (at least) two different perspectives (Denzin, 2006; Jupp, 2006). Four types of triangulation have been defined (Jupp, 2006).

- Triangulation of data means that data is drawn from different sources at different times, in different places, or from different people (possibly by using a different sampling technique).
- Investigator triangulation means that different researchers (observers, interviewers, and so on) have been involved in collecting and interpreting the data.
- Triangulation of theories means the use of different theoretical frameworks through which to view the data or findings.
- Methodological triangulation means to employ different data gathering techniques.

The last of these is the most common form of triangulation—to validate the results of some inquiry by pointing to similar results yielded through different perspectives. However, validation through true triangulation is difficult to achieve. Different data gathering methods result in different kinds of data, which may or may not be compatible. Using different theoretical frameworks may or may not result in complementary findings, but to achieve theoretical triangulation would require the theories to have similar philosophical underpinnings. Using more than one data gathering technique, and more than one data analysis approach, is good practice because it leads to insights from the different approaches even though it may not be achieving true triangulation.

Triangulation has sometimes been used to make up for the limitations of another type of data collection (Mackay and Anne-Laure Fayard, 1997). This is a different rationale from the original idea, which has more to do with the verification and reliability of data. Furthermore,

a kind of triangulation is being used increasingly in crowd sourcing and other studies involving large amounts of data to check that the data collected from the original study is real and reliable. This is known as checking for “ground truth.”

For an example of methodological triangulation, see:

<https://medium.com/design-voices/the-power-of-triangulation-in-design-research-64a0957d47d2>

For more information about ground truth and how ground truth databases are used to check data obtained in autonomous driving, see “The HCI Bench Mark Suite: Stereo and Flow Ground Truth with Uncertainties for Urban Autonomous Driving” at <https://ieeexplore.ieee.org/document/7789500/>.

8.2.5 Pilot Studies

A pilot study is a small trial run of the main study. The aim is to make sure that the proposed method is viable before embarking on the real study. For example, the equipment and instructions can be checked, the questions for an interview or in a questionnaire can be tested for clarity, and an experimental procedure can be confirmed as viable. This can identify potential problems in advance so that they can be corrected. Distributing 500 questionnaires and then being told that two of the questions were very confusing wastes time, annoys participants, and is an expensive error that could be avoided by doing a pilot study.

If it is difficult to find participants or access to them is limited, asking colleagues or peers to participate can work as an alternative for a pilot study. Note that anyone involved in a pilot study cannot be involved in the main study itself. Why? Because they will know more about the study and this can distort the results.

BOX 8.2

Data, Information, and Conclusions

There is an important difference between raw data, information, and conclusions. Data is what you collect; this is then analyzed and interpreted and conclusions drawn. Information is gained from analyzing and interpreting the data and conclusions represent the actions to be taken based on the information. For example, consider a study to determine whether a new screen layout for a local leisure center has improved the user’s experience when booking a swimming lesson. In this case, the data collected might include a set of times to complete the booking, user comments regarding the new screen layout, biometric readings of the user’s

(Continued)

heart rate while booking a lesson, and so on. At this stage, the data is raw. Information will emerge once this raw data has been analyzed and the results interpreted. For example, analyzing the data might indicate that users who have been using the leisure center for more than five years find the new layout frustrating and take longer to book, while those who have been using it for less than two years find the new layout helpful and can book lessons more quickly. This indicates that the new layout is good for newcomers but not so good for long-term users of the leisure center; this is information. A conclusion from this might be that a more extensive help system is needed for more experienced users to become used to the changes. ■

8.3 Data Recording

Capturing data is necessary so that the results of a data gathering session can be analyzed and shared. Some forms of data gathering, such as questionnaires, diaries, interaction logging, scraping, and collecting work artifacts, are self-documenting and no further recording is necessary. For other techniques, however, there is a choice in recording approaches. The most common of these are taking notes, photographs, or recording audio or video. Often, several data recording approaches are used together. For example, an interview may be voice recorded, and then to help the interviewer in later analysis, a photograph of the interviewee may be taken to remind the interviewer about the context of the discussion.

Which data recording approaches are used will depend on the goal of the study and how the data will be used, the context, the time and resources available, and the sensitivity of the situation; the choice of data recording approach will affect the level of detail collected and how intrusive the data gathering will be. In most settings, audio recording, photographs, and notes will be sufficient. In others, it is essential to collect video data so as to record in detail the intricacies of the activity and its context. Three common data recording approaches are discussed next.

8.3.1 Notes Plus Photographs

Taking notes (by hand or by typing) is the least technical and most flexible way of recording data, even if it seems old-fashioned. Handwritten notes may be transcribed in whole or in part, and while this may seem tedious, it is usually the first step in analysis, and it gives the analyst a good overview of the quality and contents of the data collected. Tools exist for supporting data collection and analysis, but the advantages of handwritten notes include that using pen and paper can be less intrusive than typing and is more flexible, for example, for drawing diagrams of work layouts. Furthermore, researchers often comment that writing notes helps them to focus on what is important and starts them thinking about what the data is telling them. The disadvantages of notes include that it can be difficult to capture the right highlights, and it can be tiring to write and listen or observe at the same time. It is easy to lose concentration, biases creep in, handwriting can be difficult to decipher, and the speed of writing is limited. Working with a colleague can reduce some of these problems while also providing another perspective.

If appropriate, photograph(s) and short videos (captured via smartphones or other hand-held devices) of artifacts, events, and the environment can supplement notes and hand-drawn sketches, providing that permission has been given to collect data using these approaches.

8.3.2 Audio Plus Photographs

Audio recording is a useful alternative to note-taking and is less intrusive than video. During observation, it allows observers to focus on the activity rather than on trying to capture every spoken word. In an interview, it allows the interviewer to pay more attention to the interviewee rather than trying to take notes as well as listening. It isn't always necessary to transcribe all of the data collected—often only sections are needed, depending on the goals of the study. Many studies do not need a great level of detail, and instead recordings are used as a reminder and as a source of anecdotes for reports. It is surprising how evocative audio recordings of people or places from the data session can be, and those memories provide added context to the analysis. If audio recording is the main or only data collection technique, then the quality needs to be good; performing interviews remotely, for example using Skype, can be compromised because of poor connections and acoustics. Audio recordings are often supplemented with photographs.

8.3.3 Video

Smartphones can be used to collect short video clips of activity. They are easy to use and less obtrusive than setting up sophisticated cameras. But there are occasions when a video is needed for long periods of time or when holding a phone is unreliable, for example, recording how designers collaborate together in a workshop or how teens interact in a “makerspace,” in which people can work on projects while sharing ideas, equipment, and knowledge. For these kinds of sessions, more professional video equipment that clearly captures both visual and audio data is more appropriate. Other ways of recording facial expressions together with verbal comments are also being used, such as GoToMeeting, which can be operated both in-person and remotely. Using such systems can create additional planning issues that have to be addressed to minimize how intrusive the recording is, while at the same time making sure that the data is of good quality (Denzin and Lincoln, 2011). When considering whether to use a camera, Heath et al. (2010) suggest the following issues to consider:

- *Deciding whether to fix the camera's position or use a roving recorder.* This decision depends on the activity being recorded and the purpose to which the video data will be put, for example, for illustrative purposes only or for detailed data analysis. In some cases, such as pervasive games, a roving camera is the only way to capture the required action. For some studies, the video on a smartphone may be adequate and require less effort to set up.
- *Deciding where to point the camera in order to capture what is required.* Heath and his colleagues suggest carrying out fieldwork for a short time before starting to video record in order to become familiar with the environment and be able to identify suitable recording locations. Involving the participants themselves in deciding what and where to record also helps to capture relevant action.
- *Understanding the impact of the recording on participants.* It is often assumed that video recording will have an impact on participants and their behavior. However, it is worth taking an empirical approach to this issue and examining the data itself to see whether there is any evidence of people changing their behavior such as orienting themselves toward the camera.

ACTIVITY 8.1

Imagine that you are a consultant who is employed to help develop a new augmented reality garden planning tool to be used by amateur and professional garden designers. The goal is to find out how garden designers use an early prototype as they walk around their clients' gardens sketching design ideas, taking notes, and asking the clients about what they like and how they and their families use the garden. What are the advantages and disadvantages of the three approaches (note-taking, audio recording with photographs, and video) for data recording in this environment?

Comment

Handwritten notes do not require specialized equipment. They are unobtrusive and flexible but difficult to do while walking around a garden. If it starts to rain, there is no equipment to get wet, but notes may get soggy and difficult to read (and write!). Garden planning is a highly visual, aesthetic activity, so supplementing notes with photographs would be appropriate.

Video captures more information, for example, continuous panoramas of the landscape, what the designers are seeing, sketches, comments, and so on, but it is more intrusive and will also be affected by the weather. Short video sequences recorded on a smartphone may be sufficient as the video is unlikely to be used for detailed analysis. Audio may be a good compromise, but synchronizing audio with activities such as looking at sketches and other artifacts later can be tricky and error prone. ■

8.4 Interviews

Interviews can be thought of as a “conversation with a purpose” (Kahn and Cannell, 1957). How much like an ordinary conversation the interview will be depends on the type of interview. There are four main types of interviews: open-ended or unstructured, structured, semi-structured, and group interviews (Fontana and Frey, 2005). The first three types are named according to how much control the interviewer imposes on the conversation by following a predetermined set of questions. The fourth type, which is often called a *focus group*, involves a small group guided by a facilitator. The facilitation may be quite informal or follow a structured format.

The most appropriate approach to interviewing depends on the purpose of the interview, the questions to be addressed, and the interaction design activity. For example, if the goal is first to gain impressions about users' reactions to a new design concept, then an informal, open-ended interview is often the best approach. But if the goal is to get feedback about a particular design feature, such as the layout of a new web browser, then a structured interview or questionnaire is often better. This is because the goals and questions are more specific in the latter case.

8.4.1 Unstructured Interviews

Open-ended or unstructured interviews are at one end of a spectrum of how much control the interviewer has over the interview process. They are exploratory and are similar to conversations around a particular topic; they often go into considerable depth. Questions posed by

the interviewer are open, meaning that there is no particular expectation about the format or content of answers. For example, the first question asked of all participants might be: “What are the pros and cons of having a wearable?” Here, the interviewee is free to answer as fully or as briefly as they want, and both the interviewer and interviewee can steer the interview. For example, often the interviewer will say: “Can you tell me a bit more about . . .” This is referred to as *probing*.

Despite being unstructured and open, the interviewer needs a plan of the main topics to be covered so that they can make sure that all of the topics are discussed. Going into an interview without an agenda should not be confused with being open to hearing new ideas (see section 8.4.5, “Planning and Conducting an Interview”). One of the skills needed to conduct an unstructured interview is getting the balance right between obtaining answers to relevant questions and being prepared to follow unanticipated lines of inquiry.

A benefit of unstructured interviews is that they generate rich data that is often interrelated and complex, that is, data that provides a deep understanding of the topic. In addition, interviewees may mention issues that the interviewer has not considered. A lot of unstructured data is generated, and the interviews will not be consistent across participants since each interview takes on its own format. Unstructured interviews can be time-consuming to analyze, but they can also produce rich insights. Themes can be identified across interviews using techniques from grounded theory and other analytic approaches, as discussed in Chapter 9, “Data Analysis, Interpretation, and Presentation.”

8.4.2 Structured Interviews

In structured interviews, the interviewer asks predetermined questions similar to those in a questionnaire (see section 8.5, “Questionnaires”), and the same questions are used with each participant so that the study is standardized. The questions need to be short and clearly worded, and they are typically closed questions, which means that they require an answer from a predetermined set of alternatives. (This may include an “other” option, but ideally this would not be chosen often.) Closed questions work well if the range of possible answers is known or if participants don’t have much time. Structured interviews are useful only when the goals are clearly understood and specific questions can be identified. Example questions for a structured interview might be the following:

- “Which of the following websites do you visit most frequently: Amazon.com, Google.com, or msn.com?”
- “How often do you visit this website: every day, once a week, once a month, less often than once a month?”
- “Do you ever purchase anything online: Yes/No? If your answer is Yes, how often do you purchase things online: every day, once a week, once a month, less frequently than once a month?”

Questions in a structured interview are worded the same for each participant and are asked in the same order.

8.4.3 Semi-structured Interviews

Semi-structured interviews combine features of structured and unstructured interviews and use both closed and open questions. The interviewer has a basic script for guidance so that the same topics are covered with each interviewee. The interviewer starts with preplanned

questions and then probes the interviewee to say more until no new relevant information is forthcoming. Here's an example:

Which music websites do you visit most frequently?

Answer: Mentions several but stresses that they prefer hottestmusic.com

Why?

Answer: Says that they like the site layout

Tell me more about the site layout.

Answer: Silence, followed by an answer describing the site's layout

Anything else that you like about the site?

Answer: Describes the animations

Thanks. Are there any other reasons for visiting this site so often that you haven't mentioned?

It is important not to pre-empt an answer by phrasing a question to suggest that a particular answer is expected. For example, "You seemed to like this use of color . . ." assumes that this is the case and will probably encourage the interviewee to answer that this is true so as not to offend the interviewer. Children are particularly prone to behave in this way (see Box 8.3, "Working with different kinds of users.") The body language of the interviewer, for example whether they are smiling, scowling, looking disapproving, and so forth, can have a strong influence on whether the interviewee will agree with a question, and the interviewee needs to have time to speak and not be rushed.

Probes are a useful device for getting more information, especially neutral probes such as "Do you want to tell me anything else?" and prompts that remind interviewees if they forget terms or names help to move the interview along. Semi-structured interviews are intended to be broadly replicable, so probing and prompting aim to move the interview along without introducing bias.

BOX 8.3

Working with Different Kinds of Users

Focusing on the needs of users and including users in the design process is a central theme of this book. But users vary considerably based on their age, educational, life, and cultural experiences, and physical and cognitive abilities. For example, children think and react to situations differently than adults. Therefore, if children are to be included in data gathering sessions, then child-friendly methods are needed to make them feel at ease so that they will communicate with you. For very young children of pre-reading or early reading age, data gathering sessions need to rely on images and chat rather than written instructions or questionnaires. Researchers who work with children have developed sets of "smileys," such as those shown in Figure 8.2, so that children can select the one that most closely represents their feelings (see Read et al., 2002).



Figure 8.2 A smileyometer gauge for early readers

Source: Read et al. (2002)

Similarly, different approaches are needed when working with users from different cultures (Winschiers-Theophilus et al., 2012). In their work with local communities in Namibia, Heike Winschiers-Theophilus and Nicola Bidwell (2013) had to find ways of communicating with local participants, which included developing a variety of visual and other techniques to communicate ideas and collect data about the collective understanding and feelings inherent in the local cultures of the people with whom they worked.

Laurianne Sitbon and Shanjana Farhin (2017) report a study in which researchers interacted with people with intellectual disabilities, where they involved caregivers who knew each participant well and could appropriately make the researchers' questions more concrete. This made it more understandable for the participants. An example of this was when the interviewer assumed that the participant understood the concept of a phone app to provide information about bus times. The caregiver made their questions more concrete for the participant by relating the concept of the phone app to familiar people and circumstances and bringing in a personal example (for instance, "So you don't have to ring your mom to say 'Mom, I am lost'").

Another group of technology users are studied by the field of Animal-Computer Interaction (Mancini et al., 2017). Data gathering with animals poses additional and different challenges. For example, in their study of dogs' attention to TV screens, Ilyena Hirskyj-Douglas et al. (2017) used a combination of observation and tracking equipment to capture when a dog turns their head. But interpreting the data, or checking that the interpretation is accurate, requires animal behavior expertise. ■

The examples in Box 8.3 demonstrate that technology developers need to adapt their data collection techniques to suit the participants with whom they work. As the saying goes, "One size doesn't fit all."

8.4.4 Focus Groups

Interviews are often conducted with one interviewer and one interviewee, but it is also common to interview people in groups. One form of group interview that is sometimes used in interaction design activities is the focus group. Normally, three to ten people are involved, and the discussion is led by a trained facilitator. Participants are selected to provide a representative sample of the target population. For example, in the evaluation of a university website, a group of administrators, faculty, and students may form three separate focus groups because they use the web for different purposes. In requirements activities, a focus group may be held in order to identify conflicts in expectations or terminology from different stakeholders.



The focus group hated it. So he showed it to an out-of-focus group.

Source: Mike Baldwin / Cartoon Stock

The benefit of a focus group is that it allows diverse or sensitive issues to be raised that might otherwise be missed, for example in the requirements activity to understand multiple points within a collaborative process or to hear different user stories (Unger and Chandler, 2012). The method is more appropriate for investigating shared issues rather than individual experiences. Focus groups enable people to put forward their own perspectives. A preset agenda is developed to guide the discussion, but there is sufficient flexibility for the facilitator to follow unanticipated issues as they are raised. The facilitator guides and prompts discussion, encourages quiet people to participate, and stops verbose ones from dominating the discussion. The discussion is usually recorded for later analysis, and participants may be invited to explain their comments more fully at a later date.

Focus groups can be useful, but only if used for the right kind of activities. For a discussion of when focus groups don't work, see the following links:

<https://www.nomensa.com/blog/2016/are-focus-groups-useful-research-technique-ux>

<http://gerrymcgovern.com/why-focus-groups-dont-work/>

The format of focus groups can be adapted to fit within local cultural settings. For example, a study with the Mbeere people of Kenya aimed to find out how water was being used, any plans for future irrigation systems, and the possible role of technology in water management (Warrick et al., 2016). The researcher met with the elders from the community, and the focus group took the form of a traditional Kenyan “talking circle,” in which the elders sit in a circle and each person gives their opinions in turn. The researcher, who was from the Mbeere community, knew that it was impolite to interrupt or suggest that the conversation needed to move along, because traditionally each person speaks for as long as they want.

8.4.5 Planning and Conducting an Interview

Planning an interview involves developing the set of questions or topics to be covered, collating any documentation to give to the interviewee (such as consent form or project description), checking that recording equipment works, structuring the interview, and organizing a suitable time and place.

Developing Interview Questions

Questions may be open-ended (or open) or closed-ended (or closed). Open questions are best suited where the goal of the session is exploratory; closed questions are best suited where the possible answers are known in advance. An unstructured interview will usually consist mainly of open questions, while a structured interview will usually consist of closed questions. A semi-structured interview may use a combination of both types.

DILEMMA

What They Say and What They Do

What users say isn't always what they do. People sometimes give the answers that they think show them in the best light, they may have forgotten what happened, or they may want to please the interviewer by answering in the way they think will satisfy them. This may be problematic when the interviewer and interviewee don't know each other, especially if the interview is being conducted remotely by Skype, Cisco Webex, or another digital conferencing system.

For example, Yvonne Rogers et al. (2010) conducted a study to investigate whether a set of twinkly lights embedded in the floor of an office building could persuade people to take the stairs rather than the lift (or elevator). In interviews, participants told the researchers that they did not change their behavior but logged data showed that their behavior did, in fact, change significantly. So, can interviewers believe all of the responses they get? Are the respondents telling the truth, or are they simply giving the answers that they think the interviewer wants to hear?

It isn't possible to avoid this behavior, but an interviewer can be aware of it and reduce such biases by choosing questions carefully, by getting a large number of participants, or by using a combination of data gathering techniques. ■

The following guidelines help in developing interview questions (Robson and McCartan, 2016):

- Long or compound questions can be difficult to remember or confusing, so split them into two separate questions. For example, instead of “How do you like this smartphone app compared with previous ones that you have used?” say, “How do you like this smartphone app?” “Have you used other smartphone apps?” If so, “How did you like them?” This is easier for the interviewee to respond to and easier for the interviewer to record.
- Interviewees may not understand jargon or complex language and might be too embarrassed to admit it, so explain things to them in straightforward ways.
- Try to keep questions neutral, both when preparing the interview script and in conversation during the interview itself. For example, if you ask “Why do you like this style of interaction?” this question assumes that the person does like it and will discourage some interviewees from stating their real feelings.

ACTIVITY 8.2

Several devices are available for reading ebooks, watching movies, and browsing photographs (see Figure 8.3). The design differs between makes and models, but they are all aimed at providing a comfortable user experience. An increasing number of people also read books and watch movies on their smartphones, and they may purchase phones with larger screens for this purpose.



Figure 8.3 (a) Sony's eReader, (b) Amazon's Kindle, (c) Apple's iPad, and (d) Apple's iPhone
 Source: (a) Sony Europe Limited, (b) Martyn Landi / PA Archive / PA Images, (c) Mark Lennihan / AP Images, and (d) Helen Sharp

The developers of a new device for reading books online want to find out how appealing it will be to young people aged 16–18, so they have decided to conduct some interviews.

1. What is the goal of this data gathering session?
2. Suggest ways of recording the interview data.

3. Suggest a set of questions for use in an unstructured interview that seeks to understand the appeal of reading books online to young people in the 16–18 year old age group.
4. Based on the results of the unstructured interviews, the developers of the new device have found that an important acceptance factor is whether the device can be handled easily. Write a set of semi-structured interview questions to evaluate this aspect based on an initial prototype and run a pilot interview with two of your peers. Ask them to comment on your questions and refine them based on their comments.

Comment

1. The goal is to understand what makes devices for reading books online appealing to people aged 16–18.
2. Audio recording will be less cumbersome and distracting than taking notes, and all important points will be captured. Video recording is not needed in this initial interview as it isn't necessary to capture any detailed interactions. However, it would be useful to take photographs of any devices referred to by the interviewee.
3. Possible questions include the following: Why do you read books online? Do you ever read print-based books? If so, what makes you choose to read online versus a print-based format? Do you find reading a book online comfortable? In what way(s) does reading online versus reading from print affect your ability to become engrossed in the story you are reading?
4. Semi-structured interview questions may be open or closed-ended. Some closed-ended questions that you might ask include the following:
 - Have you used any kind of device for reading books online before?
 - Would you like to read a book online using this device?
 - In your opinion, is the device easy to handle?
 Some open-ended questions, with follow-on probes, include the following:
 - What do you like most about the device? Why?
 - What do you like least about the device? Why?
 - Please give me an example of where the device was uncomfortable or difficult to use. ■

It is helpful when collecting answers to closed-ended questions to list possible responses together with boxes that can be checked. Here's one way to convert some of the questions from Activity 8.2:

1. Have you used a device for reading books online before? (Explore previous knowledge.)
Interviewer checks box: *Yes* *No* *Don't remember/know*
2. Would you like to read a book using a device designed for reading online? (Explore initial reaction; then explore the response.)
Interviewer checks box: *Yes* *No* *Don't know*
3. Why?
If response is "Yes" or "No," interviewer asks, "Which of the following statements represents your feelings best?"

For “Yes,” interviewer checks one of these boxes:

- I don't like carrying heavy books.
- This is fun/cool.
- My friend told me they are great.
- It's the way of the future.
- Another reason (interviewer notes the reason).

For “No,” interviewer checks one of these boxes:

- I don't like using gadgets if I can avoid it.
- I can't read the screen clearly.
- I prefer the feel of paper.
- Another reason (interviewer notes the reason).

4. In your opinion, is the device for reading online easy to handle or cumbersome?

Interviewer checks one of these boxes:

- Easy to handle
- Cumbersome
- Neither

Running the Interview

Before starting, make sure that the goals of the interview have been explained to the interviewee and that they are willing to proceed. Finding out about the interviewee and their environment before the interview will make it easier to put them at ease, especially if it is an unfamiliar setting.

During the interview, it is better to listen more than to talk, to respond with sympathy but without bias, and to appear to enjoy the interview. The following is a common sequence for an interview (Robson and McCartan, 2016):

1. An introduction in which the interviewer introduces themselves and explains why they are doing the interview, reassures interviewees regarding any ethical issues, and asks if they mind being recorded, if appropriate. This should be exactly the same for each interviewee.
2. A warm-up session where easy, nonthreatening questions come first. These may include questions about demographic information, such as “What area of the country do you live in?”
3. A main session in which the questions are presented in a logical sequence, with the more probing ones at the end. In a semi-structured interview, the order of questions may vary between participants, depending on the course of the conversation, how much probing is done, and what seems more natural.
4. A cooling-off period consisting of a few easy questions (to defuse any tension that may have arisen).
5. A closing session in which the interviewer thanks the interviewee and switches off the recorder or puts their notebook away, signaling that the interview has ended.

8.4.6 Other Forms of Interview

Conducting face-to-face interviews and focus groups can be impractical, but the prevalence of Skype, Cisco WebEx, Zoom, and other digital conferencing systems, email, and phone-based interactions (voice or chat), sometimes with screen-sharing software, make remote interviewing a good alternative. These are carried out in a similar fashion to face-to-face sessions, but poor connections and acoustics can cause different challenges, and participants may be tempted to multitask rather than focus on the session at hand. Advantages of remote focus groups and interviews, especially when done through audio-only channels, include the following:

- The participants are in their own environment and are more relaxed.
- Participants don't have to travel.
- Participants don't need to worry about what they wear.
- For interviews involving sensitive issues, interviewees can remain anonymous.

In addition, participants can leave the conversation whenever they want to by just cutting the connection, which adds to their sense of security. From the interviewer's perspective, a wider set of participants can be reached easily, but a potential disadvantage is that the facilitator does not have a good view of the interviewees' body language.

For more information and some interesting thoughts on remote usability testing, see <http://www.uxbooth.com/articles/hidden-benefits-remote-research/>.

Retrospective interviews, that is, interviews that reflect on an activity or a data gathering session in the recent past, may be conducted with participants to check that the interviewer has correctly understood what was happening. This is a common practice in observational studies where it is sometimes referred to as *member checking*.

8.4.7 Enriching the Interview Experience

Face-to-face interviews often take place in a neutral location away from the interviewee's normal environment. This creates an artificial context, and it can be difficult for interviewees to give full answers to the questions posed. To help combat this, interviews can be enriched by using props such as personas prototypes or work artifacts that the interviewee or interviewer brings along, or descriptions of common tasks (examples of these kinds of props are scenarios and prototypes, which are covered in Chapter 11, "Discovering Requirements," and Chapter 12, "Design, Prototyping, and Construction"). These props can be used to provide context for the interviewees and help to ground the data in a real setting. Figure 8.4 illustrates the use of personas in a focus group setting.



Figure 8.4 Enriching a focus group with personas displayed on the wall for all participants to see

As another example, Clara Mancini et al. (2009) used a combination of questionnaire prompts and deferred contextual interviews when investigating mobile privacy. A simple multiple-choice questionnaire was sent electronically to the participants' smartphones, and they answered the questions using these devices. Interviews about the recorded events were conducted later, based on the questionnaire answers given at the time of the event.

8.5 Questionnaires

Questionnaires are a well-established technique for collecting demographic data and users' opinions. They are similar to interviews in that they can have closed or open-ended questions, but once a questionnaire is produced, it can be distributed to a large number of participants without requiring additional data gathering resources. Thus, more data can be collected than would normally be possible in an interview study. Furthermore, participants who are located in remote locations or those who cannot attend an interview at a particular time can be involved more easily. Often a message is sent electronically to potential participants directing them to an online questionnaire.

Effort and skill are needed to ensure that questions are clearly worded and the data collected can be analyzed efficiently. Well-designed questionnaires are good for getting answers to specific questions from a large group of people. Questionnaires can be used on their own

or in conjunction with other methods to clarify or deepen understanding. For example, information obtained through interviews with a small selection of interviewees might be corroborated by sending a questionnaire to a wider group to confirm the conclusions.

Questionnaire questions and structured interview questions are similar, so which technique is used when? Essentially, the difference lies in the motivation of the respondent to answer the questions. If their motivation is high enough to complete a questionnaire without anyone else present, then a questionnaire will be appropriate. On the other hand, if the respondents need some persuasion to answer the questions, a structured interview format would be better. For example, structured interviews are easier and quicker to conduct if people will not stop to complete a questionnaire, such as at a train station or while walking to their next meeting.

It can be harder to develop good questionnaire questions compared with structured interview questions because the interviewer is not available to explain them or to clarify any ambiguities. Because of this, it is important that questions are specific; when possible, ask closed-ended questions and offer a range of answers, including a “no opinion” or “none of these” option. Finally, use negative questions carefully, as they can be confusing and may lead to false information. Some questionnaire designers, however, use a mixture of negative and positive questions deliberately because it helps to check the users’ intentions.

8.5.1 Questionnaire Structure

Many questionnaires start by asking for basic demographic information (gender, age, place of birth) and details of relevant experience (the number of hours a day spent searching on the Internet, the level of expertise within the domain under study, and so on). This background information is useful for putting the questionnaire responses into context. For example, if two responses conflict, these different perspectives may be because of their level of experience—a group of people who are using a social networking site for the first time are likely to express different opinions than another group with five years’ experience of using such sites. However, only contextual information that is relevant to the study goal needs to be collected. For example, it is unlikely that a person’s height will provide relevant context to their responses about Internet use, but it might be relevant for a study concerning wearables.

Specific questions that contribute to the data-gathering goal usually follow these demographic questions. If the questionnaire is long, the questions may be subdivided into related topics to make it easier and more logical to complete.

The following is a checklist of general advice for designing a questionnaire:

- Think about the ordering of questions. The impact of a question can be influenced by question order.
- Consider whether different versions of the questionnaire are needed for different populations.
- Provide clear instructions on how to complete the questionnaire, for example, whether answers can be saved and completed later. Aim for both careful wording and good typography.
- Think about the length of the questionnaire, and avoid questions that don’t address the study goals.
- If the questionnaire has to be long, consider allowing respondents to opt out at different stages. It is usually better to get answers to some sections than no answers at all because of dropout.
- Think about questionnaire layout and pacing; for instance, strike a balance between using white space, or individual web pages, and the need to keep the questionnaire as compact as possible.

8.5.2 Question and Response Format

Different formats of question and response can be chosen. For example, with a closed-ended question, it may be appropriate to indicate only one response, or it may be appropriate to indicate several. Sometimes, it is better to ask users to locate their answer within a range. Selecting the most appropriate question and response format makes it easier for respondents to answer clearly. Some commonly used formats are described next.

Check Boxes and Ranges

The range of answers to demographic questions is predictable. Nationality, for example, has a finite number of alternatives, and asking respondents to choose a response from a predefined list makes sense for collecting this information. A similar approach can be adopted if details of age are needed. But since some people do not like to give their exact age, many questionnaires ask respondents to specify their age as a range. A common design error arises when the ranges overlap. For example, specifying two ranges as 15–20 and 20–25 will cause confusion; that is, which box do people who are 20 years old check? Making the ranges 15–19 and 20–24 avoids this problem.

A frequently asked question about ranges is whether the interval must be equal in all cases. The answer is no—it depends on what you want to know. For example, people who might use a website about life insurance are likely to be employed individuals who are 21–65 years old. The question could, therefore, have just three ranges: under 21, 21–65, and over 65. In contrast, to see how the population’s political views vary across generations might require 10-year cohort groups for people over 21, in which case the following ranges would be appropriate: under 21, 21–30, 31–40, and so forth.

Rating Scales

There are a number of different types of rating scales, each with its own purpose (see Oppenheim, 2000). Two commonly used scales are the Likert and semantic differential scales. Their purpose is to elicit a range of responses to a question that can be compared across respondents. They are good for getting people to make judgments, such as how easy, how usable, and the like.

Likert scales rely on identifying a set of statements representing a range of possible opinions, while semantic differential scales rely on choosing pairs of words that represent the range of possible opinions. Likert scales are more commonly used because identifying suitable statements that respondents will understand consistently is easier than identifying semantic pairs that respondents interpret as intended.

Likert Scales

Likert scales are used for measuring opinions, attitudes, and beliefs, and consequently they are widely used for evaluating user satisfaction with products. For example, users’ opinions about the use of color in a website could be evaluated with a Likert scale using a range of numbers, as in question 1 here, or with words as in question 2:

1. The use of color is excellent (where 1 represents strongly agree and 5 represents strongly disagree):

| | | | | |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| 1 | 2 | 3 | 4 | 5 |
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

2. The use of color is excellent:

| | | | | |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Strongly agree | Agree | OK | Disagree | Strongly disagree |
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

In both cases, respondents would be asked to choose the right box, number, or phrase. Designing a Likert scale involves the following steps:

1. *Gather a pool of short statements about the subject to be investigated.* Examples are “This control panel is clear” and “The procedure for checking credit rating is too complex.” A brainstorming session with peers is a good way to identify key aspects to be investigated.
2. *Decide on the scale.* There are three main issues to be addressed here: How many points does the scale need? Should the scale be discrete or continuous? How can the scale be represented? See Box 8.4 What Scales to Use: Three, Five, Seven, or More? for more on this.
3. *Select items for the final questionnaire, and reword as necessary to make them clear.*

In the first example above, the scale is arranged with 1 as the highest choice on the left and 5 as the lowest choice on the right. The logic for this is that first is the best place to be in a race and fifth would be the worst place. While there is no absolute right or wrong way of ordering the numbers other researchers prefer to arrange the scales the other way around with 1 as the lowest on the left and 5 as the highest on the right. They argue that intuitively the highest number suggests the best choice and the lowest number suggests the worst choice. Another reason for going from lowest to highest is that when the results are reported, it is more intuitive for readers to see high numbers representing the best choices. The important thing is to be consistent.

Semantic Differential Scales

Semantic differential scales explore a range of bipolar attitudes about a particular item, each of which is represented as a pair of adjectives. The participant is asked to choose a point between the two extremes to indicate agreement with the poles, as shown in Figure 8.5. The score for the investigation is found by summing the scores for each bipolar pair. Scores are then computed across groups of participants. Notice that in this example the poles are mixed so that good and bad features are distributed on the right and the left. In this example, there are seven positions on the scale.




Figure 8.5 An example of a semantic differential scale

BOX 8.4

What Scales to Use: Three, Five, Seven, or More?

Issues to address when designing Likert and semantic differential scales include the following: how many points are needed on the scale, how should they be presented, and in what form?

Many questionnaires use seven- or five-point scales, and there are also three-point scales. Some even use nine-point scales. Arguments for the number of points go both ways. Advocates of long scales argue that they help to show discrimination. Rating features on an interface is more difficult for most people than, say, selecting among different flavors of ice cream, and when the task is difficult, there is evidence to show that people “hedge their bets.” Rather than selecting the poles of the scales if there is no right or wrong, respondents tend to select values nearer the center. The counterargument is that people cannot be expected to discern accurately among points on a large scale, so any scale of more than five points is unnecessarily difficult to use.

Another aspect to consider is whether to give the scale an even or odd number of points. An odd number provides a clear central point, while an even number forces participants to decide and prevents them from sitting on the fence.

We suggest the following guidelines:

How many points on the scale?

Use a small number, three, for example, when the possibilities are very limited, as in Yes/No type answers.

Yes Don't know No

Use a medium-sized range, five, for example, when making judgments that involve like/dislike or agree/disagree statements.

Strongly agree Agree OK Disagree Strongly disagree

Use a longer range, seven or nine, for example, when asking respondents to make subtle judgments, such as when asking about a user experience dimension such as “level of appeal” of a character in a video game.

┌──────────┴──────────┬──────────┬──────────┬──────────┬──────────┬──────────┐
very appealing ok repulsive

Discrete or continuous?

Use boxes for discrete choices and scales for finer judgments.

What order?

Decide which way to order the scale, and be consistent. ■

ACTIVITY 8.3

Spot four poorly designed features in the excerpt from a questionnaire in Figure 8.6.

2. State your age in years

3. How many hours a day do you spend searching online? <1 hour
 1–3 hours
 3–5 hours
 >5 hours

4. Which of the following do you do online?

purchase goods

send e-mail

visit chatrooms

use bulletin boards

find information

read the news

5. How useful is the Internet to you?

Figure 8.6 A questionnaire with poorly designed features

Comment

Some of the features that could be improved upon include the following:

- Question 2 requests an exact age. Many people prefer not to give this information and would rather position themselves within a range.
- In question 3, the number of hours spent searching is indicated with overlapping scales, that is, 1–3 and 3–5. How would someone answer if they spend 3 hours a day searching online?
- For question 4, the questionnaire doesn't say how many boxes to check.
- The space left for people to answer open-ended question 5 is too small, which will annoy some people and deter them from giving their opinions.

Many online survey tools prevent users from making some of these design errors. It is important, however, to be aware of such things because paper is still sometimes used. ■

8.5.3 Administering Questionnaires

Two important issues when using questionnaires are reaching a representative sample of participants and ensuring a reasonable response rate. For large surveys, potential respondents need to be selected using a sampling technique. However, interaction designers commonly use a small number of participants, often fewer than 20 users. Completion rates of 100 percent are often achieved with these small samples, but with larger or more remote populations, ensuring that surveys are returned is a well-known problem. A 40 percent return is generally acceptable

for many surveys, but much lower rates are common. Depending on your audience, you might want to consider offering incentives (see section 8.2.3, “Relationship with Participants”).

While questionnaires are often online, paper questionnaires may be more convenient in some situations, for example, if participants do not have Internet access or if it is expensive to use. Occasionally, short questionnaires are sent within the body of an email, but more often the advantages of the data being compiled automatically and either partly or fully analyzed make online questionnaires attractive. Online questionnaires are interactive and can include check boxes, radio buttons, pull-down and pop-up menus, help screens, graphics, or videos (see Figure 8.7). They can also provide immediate data validation; for example, the entry must be a number between 1 and 20, and automatically skip questions that are irrelevant to some respondents, such as questions aimed only at teenagers. Other advantages of online questionnaires include faster response rates and automatic transfer of responses into a database for analysis (Toepoel, 2016).

D. Internationally-agreed development goals outlined in the Millennium Declaration :

Is this activity relevant to achieving the MDGs listed below? (see www.un.org/millenniumgoals/ and the targets for each goal) Yes No
If yes, please tick all goals that apply

1. Eradicate poverty and hunger
2. Achieve Universal Primary Education
3. Promote gender equality & empower women
4. Reduce child mortality
5. Improve maternal health
6. Combat HIV/AIDS, Malaria and other diseases
7. Ensure environmental sustainability
8. Develop a global partnership for development

E. More Information :

Please provide a website for this activity
Website (URL) :

F. Geographical Coverage* :

Please tick a box to indicate the geographical coverage
 Local National Regional International
Please specify coverage :

G. Timescale* :

Please tick a box to indicate the timescale of the activity
 Completed Planned for future Ongoing

Specify dates using the format day/month/year (dd/mm/yyyy) :
From: To:

H. Activity Type* :

Please tick one or more boxes to indicate the type of activity described above
 Project Programme WSIS Thematic Meeting Conference Publication Training initiative
 Guidelines Tool-kit Website Database
Other (please specify) :

Figure 8.7 An excerpt from a web-based questionnaire showing check boxes, radio buttons, and pull-down menus

The main problem with online questionnaires is the difficulty of obtaining a random sample of respondents; online questionnaires usually rely on convenience sampling, and hence their results cannot be generalized. In some countries, online questions, often delivered via smartphones, are frequently used in conjunction with television to elicit viewers’ opinions of programs and political events.

Deploying an online questionnaire involves the following steps (Toepoel, 2016, Chapter 10):

1. *Plan the survey timeline.* If there is a deadline, work backward from the deadline and plan what needs to be done on a weekly basis.
2. *Design the questionnaire offline.* Using plain text is useful as this can then be copied more easily into the online survey tool.
3. *Program the online survey.* How long this will take depends on the complexity of the design, for example, how many navigational paths it contains or if it has a lot of interactive features.
4. *Test the survey, both to make sure that it behaves as envisioned and to check the questions themselves.* This includes getting feedback from content experts, survey experts, and potential respondents. This last group forms the basis of a pilot study.
5. *Recruit respondents.* As mentioned earlier, participants may have different reasons for taking part in the survey, but especially when respondents need to be encouraged, make the invitations intriguing, simple, friendly, respectful, trustworthy, motivating, interesting, informative, and short.

There are many online questionnaire templates available that provide a range of options, including different question types (for example open-ended, multiple choice), rating scales (such as Likert, semantic differential), and answer types (for example, radio buttons, check boxes, drop-down menus).

The following activity asks you to make use of one of these templates. Apart from being able to administer an online questionnaire widely, these templates also enable the questionnaire to be segmented. For example, airline satisfaction questionnaires often have different sections for check-in, baggage handling, airport lounge, inflight movies, inflight food service, and so forth. If you didn't use an airport lounge or check your baggage, you can skip those sections. This avoids respondents getting frustrated by having to go through questions that are not relevant to them. It is also a useful technique for long questionnaires, as it ensures that if a respondent opts out for lack of time or gets tired of answering the questions, the data that has been provided already is available to be analyzed.

ACTIVITY 8.4

Go to questionpro.com, surveymonkey.com, or a similar survey site and design your own questionnaire using the set of widgets that is available for a free trial period.

Create an online questionnaire for the set of questions that you developed for Activity 8.2. For each question, produce two different designs; for example, use radio buttons and drop-down menus for one question, and provide a 10-point semantic differential scale and a 5-point scale for another question.

What differences (if any) do you think the two designs will have on a respondent's behavior? Ask a number of people to answer one or the other of your questions and see whether the answers differ for the two designs.

(Continued)

Comment

Respondents may have used the response types in different ways. For example, they may select the end options more often from a drop-down menu than from a list of options that are chosen via radio buttons. Alternatively, you may find no difference and that people's opinions are not affected by the widget style used. Some differences, of course, may be due to the variation between individual responses rather than being caused by features in the questionnaire design. To tease the effects apart, you would need to ask a large number of participants (for instance, in the range 50–100) to respond to the questions for each design. ■

BOX 8.5**Do people answer online questionnaires differently than paper and pencil? If so, why?**

There has been much research examining how people respond to surveys when using a computer compared with paper and pencil methods. Some studies suggest that people are more revealing and consistent in their responses when using a computer to report their habits and behaviors, such as eating, drinking, and amount of exercise see Luce et al. (2003). Students have also been found to rate their instructors less favorably when online (Chang, 2004).

In a Danish study in which 3,600 people were invited to participate, the researchers concluded that although response rates for web-based invitations were lower, they were more cost-effective (by a factor of 10) and had only slightly lower numbers of missing values than questionnaires sent via paper (Ebert et al., 2018). Similarly, a study by Diaz de Rada and Dominguez-Alvarez (2014), in which the quality of the information collected from a survey given to citizens of Andalusia in Spain was analyzed, several advantages of using online versus paper-based questionnaires were identified. These included a low number of unanswered questions, more detailed answers to open-ended questions, and longer answers to questions in the online questionnaires than in the paper questionnaires. In the five open-ended questions, respondents wrote 63 characters more on average on the online questionnaires than on the paper questionnaires. For the questions in which participants had to select from a drop-down menu, there was a better response rate than when the selection was presented on paper with blank spaces.

One factor that can influence how people answer questions is the way the information is structured, such as the use of headers, the ordering, and the placement of questions. Online questionnaires provide more options for presenting information, including the use of drop-down menus, radio buttons, and jump-to options, which may influence how people read and navigate a questionnaire. But do these issues affect respondents' answers? Smyth et al. (2005) have found that providing forced choice formats results in more options being selected. Another example is provided by Funcke et al. (2011), who found that continuous sliders enabled researchers to collect more accurate data because they support continuous rather than discrete scales. They also encouraged higher response rates. What can be concluded from these investigations is that the details of questionnaire design can impact how respondents react. ■

8.6 Observation

Observation is useful at any stage during product development. Early in design, observation helps designers understand the users' context, tasks, and goals. Observation conducted later in development, for example, in evaluation, may be used to investigate how well a prototype supports these tasks and goals.

Users may be observed directly by the investigator as they perform their activities or indirectly through records of the activity that are studied afterward (Bernard, 2017). Observation may also take place in the field or in a controlled environment. In the former case, individuals are observed as they go about their day-to-day tasks in the natural setting. In the latter case, individuals are observed performing specified tasks within a controlled environment such as a usability laboratory.

ACTIVITY 8.5

To appreciate the different merits of observation in the field and observation in a controlled environment, read the following scenarios and answer the questions that appear after.

Scenario 1 A usability consultant joins a group of tourists who have been given a wearable navigation device that fits onto a wrist strap to test on a visit to Stockholm. After sight-seeing for the day, they use the device to find a list of restaurants within 2 kilometers of their current position. Several are listed, and they find the phone numbers of a few, call them to ask about their menus, select one, make a booking, and head off to the restaurant. The usability consultant observes some difficulty operating the device, especially on the move. Discussion with the group supports the evaluator's impression that there are problems with the interface, but on balance the device is useful, and the group is pleased to get a table at a good restaurant nearby.

Scenario 2 A usability consultant observes how participants perform a preplanned task using the wearable navigation device in a usability laboratory. The task requires the participants to find the phone number of a restaurant called Matisse. It takes them several minutes to do this, and they appear to have problems. The video recording and interaction log suggest that the interface is quirky and the audio interaction is of poor quality. This is supported by participants' answers on a user satisfaction questionnaire.

1. What are the advantages and disadvantages of these two types of observation?
2. When might each type of observation be useful?

Comment

1. The advantages of the field study are that the observer saw how the device could be used in a real situation to solve a real problem. They experienced the delight expressed with the overall concept and the frustration with the interface. By watching how the group used the device on the move, they gained an understanding of what the participants liked and what was lacking. The disadvantage is that the observer was an insider in the group,

(Continued)

so how objective could they be? The data is qualitative, and while anecdotes can be very persuasive, how useful are they? Maybe they were having such a good time that their judgment was clouded and they missed hearing negative comments and didn't notice some of the participant's annoyance. Another study could be done to find out more, but it is not possible to replicate the exact conditions of this study. The advantages of the lab study are that it is easier to replicate, so several users could perform the same task, specific usability problems can be identified, users' performance can be compared, and averages for such measures as the time it took to do a specific task and the number of errors can be calculated. The observer could also be more objective as an outsider. The disadvantage is that the study is artificial and says nothing about how the device would be used in the real environment.

2. Both types of study have merits. Which is better depends on the goals of the study. The lab study is useful for examining details of the interaction style to make sure that usability problems with the interface and button design are diagnosed and corrected. The field study reveals how the navigation device is used in a real-world context and how it integrates with or changes users' behavior. Without this study, it is possible that developers might not have discovered the enthusiasm for the device because the reward for doing laboratory tasks is not as compelling as a good meal! In fact, according to Kjeldskov and Skov (2014), there is no definitive answer to which kind of study is preferable for mobile devices. They suggest that the real question is when and how to engage with longitudinal field studies. ■

8.6.1 Direct Observation in the Field

It can be difficult for people to explain what they do or to describe accurately how they achieve a task. It is unlikely that an interaction designer will get a full and true story using interviews or questionnaires. Observation in the field can help fill in details about how users behave and use technology, and nuances that are not elicited from other forms of investigation may be observed. Understanding the context provides important information about why activities happen the way that they do. However, observation in the field can be complicated and harder to do well than at first appreciated. Observation can also result in a lot of data, some of which may be tedious to analyze and not very relevant.

All data gathering should have a clearly stated goal, but it is particularly important to have a focus for an observation session because there is always so much going on. On the other hand, it is also important to be prepared to change the plan if circumstances change. For example, the plan may be to spend one day observing an individual performing a task, but an unexpected meeting crops up, which is relevant to the observation goal and so it makes sense to attend the meeting instead. In observation, there is a careful balance between being guided by goals and being open to modifying, shaping, or refocusing the study as more is learned about the situation. Being able to keep this balance is a skill that develops with experience.

Structuring Frameworks for Observation in the Field

During an observation, events can be complex and rapidly changing. There is a lot for observers to think about, so many experts have a framework to structure and focus their observation. The framework can be quite simple. For example, this is a practitioner's framework for use in evaluation studies that focuses on just three easy-to-remember items:

The person: Who is using the technology at any particular time?

The place: Where are they using it?

The thing: What are they doing with it?

Even a simple framework such as this one based on who, where, and what can be surprisingly effective to help observers keep their goals and questions in sight. Experienced observers may prefer a more detailed framework, such as the following (Robson and McCarten, 2016, p. 328), which encourages them to pay greater attention to the context of the activity:

Space: What is the physical space like, and how is it laid out?

Actors: What are the names and relevant details of the people involved?

Activities: What are the actors doing, and why?

Objects: What physical objects are present, such as furniture?

Acts: What are specific individual actions?

Events: Is what you observe part of a special event?

Time: What is the sequence of events?

Goals: What are the actors trying to accomplish?

Feelings: What is the mood of the group and of individuals?

This framework was devised for any type of observation, so when used in the context of interaction design, it might need to be modified slightly. For example, if the focus is going to be on how some technology is used, the framework could be modified to ask the following:

Objects: What physical objects, in addition to the technology being studied, are present, and do they impact on the technology use?

Both of these frameworks are relatively general and could be used in many different types of study, or as a basis for developing a new framework for a specific study.

ACTIVITY 8.6

1. Find a small group of people who are using any kind of technology, for example, smartphones, household appliances, or video game systems, and try to answer the question, "What are these people doing?" Watch for three to five minutes, and write down what you observe. When finished, note down how it felt to be doing this and any reactions in the group of people observed.
2. If you were to observe the group again, what would you do differently?
3. Observe this group again for about 10 minutes using the detailed framework given above.

(Continued)

Comment

1. What problems did this exercise highlight? Was it hard to watch everything and remember what happened? How did the people being watched feel? Did they know they were being watched? Perhaps some of them objected and walked away. If you didn't tell them that they were being watched, should you have?
2. The initial goal of the observation, that is, to find out what the people are doing, was vague, and chances are that it was quite a frustrating experience not knowing what was significant and what could be ignored. The questions used to guide observation need to be more focused. For example, you might ask the following: What are the people doing with the technology? Is everyone in the group using it? Are they looking pleased, frustrated, serious, happy? Does the technology appear to be central to the users' goals?
3. Ideally, you will have felt more confident this second time, partly because it is the second time doing some observation and partly because the framework provided a structure for what to look at. ■

Degree of Participation

Depending on the type of study, the degree of participation within the study environment varies across a spectrum, which can be characterized as insider at one end and outsider at the other. Where a particular study falls along this spectrum depends on its goal and on the practical and ethical issues that constrain and shape it.

An observer who adopts an approach right at the outsider end of the spectrum is called a *passive observer*, and they will not take any part in the study environment at all. It is difficult to be a truly passive observer in the field, simply because it's not possible to avoid interacting with the activities. Passive observation is more appropriate in lab studies.

An observer who adopts an approach at the insider end of this spectrum is called a *participant observer*. This means that they attempt, at various levels depending on the type of study, to become a member of the group being studied. This can be a difficult role to play since being an observer also requires a certain level of detachment, while being a participant assumes a different role. As a participant observer, it is important to keep the two roles clear and separate so that observation notes are objective while participation is also maintained. It may not be possible to take a full participant observer approach for other reasons. For example, the observer may not be skilled enough in the task at hand, the organization/group may not be prepared for an outsider to take part in their activities, or the timescale may not provide sufficient opportunity to become familiar enough with the task to participate fully. Similarly, if observing activity in a private place such as the home, full participation would be difficult even if, as suggested by some researchers (for example, Bell et al., 2005), you have spent time getting to know the family before starting the study. Chandrika Cycil et al. (2013) overcame this issue in their study of in-car conversations between parents and children by traveling with the families initially for a week and then asking family members to video relevant episodes of activity. In this way, they had gained an understanding of the context and family dynamics and then collected more detailed data to study activity in depth.

Planning and Conducting an Observation in the Field

The frameworks introduced in the previous section are useful for providing focus and also for organizing the observation and data gathering activity. Choosing a framework is important, but there are other decisions that need to be made, including the level of participation to adopt, how to make a record of the data, how to gain acceptance in the group being studied, how to handle sensitive issues such as cultural differences or access to private spaces, and how to ensure that the study uses different perspectives (people, activities, job roles, and so forth).

One way to achieve this last point is to work as a team. This can have several benefits.

- Each person can agree to focus on different people or different parts of the context, thereby covering more ground.
- Observation and reflection can be interweaved more easily when there is more than one observer.
- More reliable data is likely to be generated because observations can be compared.
- Results will reflect different perspectives.

Once in the throes of an observation, there are other issues that need to be considered. For example, it will be easier to relate to some people more than others. Although it will be tempting to pay attention to them more than others, attention needs to be paid to everyone in the group. Observation is a fluid activity, and the study will need to be refocused as it progresses in response to what is learned. Having observed for a while, interesting phenomena that seem relevant will start to emerge. Gradually, ideas will sharpen into questions that guide further observation.

Observing is also an intense and tiring activity, but checking notes and records and reviewing observations and experiences at the end of each day is important. If this is not done, then valuable information may be lost as the next day's events override the previous day's findings. Writing a diary or private blog is one way of achieving this. Any documents or other artifacts that are collected or copied (such as minutes of a meeting or discussion items) can be annotated, describing how they are used during the observed activity. Where an observation lasts several days or weeks, time can be taken out of each day to go through notes and other records.

As notes are reviewed, separate personal opinion from observation and mark issues for further investigation. It is also a good idea to check observations and interpretations with an informant or members of the participant group for accuracy.

DILEMMA

When to Stop Observing?

Knowing when to stop doing any type of data gathering can be difficult for novices, but it is particularly tricky in observational studies because there is no obvious ending. Schedules often dictate when your study ends. Otherwise, stop when nothing new is emerging. Two indications of having done enough are when similar patterns of behavior are being seen and when all of the main stakeholder groups have been observed and a good understanding of their perspectives has been achieved. ■

Ethnography

Ethnography has traditionally been used in the social sciences to uncover the organization of societies and their activities. Since the early 1990s, it has gained credibility in interaction design, and particularly in the design of collaborative systems; see Box 8.6, “Ethnography in Requirements” and Crabtree (2003). A large part of most ethnographic studies is direct observation, but interviews, questionnaires, and studying artifacts used in the activities also feature in many ethnographic studies. A distinguishing feature of ethnographic studies compared with other data gathering is that a situation is observed without imposing any *a priori* structure or framework upon it, and everything is viewed as “strange.” In this way, the aim is to capture and articulate the participants’ perspective of the situation under study.

BOX 8.6

Ethnography in Requirements

The MERboard is a tool scientists and engineers use to display, capture, annotate, and share information in support of the operation of two Mars Exploration Rovers (MERs) on the surface of Mars. The MER (see Figure 8.8) acts like a human geological explorer by collecting and analyzing samples and then transmitting the results to the scientists on Earth. The scientists and engineers collaboratively analyze the data received, decide what to study next, create plans of action, and send commands to the robots on the surface of Mars.



Figure 8.8 Mars Exploration Rover

Source: NASA Jet Propulsion Laboratory (NASA-JPL)

The requirements for MERboard were identified partly through ethnographic fieldwork, observations, and analysis (Trimble et al., 2002). The team of scientists and engineers ran a series of field tests that simulated the process of receiving data, analyzing it, creating plans, and transmitting them to the MERs. The main problems they identified stemmed from the scientists’ limitations in displaying, sharing, and storing information (see Figure 8.9a).



Figure 8.9 (a) The situation before MERboard; (b) a scientist using MERboard to present information

Source: Trimble et al. (2002)

These observations led to the development of MERboard (see Figure 8.9b), which contains four core applications: a whiteboard for brainstorming and sketching, a browser for displaying information from the web, the capability to display personal information and information across several screens, and a file storage space linked specifically to MERboard. ■

Ethnography has become popular within interaction design because it allows designers to obtain a detailed and nuanced understanding of people’s behavior and the use of technology that cannot be obtained by other methods of data gathering (Lazar et al., 2017). While there has been much discussion of how big data can address many design issues, big data is likely to be most powerful when combined with ethnography to explain how and why people do what they do (Churchill, 2018).

The observer in an ethnographic study adopts a participant observer (insider) role as much as possible (Fetterman, 2010). While participant observation is a hallmark of ethnographic studies, it is also used within other methodological frameworks such as action research (Hayes, 2011), where one of the goals is to improve the current situation.

Ethnographic data is based on what is available, what is “ordinary,” what it is that people do, say, and how they work. The data collected therefore has many forms: documents, notes taken by the observer(s), pictures, and room layout sketches. Notes may include snippets of conversations and descriptions of rooms, meetings, what someone did, or how people reacted to a situation. Data gathering is opportunistic, and observers make the most of opportunities as they present themselves. Often, interesting phenomena do not reveal themselves immediately but only later, so it is important to gather as much as possible within the framework of observation. Initially, spend time getting to know people in the participant group and bonding with them. Participants need to understand why the observers are there, what they hope

to achieve, and how long they plan to be there. Going to lunch with them, buying coffee, and bringing small gifts, for example, cookies, can greatly help this socialization process. Moreover, key information may be revealed during one of these informal gatherings.

It is important to show interest in the stories, gripes, and explanations that are provided and to be prepared to step back if a participant's phone rings or someone else enters the workspace. A good tactic is to explain to one of the participants during a quiet moment what you think is happening and then let them correct any misunderstandings. However, asking too many questions, taking pictures of everything, showing off your knowledge, and getting in their way can be very off-putting. Putting up cameras on tripods on the first day may not be a good idea. Listening and watching while sitting on the sidelines and occasionally asking questions is a better approach.

The following is an illustrative list of materials that might be recorded and collected during an ethnographic study (adapted from Crabtree, 2003, p. 53):

- Activity or job descriptions
- Rules and procedures (and so on) that govern particular activities
- Descriptions of activities observed
- Recordings of the talk taking place between parties involved in observed activities
- Informal interviews with participants explaining the detail of observed activities
- Diagrams of the physical layout, including the position of artifacts
- Photographs of artifacts (documents, diagrams, forms, computers, and so on) used in the course of observed activities
- Videos of artifacts as used in the course of observed activities
- Descriptions of artifacts used in the course of observed activities
- Workflow diagrams showing the sequential order of tasks involved in observed activities
- Process maps showing connections between activities

Traditionally, ethnographic studies in this field aim to understand what people do and how they organize action and interaction within a particular context of interest to designers. However, recently there has been a trend toward studies that draw more on ethnography's anthropological roots and the study of culture. This trend has been brought about by the perceived need to use different approaches because the computers and other digital technologies, especially mobile devices, are embedded in everyday activity, and not just in the workplace as in the 1990s.

BOX 8.7

Doing Ethnography Online

As collaboration and social activity online have increased, ethnographers have adapted their approach to study social media and the various forms of computer-mediated communication (Rotman et al., 2013; Bauwens and Genoud, 2014). This practice has various names, the most common of which are online ethnography (Rotman et al., 2012), virtual ethnography (Hine, 2000), and netnography (Kozinets, 2010). Where a community or activity has both an online and offline presence, it is common to incorporate both online and offline techniques within the data gathering program. However, where the community or activities of interest exist almost exclusively online, then mostly online techniques are used and virtual ethnography becomes central.

Why is it necessary to distinguish between online and face-to-face ethnography? It is important because interaction online is different from interaction in person. For example, communication in person is richer (through gesture, facial expression, tone of voice, and so on) than online communication, and anonymity is more easily achieved when communicating online. In addition, virtual worlds have a persistence, due to regular archiving, that does not typically occur in face-to-face situations. This makes characteristics of the communication different, which often includes how ethnographers introduce themselves to the community, how they act within the community, and how they report their findings. For these reasons, some researchers who work primarily online also try to meet with some of the participants face-to-face, particularly when working on sensitive topics (Lingel, 2012).

Special tools may be developed to support ethnographic data collection. Mobilab is an online collaborative platform that was developed for citizens living in Switzerland to report and discuss their daily mobility during an eight-week period using their mobile phones, tablets, and computers (Bauwens and Genoud, 2014). Mobilab enabled the researchers to more easily engage in discussion with participants on a variety of topics, including trucks parking on a bikeway.

For observational studies in large social spaces, such as digital libraries or Facebook, there are different ethical issues to consider. For example, it is unrealistic to ask everyone using a digital library to sign any kind of form agreeing to be involved in the study, yet participants do need to understand the observer's role and the purpose of their study. The presentation of results needs to be modified too. Quotes from participants in the community, even if anonymized in the report, can easily be attributed by a simple search of the community archive or the IP address of the sender, so care is needed to protect their privacy. ■

8.6.2 Direct Observation in Controlled Environments

Observing users in a controlled environment may occur within a purposely built usability lab, but portable labs that can be set up in any room are quite common. Portable laboratories can mean that more participants take part because they don't have to travel away from their normal environment. Observation in a controlled environment inevitably takes on a more formal character than observation in the field, and the user may feel more apprehensive. As with interviews, it is a good idea to prepare a script to guide how the participants will be greeted, be told about the goals of the study and how long it will last, and have their rights explained. Use of a script ensures that each participant will be treated in the same way, which brings more credibility to the results obtained from the study.

The same basic data recording techniques are used for direct observation in the laboratory and field studies (that is, capturing photographs, taking notes, collecting video, and so on), but the way in which these techniques are used is different. In the lab the emphasis is on the details of what individuals do, while in the field the context is important, and the focus is on how people interact with each other, the technology, and their environment.

The arrangement of equipment with respect to the participant is important in a controlled study because details of the person's activity need to be captured. For example, one camera might record facial expressions, another might focus on mouse and keyboard activity, and another might record a broad view of the participant and capture body language. The stream of data from the cameras can be fed into a video editing and analysis suite where it is coordinated and time-stamped, annotated, and partially edited.

The Think-Aloud Technique

One of the problems with observation is that the observer doesn't know what users are thinking and can only guess from what they see. Observation in the field should not be intrusive, as this will disturb the context the study is trying to capture. This limits the questions being asked of the participant. However, in a controlled environment, the observer can afford to be a little more intrusive. The think-aloud technique is a useful way of understanding what is going on in a person's head.

Imagine observing someone who has been asked to evaluate the interface of the web search engine Lycos.com. The user, who does not have much experience of web searches, is told to look for a phone for a 10-year-old child. They are told to type www.lycos.com and then proceed however they think best. They type the URL and get a screen similar to the one in Figure 8.10.



Figure 8.10 Home page of Lycos search engine

Source: <https://www.lycos.com>

Next, they type **child's phone** in the search box. They get a screen similar to the one shown in Figure 8.11. They are silent. What is going on? What are they thinking? One way around the problem of knowing what they are doing is to collect a think-aloud protocol, a technique developed by Anders Ericsson and Herbert Simon (1985) for examining people's problem-solving strategies. The technique requires people to say out loud everything that they are thinking and trying to do so that their thought processes are externalized.

So, let's imagine an action replay of the situation just described, as follows, but this time the user has been instructed to think aloud:

"I'm typing in www.lycos.com, as you told me." <types>

*"Now I am typing **child's phone** and then clicking the search button.*

<pause and silence>

"It's taking a few seconds to respond."

"Oh! Now I have a choice of other websites to go to. Hmm, I wonder which one I should select. Well, it's for a young child so I want a 'child-safe phone.' This one mentions safe phones <He clicks on 7 Best Cell Phones for Kids - Mashable>

"Gosh, there's a lot more models to select from than I expected! Hmm, some of these are for older children. I wonder what I do next to find one for a 10-year-old."




Figure 8.11 The screen that appears in response to searching for “child’s phone”

Source: <https://www.lycos.com>

<pauses and looks at the screen>

I guess I should scroll through them and identify those that might be appropriate.”

<silence . . . >

Now you know more about what the user is trying to achieve, but they are silent again. They are looking at the screen, but what are they thinking now? What are they looking at?

The occurrence of these silences is one of the biggest problems with the think-aloud technique.

ACTIVITY 8.7

Try a think-aloud exercise yourself. Go to a website, such as Amazon or eBay, and look for something to buy. Think aloud as you search and notice how you feel and behave.

Afterward, reflect on the experience. Was it difficult to keep speaking all the way through the task? Did you feel awkward? Did you stop talking when you got stuck?

Comment

Feeling self-conscious and awkward doing this is a common response, and some people say they feel really embarrassed. Many people forget to speak out loud and find it difficult to do so when the task becomes difficult. In fact, you probably stopped speaking when the task became demanding, and that is exactly the time when an observer is most eager to hear what’s happening.

(Continued)

If a user is silent during a think-aloud protocol, the observer could interrupt and remind them to think out loud, but that would be intrusive. Another solution is to have two people work together so that they talk to each other. Working with another person (called *constructive interaction* [Miyake, 1986]) is often more natural and revealing because participants talk in order to help each other along. This technique has proved to be particularly successful with children, and it also avoids possible cultural influences on concurrent verbalization (Clemmensen et al., 2008). ■

8.6.3 Indirect Observation: Tracking Users' Activities

Sometimes direct observation is not possible because it is too intrusive or observers cannot be present over the duration of the study, and so activities are tracked indirectly. Diaries and interaction logs are two techniques for doing this.

Diaries

Participants are asked to write a diary of their activities on a regular basis, including things like what they did, when they did it, what they found hard or easy, and what their reactions were to the situation. For example, Sohn et al. (2008) asked 20 participants to record their mobile information needs through text messages and then to use these messages as prompts to help them answer six questions on a website at the end of each day. From the data collected, they identified 16 categories of mobile information needs, the most frequent of which was “trivia.”

Diaries are useful: when participants are scattered and unreachable in person; when the activity is private, for example, in the home; or when it relates to feelings, for instance, emotions or motivation. For example, Jang et al. (2016) used diaries with interviews to collect data about users' experiences with smart TVs in the home as compared to within a controlled lab setting. The study in the home was conducted over several weeks during which participants were asked to keep a diary of their experiences and feelings. Surveys were also collected. This mixed-methods study informed the user experience design of future systems.

Diaries have several advantages: they do not take up much researcher time to collect data; they do not require special equipment or expertise; and they are suitable for long-term studies. In addition, templates, like those used in open-ended online questionnaires, can be created online to standardize the data entry format so that the data can be entered directly into a database for analysis. However, diary studies rely on participants being reliable and remembering to complete them at the assigned time and as instructed, so incentives may be needed, and the process has to be straightforward.

Determining how long to run a diary study can be tricky. If the study goes on for too long, participants may lose interest and need incentives to continue. In contrast, if the study is too short, important data may be missed. For example, in a study of children's experiences of a game, Elisa Mekler et al. (2014) used diaries to collect data after each gaming session in a series. After the first few sessions, all of the children in the study showed loss of motivation for the game. However, by the end of the study, those who completed the game were more motivated than those who did not complete the game. Had the data been collected only once, the researchers may not have observed the impact of game completion on the children's motivation.

Another problem is that the participants' memories of events may be exaggerated or detail is forgotten; for example, they may remember them as better or worse than they really were or as taking more or less time than they actually did take. One way of mitigating this problem is to collect other data in diaries (such as photographs including selfies, audio and video clips, and so on). Scott Carter and Jennifer Mankoff (2005) considered whether capturing events through pictures, audio, or artifacts related to the event affects the results of the diary study. They found that images resulted in more specific recall than other media, but audio was useful for capturing events when taking a photo was too awkward. Tangible artifacts, such as those shown in Figure 8.12, also encouraged discussion about wider beliefs and attitudes.



Figure 8.12 Some tangible objects collected by participants involved in a study about a jazz festival
Source: Carter and Mankoff (2005). Reproduced with permission of ACM Publications

The experience sampling method (ESM) is similar to a diary in that it relies on participants recording information about their everyday activities. However, it differs from more traditional diary studies because participants are prompted at random times via email, text message, or similar means to answer specific questions about their context, feelings, and actions (Hektner et al., 2006). These prompts have the benefit of encouraging immediate data capture. Niels van Berkel et al. (2017) provide a comprehensive survey of ESM and its evolution, tools, and uses across a wide range of studies.

Interaction Logs, Web Analytics, and Data Scraping

Interaction logging uses software to record users' activity in a log that can be examined later. A variety of actions may be recorded, such as key presses and mouse or other device movements, time spent searching a web page, time spent looking at help systems, and task flow

through software modules. A key advantage of logging activity is that it is unobtrusive provided system performance is not affected, but it also raises ethical concerns about observing participants if this is done without their knowledge. Another advantage is that large volumes of data can be logged automatically. Visualization tools are therefore helpful for exploring and analyzing this data quantitatively and qualitatively. Algorithmic and statistical methods may also be used.

Examining the trail of activity that people leave behind when they are active on websites, Twitter, or Facebook is also a form of indirect observation. You can see an example of this by looking at a Twitter feed to which you have access, for example, that of a friend, president, prime minister, or some other leader. These trails allow examination of discussion threads on a particular topic, such as climate change, or reactions to comments made by a public figure or to a topic that is trending today. If there are just a few posts, then it is easy to see what is going on, but often the most interesting posts are those that generate a lot of comments. Examining thousands, tens of thousands, and even millions of posts requires automated techniques. Web analytics and data scraping are discussed further in Chapter 10.

8.7 Choosing and Combining Techniques

Combining data gathering techniques into a single data gathering program is common practice, for example, when collecting case study data (see Box 8.8). The benefit of using a combination of methods is to provide multiple perspectives. Choosing which data gathering techniques to use depends on a variety of factors related to the study goals. There is no right technique or combination of techniques, but some will undoubtedly be more appropriate than others. The decision about which to use will need to be made after taking all of the factors into account.

Table 8.1 provides an overview to help choose a set of techniques for a specific project. It lists the kind of information obtained (such as answers to specific questions) and the type of data (for example, mostly qualitative or mostly quantitative). It also includes some advantages and disadvantages for each technique. Note that different modalities can be used for some of these techniques. For example, interviews and focus groups can be conducted face-to-face, by phone, or through teleconferencing, so when considering advantages and disadvantages of the techniques, this should also be taken into account.

In addition, technique choice is influenced by practical issues.

- *The focus of the study.* What kind of data will support the focus and goal of the study? This will be influenced by the interaction design activity and the level of maturity of the design.
- *The participants involved.* Characteristics of the target user group including their location and availability.
- *The nature of the technique.* Does the technique require specialist equipment or training, and do the investigators have the appropriate knowledge and experience?
- *Available resources.* Expertise, tool support, time, and money.

| Technique | Good for | Kind of data | Advantages | Disadvantages |
|--|--|--|--|--|
| Interviews | Exploring issues | Some quantitative but mostly qualitative | Interviewer can guide interviewee if necessary. Encourages contact between developers and users. | Artificial environment may intimidate interviewee. It also removes them from the environment where work is typically being done. |
| Focus groups | Collecting multiple viewpoints | Some quantitative but mostly qualitative | Highlights areas of consensus and conflict. Encourages contact between developers and users. | Possibility of dominant characters. |
| Questionnaires | Answering specific questions | Quantitative and qualitative | Can reach many people with low resource requirements. | The design is key. Response rates may be low. Unless carefully designed, the responses may not provide suitable data. |
| Direct observation in the field | Understanding context of user activity | Mostly qualitative | Observing gives insights that other techniques don't provide. | Very time-consuming. Huge amounts of data are produced. |
| Direct observation in a controlled environment | Capturing the detail of what individuals do | Quantitative and qualitative | Can focus on the details of a task without interruption. | Results may have limited use in the normal environment because the conditions were artificial. |
| Indirect observation | Observing users without disturbing their activity; data captured automatically | Quantitative (logging) and qualitative (diary) | User doesn't get distracted by the data gathering; automatic recording means that it can extend over long periods of time. | A large amount of quantitative data needs tool support to analyze (logging); participants' memories may exaggerate (diary). |

Table 8.1 Overview of data gathering techniques and their use

BOX 8.8

Collecting case study data

Case studies often use a combination of methods, for example, direct and indirect observations and interviews. Although people frequently use the term *case study* colloquially to refer to a study that they are using as a case example, there is also a case study methodology that collects field study data over days, months, or even years. There is a body of literature that provides advice on how to do good case studies. Robert Yin (2013), for example, identifies these data collection sources: documentation, archival records, interviews, direct observations, participant observation, and physical artifacts. Case studies are good for integrating multiple perspectives, for example, studying new technology in the wild, and for giving meaning to first impressions. The data collection process tends to be intensive, concurrent, interactive, and iterative.

In a study of how local communities organize and adapt technology for managing their local rivers and streams, approaching it as a case study allowed a detailed contextual analysis of events and relationships that occurred over multiple groups of volunteers during a two-year period (Preece et al., 2019). From this study, the researchers learned about the volunteers' needs for highly flexible software to support the diverse groups of participants working on a wide range of water-related topics. ■

ACTIVITY 8.9

For each of the following products, consider what kinds of data gathering would be appropriate and how to use the different techniques introduced earlier. Assume that product development is just starting and that there is sufficient time and resources to use any of the techniques.

1. A new software app to support a small organic produce shop. There is a system running already with which the users are reasonably happy, but it is looking dated and needs upgrading.
2. An innovative device for diabetes sufferers to help them record and monitor their blood sugar levels.
3. An ecommerce website that sells fashion clothing for young people.

Comment

1. As this is a small shop, there are likely to be few stakeholders. Some period of observation would be important to understand the context of the new and the old systems. Interviewing the staff rather than giving them questionnaires is likely to be appropriate because there aren't very many of them, and this will yield richer data and give the developers a chance to meet the users. Organic produce is regulated by a variety of laws, so looking at this documentation will help you understand any legal constraints that have to be taken into account. This suggests a series of interviews with the main users to understand the positive and negative features of the existing system, a short observation session to understand the context of the system, and a study of documentation surrounding the regulations.
2. In this case, the user group is quite large and spread out geographically, so talking to all of them is not feasible. However, interviewing a representative sample of potential users,

possibly at a local diabetic clinic, is feasible. Observing current practices to monitor blood sugar levels will help you understand what is required. An additional group of stakeholders would be those who use or have used the other products on the market. These stakeholders can be questioned about their experience with their existing devices so that the new device can be an improvement. A questionnaire sent to a wider group in order to confirm the findings from the interviews would be appropriate, as might a focus group where possible.

3. Again, the user group is quite large and spread out geographically. In fact, the user group may not be very well defined. Interviews backed up by questionnaires and focus groups would be appropriate. In this case, identifying similar or competing sites and evaluating them will help provide information for an improved product. ■

In-Depth Activity

The aim of this in-depth activity is to practice data gathering. Assume that you have been employed to improve the user experience of an interactive product such as a smartphone app, a digital media player, a Blu-ray player, computer software, or some other type of technology. This existing product may be redesigned, or a completely new product may be created. To do the assignment, find a group of people or a single individual prepared to be the user group. These could be your family, friends, peers, or people in a local community group.

For this assignment:

- (a) Clarify the basic goal of improving the product by considering what this means in your circumstances.
- (b) Watch the group (or person) casually to get an understanding of any issues that might create challenges for this activity and any information to help refine the study goals.
- (c) Explain how you would use each of the three data gathering techniques: interview, questionnaire, and observation in your data gathering program. Explain how your plan takes account of triangulation.
- (d) Consider your relationship with the user group and decide if an informed consent form is required. (Figure 8.1 will help you to design one if needed.)
- (e) Plan your data gathering program in detail.
 - Decide what kind of interview to run and design a set of interview questions. Decide how to record the data, then acquire and test any equipment needed and run a pilot study.
 - Decide whether to include a questionnaire in your data gathering program, and design appropriate questions for it. Run a pilot study to check the questionnaire.
 - Decide whether to use direct or indirect observation and where on the outsider/insider spectrum should the observers be. Decide how to record the data, then acquire and test any equipment needed and run a pilot study.
- (f) Carry out the study, but limit its scope. For example, interview only two or three people or plan only two half-hour observation periods.
- (g) Reflect on this experience and suggest what you would do differently next time.
Keep the data gathered, as this will form the basis of the in-depth activity in Chapter 9.

Summary

This chapter has focused on three main data gathering methods that are commonly used in interaction design: interviews, questionnaires, and observation. It has described in detail the planning and execution of each. In addition, five key issues of data gathering were presented, and how to record the data gathered was discussed.

Key Points

- All data gathering sessions should have clear goals.
- Depending on the study context, an informed consent form and other permissions may be needed to run the study.
- Running a pilot study helps to test out the feasibility of a planned data gathering session and associated instruments such as questions.
- Triangulation involves investigating a phenomenon from different perspectives.
- Data may be recorded using handwritten notes, audio or video recording, a camera, or any combination of these.
- There are three styles of interviews: structured, semi-structured, and unstructured.
- Questionnaires may be paper-based, via email, or online.
- Questions for an interview or questionnaire can be open or closed-ended. Closed-ended questions require the interviewee to select from a limited range of options. Open-ended questions accept a free-range response.
- Observation may be direct or indirect.
- In direct observation, the observer may adopt different levels of participation, ranging from insider (participant observer) to outsider (passive observer).
- Choosing appropriate data gathering techniques depends on the focus of the study, participants involved, nature of the technique, and resources available.

Further Reading

FETTERMAN, D. M. (2010). *Ethnography: Step by Step* (3rd ed.) Applied Social Research Methods Series, Vol. 17. Sage. This book introduces the theory and practice of ethnography, and it is an excellent guide for beginners. It covers both data gathering and data analysis in the ethnographic tradition.

FULTON SURI, J. (2005) *Thoughtless Acts?* Chronicle Books, San Francisco. This intriguing little book invites you to consider how people react to their environment. It is a good introduction to the art of observation.

HEATH, C., HINDMARSH, J. AND LUFF, P. (2010) *Video in Qualitative Research: Analyzing Social Interaction in Everyday Life*. Sage. This is an accessible book that provides practical advice and guidance about how to set up and perform data gathering using video recording. It also covers data analysis, presenting findings, and potential implications from video research based on their own experience.

OLSON, J. S. AND KELLOGG, W. A. (eds) (2014) *Ways of Knowing in HCI*. Springer. This edited collection contains useful chapters on a wide variety of data collection and analysis techniques. Some topics that are particularly relevant to this chapter include: ethnography, experimental design, log data collection and analysis, ethics in research, and more.

ROBSON, C. AND McCARTAN, K. (2016) *Real World Research* (4th edn). John Wiley & Sons. This book provides comprehensive coverage of data gathering and analysis techniques and how to use them. Early books and related books by Robson also address topics discussed in this chapter.

TOEPOEL, V. (2016) *Doing Surveys Online*. Sage. This book is a “hands-on guide” for preparing and conducting a wide range of surveys including surveys for mobile devices, opt-in surveys, panels, polls, and more. It also discusses details about sampling that can be applied to other data gathering techniques.

Chapter 9

DATA ANALYSIS, INTERPRETATION, AND PRESENTATION

- 9.1 Introduction
- 9.2 Qualitative and Quantitative
- 9.3 Basic Quantitative Analysis
- 9.4 Basic Qualitative Analysis
- 9.5 What Kind of Analytic Framework to Use
- 9.6 Tools to Support Data Analysis
- 9.7 Interpreting and Presenting the Findings

Objectives

The main goals of this chapter are to accomplish the following:

- Discuss the difference between qualitative and quantitative data and analysis.
- Enable you to analyze data gathered from questionnaires.
- Enable you to analyze data gathered from interviews.
- Enable you to analyze data gathered from observation studies.
- Make you aware of software packages that are available to help your analysis.
- Identify some of the common pitfalls in data analysis, interpretation, and presentation.
- Enable you to interpret and present your findings in a meaningful and appropriate manner.

9.1 Introduction

The kind of analysis that can be performed on a set of data will be influenced by the goals identified at the outset and the data gathered. Broadly speaking, a qualitative analysis approach, a quantitative analysis approach, or a combination of qualitative and quantitative approaches may be taken. The last of these is very common, as it provides a more comprehensive account of the behavior being observed or the performance being measured.

Most analysis, whether it is quantitative or qualitative, begins with the initial reactions or observations from the data. This may involve identifying patterns or calculating simple numerical values such as ratios, averages, or percentages. For all data, but especially when dealing with large volumes of data (that is, Big Data), it is useful to look over the data to check for any anomalies that might be erroneous. For example, people who are 999 years old. This process is known as *data cleansing*, and there are often digital tools to help with the process. This initial analysis is followed by more detailed work using structured frameworks or theories to support the investigation.

Interpretation of the findings often proceeds in parallel with analysis, but there are different ways to interpret results, and it is important to make sure that the data supports any conclusions. A common mistake is for the investigator's existing beliefs or biases to influence the interpretation of results. Imagine that an initial analysis of the data has revealed a pattern of responses to customer care questionnaires that indicates that inquiries from customers routed through the Sydney office of an organization take longer to process than those routed through the Moscow office. This result can be interpreted in many different ways. For example, the customer care operatives in Sydney are less efficient, they provide more detailed responses, the technology supporting the inquiry process in Sydney needs to be updated, customers reaching the Sydney office demand a higher level of service, and so on. Which one is correct? To determine whether any of these potential interpretations is accurate, it would be appropriate to look at other data such as customer inquiry details and maybe to interview staff.

Another common mistake is to make claims that go beyond what the data can support. This is a matter of interpretation and of presentation. Using words such as *many* or *often* or *all* when reporting conclusions needs to be carefully considered. An investigator needs to remain as impartial and objective as possible if the conclusions are to be trusted. Showing that the conclusions are supported by the results is an important skill to develop.

Finally, finding the best way to present findings is equally skilled, and it depends on the goals but also on the audience for whom the study was performed. For example, a formal notation may be used to report the results for the requirements activity, while a summary of problems found, supported by video clips of users experiencing those problems, may be better for presentation to the team of developers.

This chapter introduces a variety of methods, and it describes in more detail how to approach data analysis and presentation using some of the common approaches taken in interaction design.

9.2 Quantitative and Qualitative

Quantitative data is in the form of numbers, or data that can easily be translated into numbers. Examples are the number of years' experience the interviewees have, the number of projects a department handles at a time, or the number of minutes it takes to perform a task. *Qualitative data* is in the form of words and images, and it includes descriptions, quotes from interviewees, vignettes of activity, and photos. It is possible to express qualitative data in numerical form, but it is not always meaningful to do so (see Box 9.1).

It is sometimes assumed that certain forms of data gathering can only result in quantitative data and that others can only result in qualitative data. However, this is a fallacy. All forms of data gathering discussed in the previous chapter may result in qualitative and quantitative

data. For example, on a questionnaire, questions about the participant's age or number of software apps they use in a day will result in quantitative data, while any comments will result in qualitative data. In an observation, quantitative data that may be recorded includes the number of people involved in a project or how many hours someone spends sorting out a problem, while notes about feelings of frustration, or the nature of interactions between team members, are qualitative data.

Quantitative analysis uses numerical methods to ascertain the magnitude, amount, or size of something; for example, the attributes, behavior, or strength of opinion of the participants. For example, in describing a population, a quantitative analysis might conclude that the average person is 5 feet 11 inches tall, weighs 180 pounds, and is 45 years old. *Qualitative analysis* focuses on the nature of something and can be represented by themes, patterns, and stories. For example, in describing the same population, a qualitative analysis might conclude that the average person is tall, thin, and middle-aged.

BOX 9.1

Use and Abuse of Numbers

Numbers are infinitely malleable and can make a convincing argument, but it is important to justify the manipulation of quantitative data and what the implications will be. Before adding a set of numbers together, finding an average, calculating a percentage, or performing any other kind of numerical translation, consider whether the operation is meaningful in the specific context.

Qualitative data can also be turned into a set of numbers. Translating non-numerical data into a numerical or ordered scale is appropriate at times, and this is a common approach in interaction design. However, this kind of translation also needs to be justified to ensure that it is meaningful in the given context. For example, assume you have collected a set of interviews from sales representatives about their use of a new mobile app for reporting sales queries. One way of turning this data into a numerical form would be to count the number of words uttered by each interviewee. Conclusions might then be drawn about how strongly the sales representatives feel about the app; for example, the more they had to say about the product, the stronger they felt about it. But do you think this is a wise way to analyze the data? Does it help to answer the study questions?

Other, less obvious, abuses include translating small population sizes into percentages. For example, saying that 50 percent of users take longer than 30 minutes to place an order through an e-commerce website carries a different meaning than saying that two out of four users had the same problem. It is better not to use percentages unless the number of data points is at least 10, and even then it is appropriate to use both percentages and raw numbers to make sure that the claim is not misunderstood.

It is possible to perform legitimate statistical calculations on a set of data and still present misleading results by not making the context clear or by choosing the particular calculation that gives the most favorable result (Huff, 1991). In addition, choosing and applying the best statistical test requires careful thinking (Cairns, 2019), as using an inappropriate test can unintentionally misrepresent the data. ■

9.2.1 First Steps in Analyzing Data

Having collected the data, some initial processing is normally required before data analysis can begin in earnest. For example, audio data may be transcribed by hand or by using an automated tool, such as Dragon; quantitative data, such as time taken or errors made, is usually entered into a spreadsheet, like Excel. Initial analysis steps for data typically collected through interviews, questionnaires, and observation are summarized in Table 9.1.

| | Usual raw data | Example qualitative data | Example quantitative data | Initial processing steps |
|----------------|---|---|---|---|
| Interviews | Audio recordings. Interviewer notes. Video recordings. | Responses to open-ended questions. Video pictures. Respondent's opinions. | Age, job role, years of experience. Responses to close-ended questions. | Transcription of recordings. Expansion of notes. Entry of answers to close-ended questions into a spreadsheet |
| Questionnaires | Written responses. Online database. | Responses to open-ended questions. Responses in "further comments" fields. Respondent's opinions. | Age, job role, years of experience. Responses to close-ended questions. | Clean up data. Filter into different data sets. Synchronization between data recordings. |
| Observation | Observer's notes. Photographs. Audio and video recordings. Data logs. Think-aloud Diaries. | Records of behavior. Description of a task as it is undertaken. Copies of informal procedures. | Demographics of participants. Time spent on a task. The number of people involved in an activity. How many different types of activity are undertaken. | Expansion of notes. Transcription of recordings. |

Table 9.1 Data gathered and typical initial processing steps for interviews, questionnaires, and observation

Interviews

Interviewer notes need to be written up and expanded as soon as possible after the interview has taken place so that the interviewer's memory is clear and fresh. An audio or video recording may be used to help in this process, or it may be transcribed for more detailed analysis.

Transcription takes significant effort, as people talk more quickly than most people can type (or write), and the recording is not always clear. It is worth considering whether to transcribe the whole interview or just sections of it that are relevant. Deciding what is relevant, however, can be difficult. Revisiting the goals of the study to see which passages address the research questions can guide this process.

Close-ended questions are usually treated as quantitative data and analyzed using basic quantitative analysis (see Section 9.3 “Basic Quantitative Analysis”). For example, a question that asks for the respondent’s age range can easily be analyzed to find out the percentage of respondents in each. More complicated statistical techniques are needed to identify relationships between responses that can be generalized, such as whether there is an interaction between the condition being tested and a demographic. For example, do people of different ages use Facebook for different lengths of time when first logging on in the morning or at night before they go to bed? Open-ended questions typically result in qualitative data that might be searched for categories or patterns of response.

Questionnaires

Increasingly, questionnaire responses are provided using online surveys, and the data is automatically stored in a database. The data can be filtered according to respondent subpopulations (for instance, everyone under 16) or according to a particular question (for example, to understand respondents’ reactions to one kind of robot personality rather than another). This allows analyses to be conducted on subsets of the data and hence to draw specific conclusions for more targeted goals. To conduct this kind of analysis requires sufficient data from a large enough sample of participants.

Observation

Observation can result in a wide variety of data including notes, photographs, data logs, think-aloud recordings (often called *protocols*), video, and audio recordings. Taken together, these different types of data can provide a rich picture of the observed activity. The difficult part is working out how to combine the different sources to create a coherent narrative of what has been recorded; analytic frameworks, discussed in section 9.5, can help with this. Initial data processing includes writing up and expanding notes and transcribing elements of the audio and video recordings and the think-aloud protocols. For observation in a controlled environment, initial processing might also include synchronizing different data recordings.

Transcriptions and the observer’s notes are most likely to be analyzed using qualitative approaches, while photographs provide contextual information. Data logs and some elements of the observer’s notes would probably be analyzed quantitatively.

9.3 Basic Quantitative Analysis

Explaining statistical analysis requires a whole book on its own (for example, see Cairns, 2019). Here, we introduce two basic quantitative analysis techniques that can be used effectively in interaction design: averages and percentages. Percentages are useful for standardizing the data, particularly to compare two or more large sets of responses.

Averages and percentages are fairly well-known numerical measures. However, there are three different types of average, and using the wrong one can lead to the misinterpretation of the results. These three are: mean, median, and mode. *Mean* refers to the commonly understood interpretation of average; that is, add together all the figures and divide by the number of figures with which you started. Median and mode averages are less well-known but are very useful. The *median* is the middle value of the data when the numbers are ranked. The *mode* is the most commonly occurring number. For example, in a set of data (2, 3, 4, 6, 6, 7, 7, 7, 8), the median is 6 and the mode is 7, while the mean is $50/9 = 5.56$. In this case, the difference between the different averages is not that great. However, consider the set (2, 2, 2, 2, 450). Now the median is 2, the mode is 2, and the mean is $458/5 = 91.6$!



"Looks good. Let me run it past the number-crunchers."

Source: Mike Baldwin / Cartoon Stock

Use of simple averages can provide useful overview information, but they need to be used with caution. Evangelos Karapanos et al. (2009) go further and suggest that averaging treats diversity among participants as error and proposes the use of a multidimensional scaling approach instead.

Before any analysis can take place, the data needs to be collated into analyzable data sets. Quantitative data can usually be translated into rows and columns, where one row equals one record, such as respondent or interviewee. If these are entered into a spreadsheet such as Excel, this makes simple manipulations and data set filtering easier. Before entering data in this way, it is important to decide how to represent the different possible answers. For example, "don't know" represents a different response from no answer at all, and they need to be distinguished, for instance, with separate columns in the spreadsheet. Also, if dealing with options from a close-ended question, such as job role, there are two different possible approaches that affect the analysis. One approach is to have a column headed "Job role" and to enter the job role as it is given by the respondent or interviewee. The alternative approach is to have a column for each possible answer. The latter approach lends itself more easily to automatic summaries. Note, however, that this option will be open only if the original question was designed to collect the appropriate data (see Box 9.2).

BOX 9.2**How Question Design Affects Data Analysis**

Different question designs affect the kinds of analyses that can be performed and the kinds of conclusions that can be drawn. To illustrate this, assume that some interviews have been conducted to evaluate a new app that lets you try on virtual clothes and see yourself in real time as a 3D holograph. This is an extension of the Memory Mirror described at <http://memorymirror.com>.

Assume that one of the questions asked is: “How do you feel about this new app?” Responses to this will be varied and may include that it is cool, impressive, realistic, clunky, technically complex, and so on. There are many possibilities, and the responses would need to be treated qualitatively. This means that analysis of the data must consider each individual response. If there are only 10 or so responses, then this may not be too bad, but if there are many more, it becomes harder to process the information and harder to summarize the findings. This is typical of open-ended questions; that is, answers are not likely to be homogeneous and so they will need to be treated individually. In contrast, answers to a close-ended question, which gives respondents a fixed set of alternatives from which to choose, can be treated quantitatively. So, for example, instead of asking “How do you feel about the virtual try-on holograph?” assume that you have asked “In your experience, are virtual try-on holographs realistic, clunky, or distorted?” This clearly reduces the number of options and the responses would be recorded as “realistic,” “clunky,” or “distorted.”

When entered in a spreadsheet, or a simple table, initial analysis of this data might look like the following:

| Respondent | Realistic | Clunky | Distorted |
|------------|-----------|--------|-----------|
| A | 1 | | |
| B | | 1 | |
| C | | 1 | |
| ... | | | |
| Z | | | 1 |
| Total | 14 | 5 | 7 |

Based on this, we can then say that 14 out of 26 (54 percent) of the respondents think virtual try-on holographs are realistic, 5 out of 26 (19 percent) think they are clunky, and 7 out of 26 (27 percent) think they are distorted. Note also that in the table, respondents’ names are replaced by letters so that they are identifiable but anonymous to any onlookers. This strategy is important for protecting participants’ privacy.

Another alternative that might be used in a questionnaire is to phrase the question in terms of a Likert scale, such as the following one. This again alters the kind of data and hence the kind of conclusions that can be drawn.

(Continued)

Virtual try-on holographs are realistic:

strongly agree agree neither disagree strongly disagree

The data could then be analyzed using a simple spreadsheet or table:

| Respondent | Strongly agree | Agree | Neither | Disagree | Strongly disagree |
|------------|----------------|-------|---------|----------|-------------------|
| A | 1 | 1 | | | |
| B | | | | | |
| C | | | | 1 | |
| ... | | | | | |
| Z | | | | | 1 |
| Total | 5 | 7 | 10 | 1 | 3 |

In this case, the kind of data being collected has changed. Based on this second set, nothing can be said about whether respondents think the virtual try-on holographs are clunky or distorted, as that question has not been asked. We can only say that, for example, 4 out of 26 (15 percent) disagreed with the statement that virtual try-on holographs are realistic, and of those, 3 (11.5 percent) strongly disagreed. ■

For simple collation and analysis, spreadsheet software such as Excel or Google Sheets is often used as it is commonly available, is well understood, and offers a variety of numerical manipulations and graphical representations. Basic analysis might involve finding out averages and identifying *outliers*, in other words, values that are significantly different from the majority, and hence not common. Producing a graphical representation provides an overall view of the data and any patterns it contains. Other tools are available for performing specific statistical tests, such as online t-tests and A/B testing tools. Data visualization tools can create more sophisticated representations of the data such as heatmaps.


For example, consider the set of data shown in Table 9.2, which was collected during an evaluation of a new photo sharing app. This data shows the users' experience of social media and the number of errors made while trying to complete a controlled task with the new app. It was captured automatically and recorded in a spreadsheet; then the totals and averages were calculated. The graphs in Figure 9.1 were generated using the spreadsheet package. They show an overall view of the data set. In particular, it is easy to see that there are no significant outliers in the error rate data.

Adding one more user to Table 9.2 with an error rate of 9 and plotting the new data as a scatter graph (see Figure 9.2) illustrates how graphs can help to identify outliers. Outliers are usually removed from the main data set because they distort the general patterns. However, outliers may also be interesting cases to investigate further in case there are special circumstances surrounding those users and their session.


These initial investigations also help to identify other areas for further investigation. For example, is there something special about users with error rate 0 or something distinctive about the performance of those who use the social media only once a month?

| Social media use | | | | | | |
|------------------|----------------------|------------|-------------|---------------------------|--------------|-----------------------|
| User | More than once a day | Once a day | Once a week | Two or three times a week | Once a month | Number of errors made |
| 1 | | 1 | | | | 4 |
| 2 | 1 | | | | | 2 |
| 3 | | | 1 | | | 1 |
| 4 | 1 | | | | | 0 |
| 5 | | | | 1 | | 2 |
| 6 | | 1 | | | | 3 |
| 7 | 1 | | | | | 2 |
| 8 | | 1 | | | | 0 |
| 9 | | | | | 1 | 3 |
| 10 | 1 | | | | | 2 |
| 11 | | | | 1 | | 1 |
| 12 | | | 1 | | | 2 |
| 13 | | 1 | | | | 4 |
| 14 | | 1 | | | | 2 |
| 15 | | | | | | 1 |
| 16 | | | | 1 | | 1 |
| 17 | | 1 | | | 1 | 0 |
| 18 | | 1 | | | | 0 |
| Totals | 4 | 7 | 2 | 3 | 2 | 30 |
| | | | | | Mean | 1.67 |
| | | | | | | (to 2 decimal places) |

Table 9.2 Data gathered during a study of a photo sharing app



(a)



(b)

Figure 9.1 Graphical representations of the data in Table 9.2 (a) The distribution of errors made (take note of the scale used in these graphs, as seemingly large differences may be much smaller in reality). (b) The spread of social media experience within the participant group.




Figure 9.2 Using a scatter diagram helps to identify outliers in your data quite quickly

ACTIVITY 9.1


The data in the following table represents the time taken for a group of users to select and buy an item from an online shopping website.

Using a spreadsheet application to which you have access, generate a bar graph and a scatter diagram to provide an overall view of the data. From this representation, make two initial observations about the data that might form the basis of further investigation.


| User | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S |
|-------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Time to complete (mins) | 15 | 10 | 12 | 10 | 14 | 13 | 11 | 18 | 14 | 17 | 20 | 15 | 18 | 24 | 12 | 16 | 18 | 20 | 26 |

Comment

The bar graph and scatter diagram are shown here.



(Continued)



From these two diagrams, there are two areas for further investigation. First, the values for user N (24) and user S (26) are higher than the others and could be looked at in more detail. In addition, there appears to be a trend that the users at the beginning of the testing time (particularly users B, C, D, E, F, and G) performed faster than those toward the end of the testing time. This is not a clear-cut situation, as O also performed well, and I, L, and P were almost as fast, but there may be something about this later testing time that has affected the results, and it is worth investigating further. ■

It is fairly straightforward to compare two sets of results, for instance from the evaluation of two interactive products, using these kinds of graphical representations of the data. Semantic differential data can also be analyzed in this way and used to identify trends, provided that the format of the question is appropriate. For example, the following question was asked in a questionnaire to evaluate two different smartphone designs:

For each pair of adjectives, place a cross at the point between them that reflects the extent to which you believe the adjectives describe the smartphone design. Please place only one cross between the marks on each line.

| | | |
|------------------------|---|-------------------------|
| <i>Annoying</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Pleasing</i> |
| <i>Easy to use</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Difficult to use</i> |
| <i>Value-for-money</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Expensive</i> |
| <i>Attractive</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Unattractive</i> |
| <i>Secure</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Not secure</i> |
| <i>Helpful</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Unhelpful</i> |
| <i>Hi-tech</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Lo-tech</i> |
| <i>Robust</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Fragile</i> |
| <i>Inefficient</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Efficient</i> |
| <i>Modern</i> | <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> | <i>Dated</i> |

Table 9.3 and Table 9.4 show the tabulated results from 100 respondents. Note that the responses have been translated into five categories, numbered from 1 to 5, based on where the respondent marked the line between each pair of adjectives. It is possible that respondents may have intentionally put a cross closer to one side of the box than the other, but it is acceptable to lose this nuance in the data, provided that the original data is not lost, and any further analysis could refer back to it.

| | 1 | 2 | 3 | 4 | 5 | |
|-----------------|----|----|----|----|----|------------------|
| Annoying | 35 | 20 | 18 | 15 | 12 | Pleasing |
| Easy to use | 20 | 28 | 21 | 13 | 18 | Difficult to use |
| Value-for-money | 15 | 30 | 22 | 27 | 6 | Expensive |
| Attractive | 37 | 22 | 32 | 6 | 3 | Unattractive |
| Secure | 52 | 29 | 12 | 4 | 3 | Not secure |
| Helpful | 33 | 21 | 32 | 12 | 2 | Unhelpful |
| Hi-tech | 12 | 24 | 36 | 12 | 16 | Lo-tech |
| Robust | 44 | 13 | 15 | 16 | 12 | Fragile |
| Inefficient | 28 | 23 | 25 | 12 | 12 | Efficient |
| Modern | 35 | 27 | 20 | 11 | 7 | Dated |

Table 9.3 Phone 1

| | 1 | 2 | 3 | 4 | 5 | |
|-----------------|----|----|----|----|----|------------------|
| Annoying | 24 | 23 | 23 | 15 | 15 | Pleasing |
| Easy | 37 | 29 | 15 | 10 | 9 | Difficult to use |
| Value-for-money | 26 | 32 | 17 | 13 | 12 | Expensive |
| Attractive | 38 | 21 | 29 | 8 | 4 | Unattractive |
| Secure | 43 | 22 | 19 | 12 | 4 | Not secure |
| Helpful | 51 | 19 | 16 | 12 | 2 | Unhelpful |
| Hi-tech | 28 | 12 | 30 | 18 | 12 | Lo-tech |
| Robust | 46 | 23 | 10 | 11 | 10 | Fragile |
| Inefficient | 10 | 6 | 37 | 29 | 18 | Efficient |
| Modern | 3 | 10 | 45 | 27 | 15 | Dated |

Table 9.4 Phone 2

The graph in Figure 9.3 shows how the two smartphone designs varied according to the respondents' perceptions of how modern the design is. This graphical notation shows clearly how the two designs compare.




Figure 9.3 A graphical comparison of two smartphone designs according to whether they are perceived as modern or dated

Data logs that capture users' interactions automatically, such as with a website or smartphone, can also be analyzed and represented graphically, thus helping to identify patterns in behavior. Also, more sophisticated manipulations and graphical images can be used to highlight patterns in collected data.

9.4 Basic Qualitative Analysis

Three basic approaches to qualitative analysis are discussed in this section: identifying themes, categorizing data, and analyzing critical incidents. *Critical incident analysis* is a way to isolate subsets of data for more detailed analysis, perhaps by identifying themes or applying categories. These three basic approaches are not mutually exclusive and are often used in combination, for example, when analyzing video material critical incidents may first be identified and then a thematic analysis undertaken. Video analysis is discussed further in Box 9.3.

As with quantitative analysis, the first step in qualitative analysis is to gain an overall impression of the data and to start looking for interesting features, topics, repeated observations, or things that stand out. Some of these will have emerged during data gathering, and this may already have suggested the kinds of pattern to look for, but it is important to confirm and re-confirm findings to make sure that initial impressions don't bias analysis. For example, you might notice from the logged data of people visiting TripAdvisor.com that they often look for reviews for hotels that are rated "terrible" first. Or, you might notice that a lot of respondents all say how frustrating it is to have to answer so many security questions when logging onto an online banking service. During this first pass, it is not necessary to capture all of the findings but instead to highlight common features and record any surprises that arise (Blandford, 2017).

For observations, the guiding framework used in data gathering will give some structure to the data. For example, the practitioner's framework for observation introduced in Chapter 8, "Data Gathering," will have resulted in a focus on who, where, and what, while

using the more detailed framework will result in patterns relating to physical objects, people's goals, sequences of events, and so on.

Qualitative data can be analyzed inductively, that is, extracting concepts from the data, or deductively, in other words using existing theoretical or conceptual ideas to categorize data elements (Robson and McCartan, 2016). Which approach is used depends on the data obtained and the goal of the study, but the underlying principle is to classify elements of the data in order to gain insights toward the study's goal. Identifying themes (thematic analysis) takes an *inductive approach*, while categorizing data takes a *deductive approach*. In practice, analysis is often performed iteratively, and it is common for themes identified inductively then to be applied deductively to new data, and for an initial, pre-existing categorization scheme, to be enhanced inductively when applied to a new situation or new data. One of the most challenging aspects of identifying themes or new categories is determining meaningful codes that are orthogonal (that is, codes which do not overlap). Another is deciding on the appropriate granularity for them, for example at the word, phrase, sentence, or paragraph level. This is also dependent on the goal of the study and the data being analyzed.

Whether an inductive or deductive approach is used, an objective is to produce a reliable analysis, that is, one that can be replicated by someone else if they were to use the same type of approach. One way to achieve this is to train another person to do the coding. When training is complete, both researchers analyze a sample of the same data. If there is a large discrepancy between the two analyses, either training was inadequate or the categorization is not working and needs to be refined. When a high level of reliability is reached between the two researchers, it can be quantified by calculating the inter-rater reliability. This is the percentage of agreement between the analyses of the two researchers, defined as the number of items of agreement, for example the number of categories or themes arising from the data that have been identified consistently by both researchers, expressed as a percentage of the total number of items examined. An alternative measure where two researchers have been used is Cohen's kappa, (κ), which considers the possibility that agreement has occurred due to chance (Cohen, 1960).

Using more sophisticated analytical frameworks to structure the analysis of qualitative data can lead to additional insights that go beyond the results of these basic techniques. Section 9.5 introduces frameworks that are commonly used in interaction design.

BOX 9.3

Analyzing Video Material

A good way to start a video analysis is to watch what has been recorded all the way through while writing a high-level narrative of what happens, noting down where in the video there are any potentially interesting events. How to decide which is an interesting event will depend on what is being observed. For example, in a study of the interruptions that occur in an open plan office, an event might be each time that a person takes a break from an ongoing activity, for instance, when a phone rings, someone walks into their cubicle, or email arrives. If it is a study of how pairs of students use a collaborative learning tool, then activities such as turn-taking, sharing of input devices, speaking over one another, and fighting over shared objects would be appropriate to record.

Chronological and video times are used to index events. These may not be the same, since recordings can run at different speeds from real time and video can be edited. Labels for certain routine events are also used, for instance lunchtime, coffee break, staff meeting, and doctor's rounds. Spreadsheets are used to record the classification and description of events, together with annotations and notes of how the events began, how they unfolded, and how they ended.

Video can be augmented with captured screens or logged data of people's interactions with a computer display, and sometimes transcription is required. There are various logging and screen capture tools available for this purpose, which enable interactions to be played back as a movie, showing screen objects being opened, moved, selected, and so on. These can then be played in parallel with the video to provide different perspectives on the talk, physical interactions, and the system's responses that occur. Having a combination of data streams can enable more detailed and fine-grained patterns of behavior to be interpreted (Heath et al., 2010). ■

9.4.1 Identifying Themes

Thematic analysis is considered an umbrella term to cover a variety of different approaches to examining qualitative data. It is a widely used analytical technique that aims to identify, analyze, and report patterns in the data (Braun and Clarke, 2006). More formally, a *theme* is something important about the data in relation to the study goal. A theme represents a pattern of some kind, perhaps a particular topic or feature found in the data set, which is considered to be important, relevant, and even unexpected with respect to the goals driving the study. Themes that are identified may relate to a variety of aspects: behavior, a user group, events, places or situations where those events happen, and so on. Each of these kinds of themes may be relevant to the study goals. For example, descriptions of typical users may be an outcome of data analysis that focuses on participant characteristics. Although thematic analysis is described in this section on qualitative analysis, themes and patterns may also emerge from quantitative data.

After an initial pass through the data, the next step is to look more systematically for themes across participants' transcripts, seeking further evidence both to confirm and disconfirm initial impressions in all of the data. This more systematic analysis focuses on checking for consistency; in other words, do the themes occur across all participants, or is it only one or two people who mention something? Another focus is on finding further themes that may not have been noticed first time. Sometimes, the refined themes resulting from this systematic analysis form the primary set of findings for the analysis, and sometimes they are just the starting point.

The study's goal provides an orienting focus for the identification and formulation of themes in the first and subsequent passes through the data. For example, consider a survey to evaluate whether the information displayed on a train travel website is appropriate and sufficient. Several of the respondents suggest that the station stops in between the origin and destination stations should be displayed. This is relevant to the study's goal and would be reported as a main theme. In another part of the survey, under further comments you might notice that several respondents say the company's logo is distracting. Although this too is a theme in the data, it is not directly relevant to the study's goals and may be reported only as a minor theme.

Once a number of themes have been identified, it is usual to step back from the set of themes to look at the bigger picture. Is an overall narrative starting to emerge, or are the themes quite disparate? Do some seem to fit together with others? If so, is there an overarching theme? Can you start to formulate a meta-narrative, that is, an overall picture of the data? In doing this, some of the original themes may not seem as relevant and can be removed. Are there some themes that contradict each other? Why might this be the case? This can be done individually, but more often this is applied in a group using brainstorming techniques with sticky notes.

A common technique for exploring data, identifying themes, and looking for an overall narrative is to create an *affinity diagram*. The approach seeks to organize individual ideas and insights into a hierarchy showing common structures and themes. Notes are grouped together when they are similar in some fashion. The groups are not predefined, but rather they emerge from the data. This process was originally introduced into the software quality community from Japan, where it is regarded as one of the seven quality processes. The affinity diagram is built gradually. One note is put up first, and then the team searches for other notes that are related in some way.

Affinity diagrams are used in Contextual Design (Beyer and Holtzblatt, 1998; Holtzblatt, 2001), but they have also been adopted widely in interaction design (Lucero, 2015). For example, Madeline Smith et al. (2018) conducted interviews to design a web app for co-watching videos across a distance, and they used affinity diagramming to identify requirements from interviewee transcripts (see Figure 9.4). Despite the prevalence of digital collaboration tools, the popularity of physical affinity diagramming using sticky notes drawn by hand, has persisted for many years (Harboe and Huang, 2015).



Figure 9.4 Section of an affinity diagram built during the design of a web application

Source: Smith (2018). Used courtesy of Madeline Smith

To read more about the use of affinity diagrams in interaction design, see the following page:

<https://uxdict.io/design-thinking-methods-affinity-diagrams-357bd8671ad4>

9.4.2 Categorizing Data

Inductive analysis is appropriate when the study is exploratory, and it is important to let the themes emerge from the data itself. Sometimes, the analysis frame (the set of categories used) is chosen beforehand, based on the study goal. In that case, analysis proceeds deductively. For example, in a study of novice interaction designer behavior in Botswana, Nicole Lotz et al. (2014) used a set of predetermined categories based on Schön (1983)'s design and reflection cycle: naming, framing, moving, and reflecting. This allowed the researchers to identify detailed patterns in the designers' behavior, which provided implications for education and support.

To illustrate categorization, we present an example derived from a set of studies looking at the use of different navigation aids in an online educational setting (Ursula Armitage, 2004). These studies involved observing users working through some online educational material (about evaluation methods), using the think-aloud technique. The think-aloud protocol was recorded and then transcribed before being analyzed from various perspectives, one of which was to identify usability problems that the participants were having with the online environment known as Nestor Navigator (Zeiliger et al., 1997). An excerpt from the transcription is shown in Figure 9.5.

I'm thinking that it's just a lot of information to absorb from the screen. I just I don't concentrate very well when I'm looking at the screen. I have a very clear idea of what I've read so far . . . but it's because of the headings I know OK this is another kind of evaluation now and before it was about evaluation which wasn't anyone can test and here it's about experts so it's like it's nice that I'm clicking every now and then coz it just sort of organizes the thoughts. But it would still be nice to see it on a piece of paper because it's a lot of text to read.

Am I supposed to, just one question, am supposed to say something about what I'm reading and what I think about it the conditions as well or how I feel reading it from the screen, what is the best thing really?

Observer: What you think about the information that you are reading on the screen . . . you don't need to give me comments . . . if you think this bit fits together.

There's so much reference to all those previously said like I'm like I've already forgotten the name of the other evaluation so it said unlike the other evaluation this one like, there really is not much contrast with the other it just says what it is may be . . . so I think I think of . . .

Maybe it would be nice to have other evaluations listed to see other evaluations you know here, to have the names of other evaluations other evaluations just to, because now when I click previous I have to click it several times so it would be nice to have this navigation, extra links.

Figure 9.5 Excerpt from a transcript of a think-aloud protocol when using an online educational environment. Note the prompt from the observer about halfway through.

Source: Armitage (2004). Used courtesy of Ursula Armitage

This excerpt was analyzed using a categorization scheme derived from a set of negative effects of a system on a user (van Rens, 1997) and was iteratively extended to accommodate the specific kinds of interaction observed in these studies. The categorization scheme is shown in Figure 9.6.

1. Interface Problems

- 1.1. Verbalizations show evidence of dissatisfaction about an aspect of the interface.
- 1.2. Verbalizations show evidence of confusion/uncertainty about an aspect of the interface.
- 1.3. Verbalizations show evidence of confusion/surprise at the outcome of an action.
- 1.4. Verbalizations show evidence of physical discomfort.
- 1.5. Verbalizations show evidence of fatigue.
- 1.6. Verbalizations show evidence of difficulty in seeing particular aspects of the interface.
- 1.7. Verbalizations show evidence that they are having problems achieving a goal that they have set themselves, or the overall task goal.
- 1.8. Verbalizations show evidence that the user has made an error.
- 1.9. The participant is unable to recover from error without external help from the experimenter.
- 1.10. The participant suggests a redesign of the interface of the electronic texts.

2. Content Problems

- 2.1. Verbalizations show evidence of dissatisfaction about aspects of the content of the electronic text.
- 2.2. Verbalizations show evidence of confusion/uncertainty about aspects of the content of the electronic text.
- 2.3. Verbalizations show evidence of a misunderstanding of the electronic text content (the user may not have noticed this immediately).
- 2.4. The participant suggests re-writing the electronic text content.

Identified problems should be coded as [UP, << problem no. >>].

Figure 9.6 Criteria for identifying usability problems from verbal protocol transcriptions

Source: Armitage (2004). Used courtesy of Ursula Armitage

This scheme developed and evolved as the transcripts were analyzed and more categories were identified inductively. Figure 9.7 shows the excerpt from Figure 9.5 coded using this categorization scheme. Note that the transcript is divided up using square brackets to indicate which element is being identified as showing a particular usability problem.

Having categorized the data, the results can be used to answer the study goals. In the earlier example, the study allowed the researchers to be able to quantify the number of usability problems encountered overall by participants, the mean number of problems per participant for each of the test conditions, and the number of unique problems of each type per participant. This also helped to identify patterns of behavior and recurring problems. Having the think-aloud protocol meant that the overall view of the usability problems could take context into account.

[I'm thinking that it's just a lot of information to absorb from the screen. **UP 1.1**] [I just I don't concentrate very well when I'm looking at the screen **UP 1.1**]. I have a very clear idea of what I've read so far . . . [but it's because of the headings **UP 1.1**] I know OK this is another kind of evaluation now and before it was about evaluation which wasn't anyone can test and here it's about experts so it's like it's nice that I'm clicking every now and then coz it just sort of organises the thoughts. [But it would still be nice to see it on a piece of paper **UP 1.10**] [because it's a lot of text to read **UP 1.1**].

Am I supposed to, just one question, am supposed to say something about what I'm reading and what I think about it the conditions as well or how I feel reading it from the screen, what is the best thing really?

Observer: What you think about the information that you are reading on the screen . . . you don't need to give me comments . . . if you think this bit fits together.

[There's so much reference to all those previously said **UP2.1**] [like I'm like I've already forgotten the name of the other evaluation so it said unlike the other evaluation this one like, there really is not much contrast with the other it just says what it is may be . . . so I think I think of . . . **UP 2.2**]

[Maybe it would be nice to have other evaluations listed to see other evaluations you know here, to have the names of other evaluations other evaluations **UP 1.10**] just to, [because now when I click previous I have to click it several times **UP 1.1, 1.7**] [so it would be nice to have this navigation, extra links **UP 1.10**].

Figure 9.7 The excerpt in Figure 9.5 coded using the categorization scheme in Figure 9.6

Source: Armitage (2004). Used courtesy of Ursula Armitage

ACTIVITY 9.2

The following is a think-aloud extract from the same study of users working through online educational material. Using the categorization scheme in Figure 9.6, code this extract for usability problems. It is useful to put brackets around the complete element of the extract that you are coding.

Well, looking at the map, again there's no obvious start point, there should be something highlighted that says 'start here.'

Ok, the next keyword that's highlighted is evaluating, but I'm not sure that's where I want to go straight away, so I'm just going to go back to the introduction.

Yeah, so I probably want to read about usability problems before I start looking at evaluation. So, I, yeah. I would have thought that the links in each one of the pages would take you to the next logical point, but my logic might be different to other people's. Just going to go and have a look at usability problems.

Ok, again I'm going to flip back to the introduction. I'm just thinking if I was going to do this myself I would still have a link back to the introduction, but I would take people through the logical sequence of each one of these bits that fans out, rather than expecting them to go back all the time.

Going back . . . to the introduction. Look at the types. Observation, didn't really want to go there. What's this bit [pointing to Types of UE on map]? Going straight to types of . . .

Ok, right, yeah, I've already been there before. We've already looked at usability problems, yep that's ok, so we'll have a look at these references.

I clicked on the map rather than going back via introduction, to be honest I get fed up going back to introduction all the time.

Comment

Coding transcripts takes practice, but this activity will give you an idea of the kinds of decisions involved in applying categories. Our coded extract is shown here:

[Well, looking at the map, again there's no obvious start point **UP 1.2, 2.2**], [there should be something highlighted that says 'start here' **UP 1.1, 1.10**].

Ok, the next keyword that's highlighted is evaluating, but [I'm not sure that's where I want to go straight away **UP 2.2**], so I'm just going to go back to the introduction.

Yeah, so I probably want to read about usability problems before I start looking at evaluation. So, I, yeah. [I would have thought that the links in each one of the pages would take you to the next logical point, but my logic might be different to other people's **UP 1.3**]. Just going to go and have a look at usability problems.

Ok, again I'm going to flip back to the introduction. [I'm just thinking if I was going to do this myself I would still have a link back to the introduction, but I would take people through the logical sequence of each one of these bits that fans out, rather than expecting them to go back all the time **UP 1.10**].

Going back . . . to the introduction. [Look at the types. Observation, didn't really want to go there. What's this bit [pointing to Types of UE on map]? **UP 2.2**] Going straight to types of . . .

Ok, right, yeah, I've already been there before. We've already looked at usability problems, yep that's ok, so we'll have a look at these references.

I clicked on the map rather than going back via introduction, [to be honest I get fed up going back to introduction all the time. **UP 1.1**]. ■

9.4.3 Critical Incident Analysis

Data gathering sessions can often result in a lot of data. Analyzing all of this data in any detail is very time-consuming and often not necessary. *Critical incident analysis* is one approach that helps to identify significant subsets of the data for more detailed analysis. This technique is a set of principles that emerged from work carried out in the United States Army Air Forces where the goal was to identify the critical requirements of good and bad performance by pilots (Flanagan, 1954). It has two basic principles: “(a) reporting facts regarding behavior is preferable to the collection of interpretations, ratings, and opinions based on general impressions; (b) reporting should be limited to those behaviors which, according to competent observers, make a significant contribution to the activity” (Flanagan, 1954, p. 355). In the interaction design context, the use of well-planned observation sessions satisfies the first principle. The second principle refers to critical incidents, that is, incidents that are significant or pivotal to the activity being observed, in either a desirable or an undesirable way.

In interaction design, critical incident analysis has been used in a variety of ways, but the main focus is to identify specific incidents that are significant and then to focus on these and analyze them in detail, using the rest of the data collected as context to inform interpretation. These may be identified by the users during a retrospective discussion of a recent event or by an observer either through studying video footage or in real time. For example, in an

evaluation study, a critical incident may be signaled by times when users were obviously stuck—usually marked by a comment, silence, looks of puzzlement, and so on. These are indications only. Whether the incident is significant enough to be worthy of further investigation depends on the severity of the problem identified.

Tuomas Kari et al. (2017) used the critical incident technique in a study of the location-based augmented reality game Pokémon GO. They were interested in identifying the types of behavior change that playing the game induced in players. To do this, they distributed a survey through social media channels asking experienced players to identify and describe one outstanding positive or negative experience. The 262 valid responses were themed and categorized into eight groups. Apart from expected behavior change such as increased physical activity, they also found that players were more social, found their routines more meaningful, expressed more positive emotions, and were more motivated to explore their surroundings. In another study, Elise Grison et al. (2013) used the critical incident technique to investigate specific factors that influence travelers' choices of transport mode in Paris in order to adapt new tools and services for mobility, such as dynamic route planners. Participants were asked to report on positive and negative real events they experienced in the context of their route to work or study, and whether they regretted or were satisfied with this choice of transport. Their findings included that contextual factors have a great influence on choice, that people were more likely to choose an alternative route to return home than when setting out, and that emotional state is important when planning a route.

To read more on critical incident analysis in usability studies where the emphasis is on understanding the cause of problems, see this site:
www.usabilitybok.org/critical-incident-technique.

ACTIVITY 9.3

Assign yourself or a friend the task of identifying the next available theater performance or movie that you'd like to attend in your area. As you perform this task, or watch your friend do it, make a note of critical incidents associated with the activity. Remember that a critical incident may be a positive or a negative event.

Comment

Information about entertainment may be available through the local newspaper, searching online, looking at social media to see what is recommended by friends, or contacting local cinemas or theaters directly. When this author asked her daughter to attempt this task, several critical incidents emerged, including the following:

1. After checking her social media channels, nothing in the recommendations appealed to her and so she decided to search online.
2. She found that one of her all-time favorite movie classics was playing at the local movie theater for just one week, and tickets were still available.
3. When trying to buy the tickets, she discovered that she needed a credit card, which she didn't have, and so she had to ask me to complete the purchase! ■

9.5 Which Kind of Analytic Framework to Use?

There are several different analytical frameworks that can be used to analyze and interpret data from a qualitative study. In this section, six different approaches are outlined, ordered roughly in terms of their granularity, that is, the level of detail involved. For example, conversation analysis has a fine level of granularity, and it allows the details of what is said and how during a short fragment of conversation to be examined, while systems-based frameworks take a broader scope and have a coarser level of granularity, such as what happens when a new digital technology is introduced into an organization, like a hospital. Conversation analysis may result in insights related to users' interactions through a collaboration technology, while systems-based analyses may result in insights related to changes in work practices, worker satisfaction, improvements in workflow, impact on an office culture, and so on. Table 9.5 lists the six approaches in terms of the main types data, focus, expected outcomes, and level of granularity.

| Framework | Data | Focus | Expected outcomes | Level of granularity |
|--------------------------|---|--|---|---|
| Conversation analysis | Recordings of spoken conversations | How conversations are conducted | Insights into how conversations are managed and how they progress | Word-level, or finer, for instance, pauses and inflection |
| Discourse analysis | Recordings of speech or writing from individuals or several participants | How words are used to convey meaning | Implicit or hidden meanings in texts | Word, phrase, or sentence-level |
| Content analysis | Any form of "text" including written pieces, video and audio recordings, or photographs | How often something is featured or is spoken about | Frequency of items appearing in a text | A wide range of levels from words, to feelings or attitudes, to artifacts or people |
| Interaction analysis | Video recordings of a naturally-occurring activity | Verbal and non-verbal interactions between people and artifacts | Insights about how knowledge and action are used within an activity | At the level of artifact, dialogue, and gesture |
| Grounded theory | Empirical data of any kind | Constructing a theory around the phenomenon of interest | A theory grounded in empirical data | Varying levels, depending on the phenomenon of interest |
| Systems-based frameworks | Large-scale and heterogeneous data | Large-scale involving people and technology, such as a hospital or airport | Insights about organizational effectiveness and efficiency | Macro-level, organizational level |

Table 9.5 Overview of analytical frameworks used in interaction Design

9.5.1 Conversation Analysis

Conversation analysis (CA) examines the semantics of a conversation in fine detail. The focus is on how a conversation is conducted (Jupp, 2006). This technique is used in sociological studies, and it examines how conversations start and how turn-taking is structured, together with other rules of conversation. It has been used to analyze interactions in a range of settings, and it has influenced designers' understanding of users' needs in these environments. It can also be used to compare conversations that take place through different media, for example, face-to-face conversations versus those conducted through social media. More recently, it has been used to analyze the conversations that take place with voice-assisted technologies and chatbots.

Voice assistants (also called smart speakers), like Amazon Echo, have become increasingly popular in domestic settings, providing a limited kind of conversational interaction, mainly by answering questions and responding to requests. But how do families orient and adapt to them? Does using this device change the way they talk, or do they talk to the device as if it was another human being?

Martin Porcheron et al. (2018) carried out a study examining how such devices were being used by families in their own homes, and they used conversation analysis with excerpts from selected conversations. A sample fragment of a conversation they analyzed is presented in Figure 9.8. This uses a particular type of syntax for marking up the minutiae of interactions and speech that took place during an approximate 10-second period. Square parentheses are used to show overlapping talk, round parentheses to indicate a pause, and physical spacing to show temporal sequencing of what is said. This level of detail enables the analysis to reveal subtle cues and mechanisms that are used during the conversations.

```

01 SUS i'd like to play beat the intro in a minute
02 LIA [ oh no: ]
03 SUS [ alexa ][( 1.1 )] beat the in[ro
04 CAR [ °yeah°; ]
05 LIA [°no:.....°
06 CAR (0.6) it's mother's day? (0.4)
07 SUS it's ( ) yep (.) listen (.) you need to keep
08 on eating your orange stuff (.) liam
09 (0.7)
10 CAR and your green stuff
11 SUS alexa (1.3) alexa (0.5)=
12 CAR =°and your brown stuff
13 SUS play beat the intro

```

Figure 9.8 An extract of the conversation between the family and Alexa, marked up for conversation analysis

Source: Porcheron et al. (2018), fragment 1. Reproduced with permission of ACM Publications

In this fragment, Susan (who is the mother) announces to Liam (her son) and Carl (the father) her desire to play a particular game (called Beat the Intro) with their Amazon Echo. Liam does not want to play (expressed by his long “no” cry in response). Susan, however, has

already called “Alexa” to wake up the device. Carl shows his support for her, as indicated by his quick “yeah” during the pause after she says “Alexa.” Alexa, however, appears not to respond. At this point, Susan returns to the ongoing family conversation, telling Liam to keep eating his “orange stuff.” Carl also chips in after her and says to Liam that he should also eat the “green stuff.” At the same time, Susan has another go at getting Alexa to wake up. She calls out “Alexa” twice in a questioning voice. In the pause between Susan’s two calls to Alexa, Carl tells Liam to keep eating again, but this time the “brown stuff.” Having succeeded in waking up Alexa, Susan then asks it to play the game.

So, what insights does this fine level of analysis provide? Martin Porcheron et al. (2018) point out that it demonstrates how a family’s interaction with the Amazon Echo is seamlessly interwoven with other ongoing activities, in this case, the parents trying to get their child to eat his food. At a more general level, it illustrates how our conversations with each other and voice-assisted technologies interleave in nuanced ways rather than being separate conversations between members of the family or the family and the device, which jump from one to another. They also show how their analysis led them to think that the term *conversational interaction* fails to distinguish between the interactional embeddedness of voice-assisted interfaces and human conversation. Instead, they suggest that current voice-assisted technologies be designed using a conceptual model, more akin to instructing rather than conversing.

9.5.2 Discourse Analysis

Discourse analysis focuses on dialogue, in other words, the meaning of what is said and how words are used to convey meaning. Discourse analysis is strongly interpretive, pays great attention to context, and views language not only as reflecting psychological and social aspects but also as constructing them (Coyle, 1995). An underlying assumption of discourse analysis is that there is no objective scientific truth. Language is a form of social reality that is open to interpretation from different perspectives. In this sense, the underlying philosophy of discourse analysis is similar to that of ethnography. Language is viewed as a constructive tool, and discourse analysis provides a way of focusing on how people use language to construct versions of their worlds (Fiske, 1994).

Small changes in wording can change meaning, as the following excerpts indicate (Coyle, 1995):

Discourse analysis is what you do when you are saying that you are doing discourse analysis. . .

According to Coyle, discourse analysis is what you do when you are saying that you are doing discourse analysis. . .

By adding just three words, “According to Coyle,” the sense of authority changes, depending on what the reader knows about Coyle’s work and reputation.

Discourse analysis is useful when trying to identify subtle and implicit meaning in what people are writing about, what is trending, what is fake news, and so on. It can be used with data from interviews; in social media such as Facebook, Twitter, and WhatsApp; and in emails. For example, Carlos Roberto Teixeira et al. (2018) proposed a taxonomy for analyzing tweets as a way of understanding the types of conversations that take place online and identifying patterns of behavior within them. The tweets they analyzed were those posted during the political scandals that were happening in Brazil during 2016. Using discourse

analysis and by understanding the cultural background and how technology was adopted and used by the Brazilian society, they were able to interpret the meaning of the tweets, beyond the words that were said.

They scraped the raw data of the tweets from the web and ranked them in order of those most retweeted. Then they selected 100 of the most influential postings, which were messages that had been retweeted between 9,000 and 47,000 times. Finally, they manually classified the tweets using Excel spreadsheets, identifying the most dominant discourse characteristics. They note, however, that since the most retweeted messages are often captured out of context (that is, outside the conversational thread), they had problems sometimes interpreting what was the context, such as the particular news story associated with a tweet.

The tweets were then classified into a number of provisional themes that were whittled down to five general ones. These were (1) “support” (tweets that promote either side of the political dispute); (2) “criticism and protest” (tweets that showed disapproval or an objection); (3) “humor” (tweets that were witty and had cartoons or jokes in them); (4) “news” (tweets that refer to the news that were neutral in tone); and (5) “neutral” (tweets that were indifferent or their position could not be inferred). For reliability, two different researchers classified each tweet using these themes. Once classified, the data was analyzed using descriptive statistics and simple visualizations, like word (or tag) clouds and pie charts. These allowed them to draw conclusions about the tweets that they had analyzed, including the relative size of the different themes and how the size changed over time. Tweets in the criticism and protest theme were the most popular overall, followed by humor. More generally, they discuss how this kind of discourse analysis shows that humor, protest, and criticism are highly dominant in this kind of online discourse. They suggest that people who tweet often express their feelings about ongoing political events, in terms of criticism and humor, in equal measures. Moreover, they found that tweets including images, videos, and animated GIFs were most frequently classified as humor.

This kind of analysis of public discourse, when done by hand, is extremely time-consuming. To help, there are new software tools being developed that can automatically process computer-mediated discourses (Ecker, 2016). The advantage, as you will see in the next chapter, is that much larger data sets can be analyzed. The downside is that the analyst is no longer “hands-on” and loses touch with the surrounding context, meaning different interpretations arise.

9.5.3 Content Analysis

Content analysis typically involves classifying the data into themes or categories and then studying the frequency of category occurrences (Krippendorff, 2013). The technique can be used for any text, where “text” refers to a range of media including video, newspapers, advertisements, survey responses, images, sounds, and so on. It can be used to analyze any online content, including the text of tweets, links, animated gifs, videos, and images. For example, Mark Blythe and Paul Cairns (2009) analyzed 100 videos from a YouTube search by relevance for “iPhone 3G” using content analysis. They categorized the videos into seven categories: review, reportage, “unboxing,” demonstration, satire, advertisement, and vlog commentaries (such as, complaints about queues). Similar to discourse analysis, an important aspect of content analysis is how it considers the wider context.

Content analysis is often used in conjunction with other analysis techniques as well. For example, Weixin Zhai and Jean-Claude Thill (2017) analyzed social media data from Weibo (the Chinese equivalent of Twitter) to investigate the emotions, attitudes, and views

of citizens around a rainstorm that hit Beijing on July 21, 2012, causing 79 deaths. They used content analysis alongside *sentiment analysis*, an approach that extracts emotional and subjective information from natural language. From their results, they found how feelings of sorrow and sadness were shared across the entire city because of the trauma associated with entrapment indoors during the deluge.

9.5.4 Interaction Analysis

Interaction analysis was developed by Brigitte Jordan and Austin Henderson (1995) as a way of investigating and understanding the interactions of human beings with each other and objects in their environment. The technique focuses on both talk and nonverbal interactions with artifacts and technologies, and it is based on video recordings. An underlying assumption of this approach is that knowledge and action are fundamentally social. The goal is to derive generalizations from videos of naturally occurring activities, focusing on how the people being observed make sense of each other's actions and their collective achievements.

Interaction analysis is an inductive process, where teams of researchers suggest statements about general patterns from multiple examples of empirical observations. Rather than individual researchers conducting separate analyses and then comparing their results for consistency, interaction analysis is conducted collaboratively; teams discuss together their observations and interpretations of the videos being watched as they watch them. The first step involves creating a content log, comprising headings and rough summaries of what has been observed. No predetermined categories are used during this stage. Instead, they emerge from repeated playing and discussion of the video material. Hypotheses are also generated by group members about what they think is happening. This process includes suggesting the intentions, motivations, and understandings of the people who are being viewed in the videos. These suggestions have to be tied to the actions of the people rather than being purely speculative. For example, if an analyst thinks someone's motivation is to take control during a board meeting, they need to provide actual examples that demonstrate how the person is achieving this (for instance, taking over the data projector, as the locus of control, for long periods of time and presenting their ideas for others to view).

The videos are then *cannibalized*, as it is called, by extracting interesting materials, reclassifying some of them in terms of what they represent, while removing others. Instances of a salient event are assembled and played one after another to determine whether a phenomenon is a robust theme or a one-off incident. An example Brigitte Jordan and Austin Henderson (1995) use to illustrate this process is their study of people around a pregnant woman who was having her first contraction. They noticed that, at the point of the first contraction, the medical staff and family all shifted their attention away from the woman to the monitoring equipment. They were able to find many more examples of this phenomenon, providing strong evidence that the presence of high-tech equipment changes the practice of caregiving, specifically that caregiving is mediated by the real-time data presented through the equipment.

An example of where interaction analysis has been used in HCI is Anna Xambo et al.'s (2013) study of how groups of musicians improvise and learn to play together when using a novel collaborative tabletop technology. Video data was collected for four groups of musicians using this technology in a number of jamming sessions. Representative video extracts were repeatedly viewed and discussed by a team of researchers, focusing on verbal communication and nonverbal communication themes. These themes were categorized into whether

they were musical, physical, or interface-related. The themes that arose during the repeated video viewing process were diverse, including error/repair situations, territory-related behaviors, emergent coordination mechanisms, and mimicking behaviors. Anna Xambo et al. (2013) then transformed these into a set of design considerations for developing tabletop technologies to support hands-on collaborative learning.

9.5.5 Grounded Theory

The goal of *grounded theory* is to develop theory from a systematic analysis and interpretation of empirical data; that is, the derived theory is grounded in the data. In this respect, it is an inductive approach to developing theory. The approach was originally developed by Barney Glaser and Anselm Strauss (1967) and since then has been adopted by several researchers, with some adaptations to different circumstances. In particular, both of them have individually (and with others) developed grounded theory in slightly different ways, but the objective of this approach remains the same. Barney Glaser (1992) provides further information about the differences and areas of controversy.

In this context, theory is: “a set of well-developed concepts related through statements of relationship, which together constitute an integrated framework that can be used to explain or predict phenomena” (Strauss and Corbin, 1998). Development of a “grounded” theory progresses through alternating data collection and data analysis: first data is collected and analyzed to identify themes, then that analysis may lead to further data collection and analysis to extend and refine the themes, and so on. During this cycle, parts of the data may be reanalyzed in more detail. Data gathering and subsequent analysis are hence driven by the emerging theory. This approach continues until no new insights emerge and the theory is well-developed. During this process, the researcher seeks to maintain a balance between objectivity and sensitivity. Objectivity is needed to maintain accurate and impartial interpretation of events; sensitivity is required to notice the subtleties in the data and identify relationships between concepts.

The thrust of the analysis undertaken is to identify and define the properties and dimensions of relevant themes called *categories* in grounded theory. According to Juliet Corbin and Anselm Strauss (2014), this coding has three aspects, which are iteratively performed through the cycle of data collection and analysis:

1. *Open coding* is the process through which categories, their properties, and dimensions are discovered in the data. This process is similar to our discussion of thematic analysis above, including the question of granularity of coding (at the word, line, sentence, conversation level, and so on).
2. *Axial coding* is the process of systematically fleshing out categories and relating them to their subcategories.
3. *Selective coding* is the process of refining and integrating categories to form a larger theoretical scheme. The categories are organized around one central category that forms the backbone of the theory. Initially, the theory will contain only an outline of the categories, but as more data is collected, they are refined and developed further.

Early books on grounded theory say little about what data collection techniques should be used but focus instead on the analysis. Some later books place more emphasis on data collection. For example, Kathy Charmaz (2014) discusses interviewing techniques and collection and analysis of documents for grounded theory analysis. When analyzing

data, Juliet Corbin and Anselm Strauss (2014) encourage the use of written records of analysis and diagrammatic representations of categories (which they call *memos* and *diagrams*). These memos and diagrams evolve as data analysis progresses. Some researchers also look to digital tools such as spreadsheets and diagramming tools, but many like to develop their own physical code books such as the one Dana Rotman et al. (2014) constructed in a study to understand the motivations of citizens to contribute to citizen science projects. The data that she analyzed was derived from in-depth semistructured interviews of 33 citizen scientists and 11 scientists from the United States, India, and Costa Rica (see Figure 9.9).

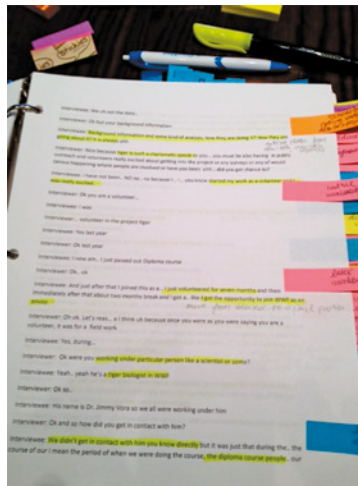


Figure 9.9 Code book used in a grounded theory analysis of citizens' motivations to contribute to citizen science

Source: Jennifer Preece

The following analytic tools are used to help stimulate the analyst's thinking and identify and characterize relevant categories:

1. **The Use of Questioning:** In this context, this refers to questioning the data, not your participants. Questions can help an analyst to generate ideas or consider different ways of looking at the data. It can be useful to ask questions when analysis appears to be in a rut.
2. **Analysis of a Word, Phrase, or Sentence:** Considering in detail the meaning of an utterance can also help to trigger different perspectives on the data.
3. **Further Analysis Through Comparisons:** Comparisons may be made between objects or between abstract categories. In either case, comparing one with the other brings alternative interpretations.

Grounded theory uses thematic analysis that is, themes are identified from the data, but as data analysis informs data collection, it also relies on categorizing new data according to

the existing thematic set and then evolving that set to accommodate new findings. As Victoria Braun and Victoria Clarke (2006) point out, “Thematic analysis differs from other analytic methods that seek to describe patterns across qualitative data . . . such as grounded theory <which is> theoretically bounded. . . the goal of a grounded theory analysis is to generate a plausible—and useful—theory of the phenomena that is grounded in the data.”

A Grounded Theory Example

So-called idle games are rising in popularity (Cutting et al., 2019). *Idle games* are minimalist games that require little or even no interaction in order for the game to progress. For example, an idle game may involve repeating a simple action like clicking an icon to accumulate resources. Example games include the Kittens Game, which is a text-based (that is, it has no graphical user interface) game that involves managing a village of kittens, while Cookie Clicker involves baking and selling cookies. Idle games also include mechanisms to automate game play so that progress may continue for extended periods of time without the player doing anything (Purkiss and Khaliq, 2015). An extreme example studied by Joe Cutting et al. (2019) is Neko Atsume, a game about collecting cats, in which progress can be made only if the game is switched off. In their study, they were interested in the notion of “engagement” and the implications of this new genre of games for current theories of engagement.

To understand more about the idle games genre, Sultan Alharthi et al. (2018) used grounded theory, specifically to develop a taxonomy and a set of characteristics for them. By defining the essential features of this genre of games and clustering them, the authors hoped to produce design implications for each type of game.

The three stages of coding, open, axial, and selective are illustrated in Figure 9.10. Note that, in this case, the research started by the researchers playing each of the games under study.

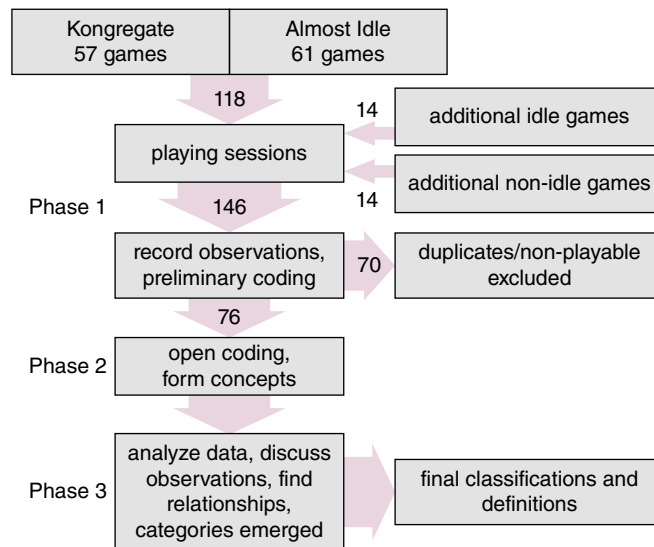


Figure 9.10 The process used by Alharthi et al., showing Phase 2 and Phase 3 using the three stages of grounded theory coding

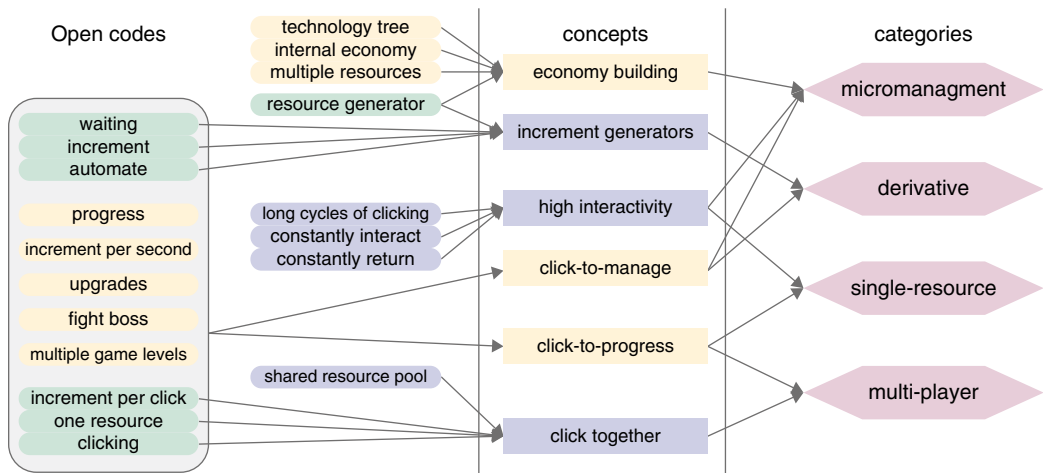
Each game was played by two researchers who recorded their observations in a spreadsheet. These observations focused on gameplay, game mechanics, rewards, interactivity, progress rate, and game interface. Then they rated the games using an 11-point interactivity scale (0–10) where 0 meant that play progressed without any interaction of the player, while 10 meant that the game progressed only slowly without player interaction. Progress through the levels of the game were also rated on the same scale.

At the end of each game session and observation, the researchers wrote a brief overview of the game and conducted preliminary open coding of their observations (see Figure 9.11 for an example of preliminary open coding).

| Game Feature | Observations |
|---------------------|--|
| Game name | <i>AdVenture Capitalist</i> [G38] |
| Play description | You start CLICKING on a lemonade stand and collect money. Spend money to make upgrades, INCREASE PRODUCTION PER CLICK . Start hiring workers and INCREASE PRODUCTION PER SECOND . When you have enough money, you can buy new businesses, automate all your businesses to INCREMENT more money, and leave the game progress. |
| Game mechanics | Click to gain money, AUTOMATE production, make upgrades to DAMAGE/SEC . |
| Rewards | ONE CURRENCY , which is money, is rewarded in return. |
| Interface | GRAPHICAL |
| Interactivity level | 7 |
| Progress rate | 9 |
| Overview | This is a SINGLE-PLAYER game, which requires LONG CYCLES OK CLICKING at the start, and making a number of upgrades. Production rate reaches \$390/sec in less that 10 minutes and you gain 1M in cash making the game progress faster. |

Figure 9.11 An illustration of preliminary open coding. Words in small capitals are identified by the researcher as potential codes.

Axial and selective coding progressed iteratively. The researchers held several discussion sessions to explore the relationship between the codes, the emergent concepts, and the initial categories. During this process, some of the games were re-observed, and related literature was drawn upon to help refine the concepts. For example, existing literature on game taxonomies, prior terms, and definitions related to idle games were incorporated into the analysis. Based on this analysis, Sultan Alharthi et al. produced a taxonomy with two basic ways to characterize the games: one based on key features and one based on interactivity. From the former, they defined incremental games as idle games in which a player selects resources to generate, waits for them to accumulate, and spends the resources to automate resource generation. Figure 9.12 illustrates the open codes, resulting concepts, and categories developed for incremental games. This shows that four categories of incremental games emerge from this analysis: micromanagement, derivative, single-resource, and multiplayer.



The analysis process that developed the incremental games super-category (each category above is part of incremental games). The process started with open coding of observations on idle games: multiple codes are created. Concepts are discovered through analyzing the open codes and identifying common features. This is an iterative process, where new codes are added, combined, or deleted. Each code is connected to one or more games and can be combined to form new concepts. Concepts are analyzed to find common relationships, and, thus, categories emerge. In the diagram, coloration is only to aid in reading. The left grouping is to show that all contained codes are part of click-to-manage and click-to-progress.

Figure 9.12 The grounded theory process showing the development of open coding, through concepts to categories

One of the surprises from their analysis is that the resulting taxonomy is based on game rules and their basic underlying structure rather than mechanisms in the game. This is in contrast to other gaming taxonomies that feature interactivity and interaction strategies. Of course, idle games are minimal, and so they don't have much interaction. However, the grounded theory approach allowed the development of a taxonomy that reflects the style and purpose of the genre.

9.5.6 Systems-Based Frameworks

For large projects where the researcher is interested in investigating how a new technology should best be introduced and what its impact is afterward, it is necessary to analyze many sources of data collected over a long period of time. Conducting analyses of small fragments of conversation or identifying themes from interviews may be useful for highlighting specific working practices, but understanding how a whole socio-technical system (for example a hospital, corporation, local council, or airport) works at scale requires a different kind of analytical framework. Two such frameworks are introduced next: socio-technical systems theory (Eason, 1987) and distributed cognition (Hutchins, 1995), as applied through the Distributed Cognition of Teamwork framework (Furniss and Blandford, 2006).

Socio-technical Systems Theory

Socio-technical systems (STS) theory makes explicit the fact that the technology and the people in a work system are interdependent (Klein, 2014). Rather than trying to optimize either the technical system or the social system independently of each other, STS suggests that this interdependency be recognized, and the "system" be treated as a whole. The ideas behind

socio-technical theory were first conceptualized around coal mining in the 1950s (see Trist and Bamford, 1951, for example), but it also has a long history of being applied in hospitals and healthcare settings (Waterson, 2014) as well as manufacturing and social media systems. Martin Maguire (2014) highlights the importance of the socio-technical perspective with the rise of virtual organizations. Ken Eason (2014) identifies five significant and enduring aspects of STS theory (Eason, 2014):

1. **Task interdependencies:** If people are focused on one large task, then the division of sub-tasks between them inevitably sets up interdependencies that are critical to understand. Understanding these interdependencies is particularly useful for recognizing the implications of change.
2. **Socio-technical systems are “open systems”:** STS are influenced by environmental factors including physical disturbances and financial, market, regulatory, and technical developments.
3. **Heterogeneity of system components:** The overall task is undertaken by humans in the social subsystem using technical resources in the technical subsystem. Both need to be resilient. Technical components can evolve while humans can learn, develop, and change the technical components to address challenges of the future.
4. **Practical contributions:** STS theory is making practical contributions in analysis of existing systems, summative evaluation of a major change, through potentially predicting challenges before changes are made, and in designing socio-technical systems that are co-optimized.
5. **Fragmentation of design processes:** In a complex socio-technical system, there are different design processes, and these can result in fragmentation. Flexibility in specification, local focus in design, user-centered design, and system evolution will help overcome these.

STS is a philosophy rather than a concrete set of methods or analytical tools. But several socio-technical design methods provide more concrete tools for using a socio-technical framework. For example, see Baxter and Sommerville (2011) and Mumford (2006).

Distributed Cognition of Teamwork

Distributed cognition and *Distributed Cognition of Teamwork* (DiCoT) were introduced in Chapter 4, “Cognitive Aspects,” as an approach to studying the nature of cognitive phenomena across individuals, artifacts, and internal and external representations. Investigating how information is propagated through different media is a key goal of this approach, and while distributed cognition provides a good theoretical framework for analyzing systems, it can be difficult to apply in practice. The DiCoT framework was developed as a method to support the application of distributed cognition. It provides a framework of models that can be constructed from a set of collected data, for example ethnographic, interview transcripts, artifacts, photographs, and so on. Underlying each model is a set of principles distilled from the distributed cognition theory. The models are as follows:

- *An information flow model that shows how information flows through the system and is transformed.* This model captures the information channels and hubs together with the sequence of activities and communication between different team roles.
- *A physical model that captures how physical structures support communication between the team roles and facilitates access to artifacts.* This model helps to describe the factors that influence the performance of the system at a physical level.

- *An artifact model that captures how artifacts in this system support cognition.* This model can be used to represent the key characteristics of an artifact and how its design, structure, and use can support team members.
- *A social structure model that examines how cognition is socially distributed.* This model maps the social structures to the goal structures, shows how work is shared, and can be used to consider the robustness of the system.
- *A system evolution model that depicts how the system has evolved over time.* This model provides some explanation for why the work is the way it is. Any design recommendations need to take this context into account.

While the form of the models is not prescribed, the underlying principles support the models' development. For example, underlying the physical model are principles such as the following:

- **Horizon of observation:** What an individual can see or hear.
- **Perceptual:** How spatial representations aid computation.
- **Arrangement of equipment:** How the physical arrangement of the environment affects access to information.

DiCoT has been used to understand collaborative work in remote and co-located software development teams, as shown in Figure 9.13 (Deshpande et al., 2016; Sharp et al. 2009), and has been found to be especially useful for studying how medical teams work and manage with ever-changing technologies that are introduced into their work environment. For example, Atish Rajkomar and Ann Blandford (2012) examined how healthcare technologies are used; specifically, they examined the use of infusion pumps by nurses in an intensive care

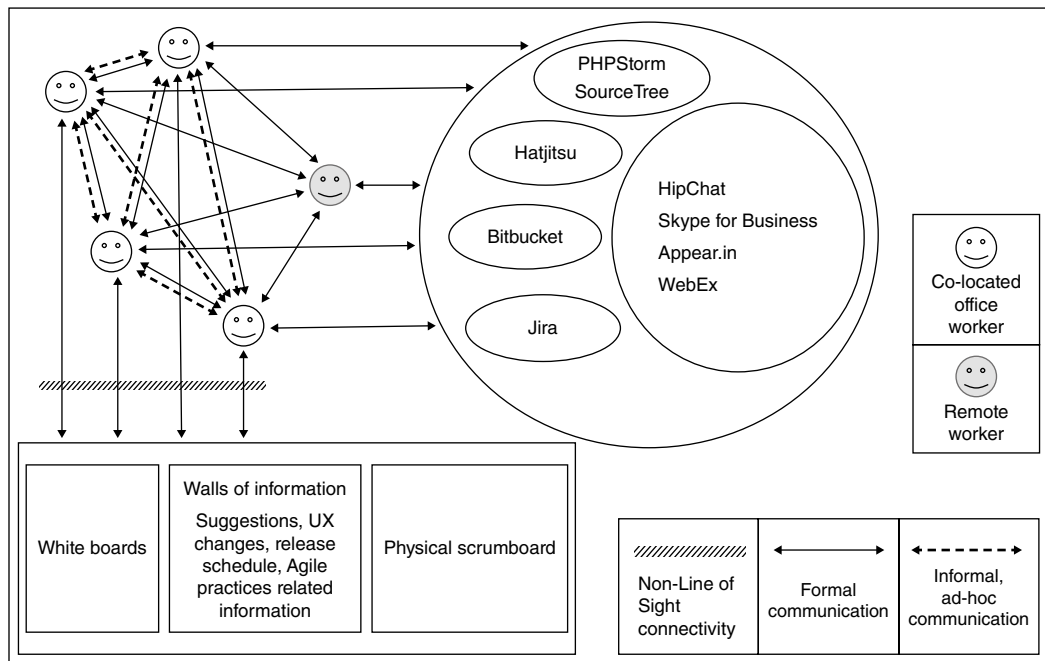


Figure 9.13 An information flow diagram from a DiCoT analysis of software development remote work, based on ethnographic data

unit (ICU). They gathered data through ethnographic observations and interviews, which they analyzed by constructing representational models that focused on information flows, physical layouts, social structures, and artifacts. They note that “the findings showed that there was significant distribution of cognition in the ICU: socially, among nurses; physically, through the material environment; and through technological artefacts.” Based on the results of this study, they were able to suggest changes that would improve the safety and efficiency of the nurses’ interactions with the infusion technology.

9.6 Tools to Support Data Analysis

While it is possible to perform these kinds of data analysis using only manual techniques, most people would agree that it is quicker, easier, and more accurate to use a software tool of some kind in the majority of cases. Using a simple spreadsheet application is surprisingly effective, but there are other more sophisticated tools available to support the organization, coding, and manipulation of data, and to perform statistical tests.

Tools in the former category (to support the organization of data) include facilities for categorization, theme-based analysis, and quantitative analysis. These typically provide facilities to associate labels (categories, themes, and so on) with sections of data, search the data for key words or phrases, investigate the relationships between different themes or categories, and help to develop the coding scheme further. Some tools can also generate graphical representations. In addition, some provide help with techniques such as content analysis and sometimes mechanisms to show the occurrence and co-occurrence of words or phrases. In addition, searching, coding, project management, writing and annotating, and report generation facilities are common.

Two well-known tools that support some of these data analysis activities are Nvivo and Dedoose. For example, Nvivo supports the annotation and coding of data including PDF documents, photos, and video and audio files. Using Nvivo, field notes can be searched for key words or phrases to support coding or content analysis; codes and data can be explored, merged, and manipulated in several ways. The information can also be printed in a variety of forms such as a list of every occasion a word or phrase is used in the data, and a tree structure showing the relationships among codes. Like all software packages, Nvivo has advantages and disadvantages, but it is particularly powerful for handling large sets of data and can generate output for statistical packages such as SAS and SPSS.

Statistical Analysis Software (SAS) and Statistical Package for the Social Sciences (SPSS) are popular quantitative analysis packages that support the use of statistical tests. SPSS, for example, is a sophisticated package offering a wide range of statistical tests such as frequency distributions, rank correlations (to determine statistical significance), regression analysis, and cluster analysis. SPSS assumes that the user knows and understands statistical analysis.

Additional tools to support the analysis of very large sets of data are discussed in Chapter 10, “Data at Scale.”

More information about software tools designed to support the analysis of qualitative data can be found through the CAQDAS Networking Project, based at the University of Surrey.

<https://www.surrey.ac.uk/computer-assisted-qualitative-data-analysis>

9.7 Interpreting and Presenting the Findings

Previous sections in this chapter have illustrated a range of different ways to present findings—as tables of numbers and text, through various graphical devices and diagrams, as a set of themes or categories, and so on. Choosing an appropriate way to present the findings of a study is as important as choosing the right analytical approach. This choice will depend on the data gathering and analysis techniques used as well as the audience and the original goals of the study. In some situations, the details of data collection and analysis will be needed, for example, when working with others to make sense of a large collection of data, or when trying to convince an audience about a controversial conclusion. This detail may include snippets of data such as photographs of the context of use or videos of participants using the product. In other situations, only the salient trends, headlines, and overall implications are needed, so the style of presentation can be leaner. Where possible, a set of different complementary representations will be chosen to communicate the findings since any one representation will emphasize some aspects and de-emphasize others.

This section focuses on three kinds of presentation styles that haven't yet been emphasized to this point: using structured notations, stories, and summarizing.

ACTIVITY 9.4

Consumer organizations and technology companies regularly conduct investigations about technology use. One such report is from DScout investigating smartphone use, available from <https://blog.dscout.com/mobile-touches>. Note that the report is a downloadable PDF from this web page. Using this report or another report that you find online, look to see how the results are reported.

1. What kinds of presentation are used?
2. What is left out of the report?
3. What is the effect of presenting findings this way?

Comment

The study from DScout found that “users tapped, swiped, and clicked a whopping 2,617 times each day on average.” Note that the web page and the report are presented differently. The web page includes more text and an edited video of participants' responses to the findings, while the PDF report is presented in a style similar to a PowerPoint presentation. Focusing on the PDF report:

1. Graphs and pie charts are used to present demographics and app use, along with lists, tables, and typographic styles to emphasize certain findings. The report uses two other representations (see Figure 9.14): a bubble diagram to show the relative use of different apps; and a timeline to represent a “heavy user” and a “average user” and how their touches are spread across the day. These two representations illustrate that developing new (or modifying old) representations to communicate the right findings you want to highlight is acceptable.

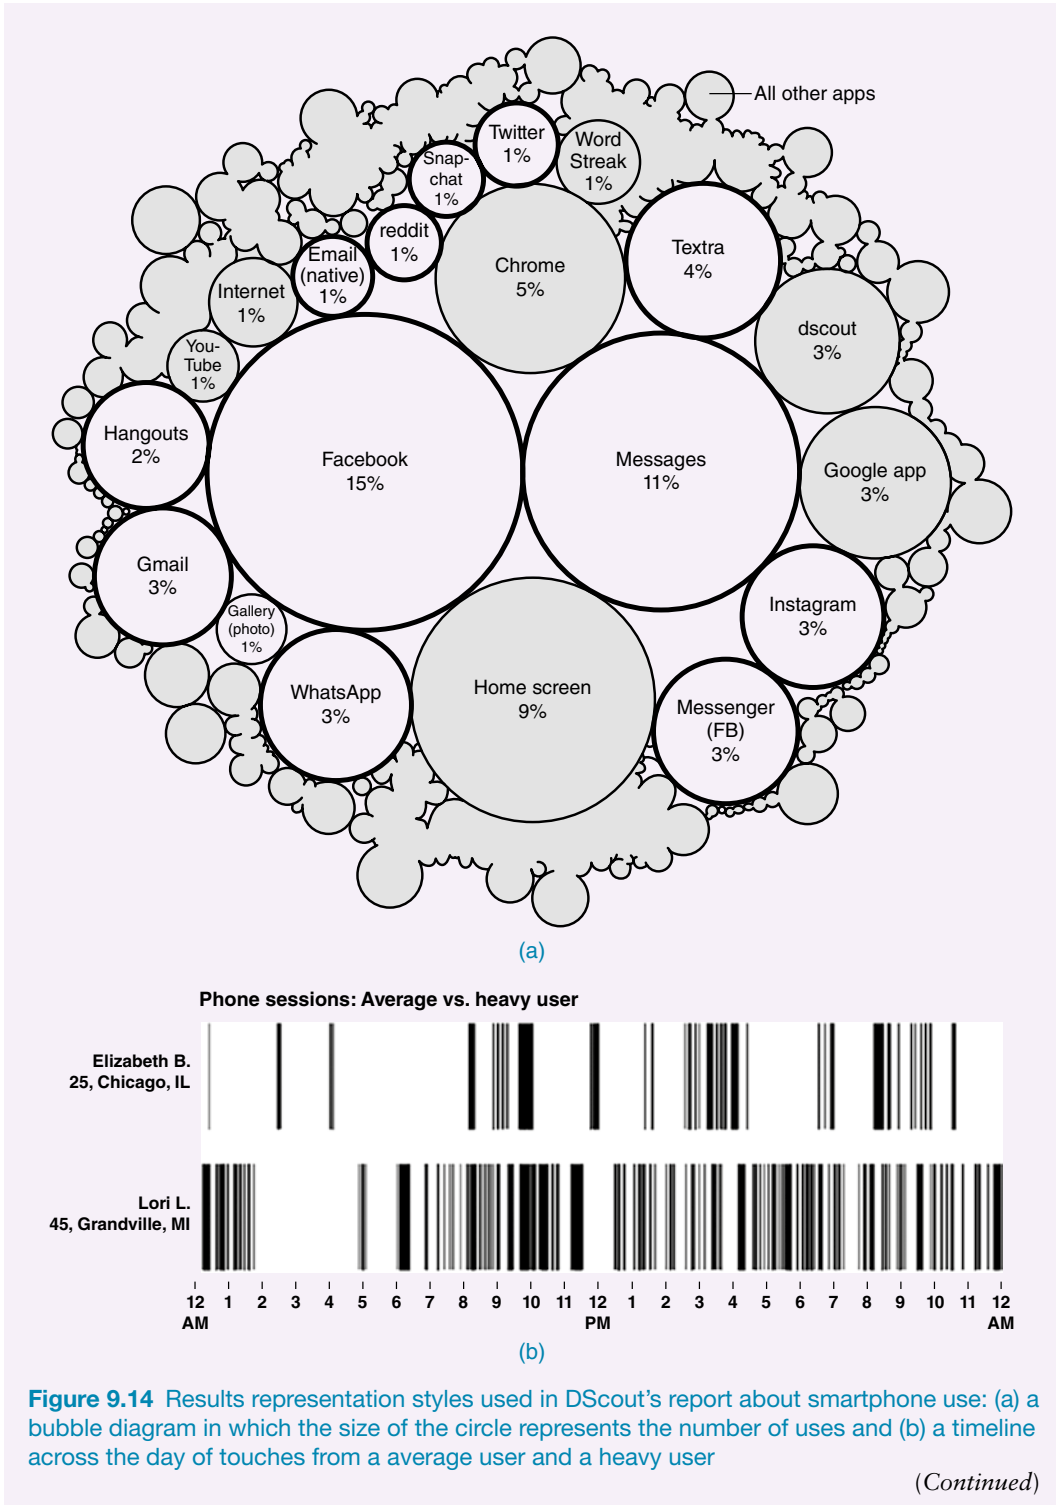


Figure 9.14 Results representation styles used in DScout's report about smartphone use: (a) a bubble diagram in which the size of the circle represents the number of uses and (b) a timeline across the day of touches from a average user and a heavy user

(Continued)

2. There is very little text or description in the PDF report and minimal details about the way in which the data was collected or analyzed. Further details are included on the web page, but not in the report.
3. The way in which the findings are presented in the report has quite an impact. The bold and clear graphical images together with minimal but highlighted text mean that the messages are communicated in a straightforward fashion. ■

9.7.1 Structured Notations

A number of *structured notations* have been developed to analyze, capture, and present information for interaction design. These notations follow a clear syntax and semantics, which have been developed to capture particular viewpoints. Some are relatively straightforward, such as the work models promoted in contextual design (Beyer and Holtzblatt, 1998) that use simple conventions for representing flows, breakdowns, individual roles, and so on. Others, such as the modeling language Unified Modeling Language (UML), have stricter and more precise syntax to be followed and are often used to represent requirements (see Chapter 11, “Discovering Requirements”); the activity diagrams, for example, are very expressive when detailed interactions need to be captured.

Advantages of using a structured notation are that the meaning of different symbols is well-defined, and so it provides clear guidance on what to look for in the data and what to highlight and that it enforces precision in expression. Disadvantages include that by highlighting specific elements, it inevitably de-emphasizes or ignores other aspects, and the precision expressed by the notation may be lost on an audience if they don't know the notation well. Producing diagrams or expressions in these notations may require further analysis of the findings in order to identify the specific characteristics and properties that the notation highlights. To overcome these disadvantages, structured notations are usually used in combination with stories or other easily accessible formats.

9.7.2 Using Stories

Storytelling is an easy and intuitive approach for people to communicate ideas and experiences. It is not surprising then that stories (also called *narratives*) are used extensively in interaction design, both to communicate findings of investigative studies and as the basis for further development, such as product design or system enhancements.

Storytelling may be employed in three different ways. First, participants (such as interviewees, questionnaire respondents, and those you have observed) may have told stories of their own during data gathering. These stories can be extracted, can be compared, and may be used to communicate findings to others, for example, to illustrate points.

Second, stories (or narratives) based on observation, such as ethnographic field studies, may be employed for further data gathering. For example, Valeria Righi et al. (2017) used stories as the basis of co-design workshops in their study to explore the design and use of technologies to support older people. The scenarios were developed on the basis of ethnographic studies and previous co-design activities and were presented through storytelling to facilitate understanding. Note that, in this case, the audience was a group of participants in the ongoing study.

Including specific stories gives authenticity to the findings, and it can add to its credibility provided the conclusions are not overstated. Making a multimedia presentation of the story by adding video or audio excerpts and photographs will illustrate the story further. This kind of approach can also be effective if presenting data from an evaluation study that involves observation, as it is hard to contest well-chosen video excerpts of users interacting with technology or extracts from interview transcripts.

Third, stories may be constructed from smaller snippets or repeated episodes that are found in the data. In this case, stories provide a way of rationalizing and collating data to form a representative account of a product's use or a certain type of event.

Any stories collected through data gathering may be used as the basis for constructing scenarios that can then be used for requirements and design activities. See Chapters 11 and 12 for more information on scenarios.

9.7.3 Summarizing the Findings

Presentation styles will usually be used in combination to produce a summary of the findings; for instance, a story may be expanded with graphical representations of activity or demographics, and data excerpts from transcripts or videos may be used to illustrate particular points. Tables of numerical data may be represented as graphs, diagrams, or rigorous notations, together with workflows or quotations.

Careful interpretation and presentation of the study results is just as important as choosing the right analysis technique so that findings are not over-emphasized, and evidence is not misrepresented. Over-generalizing results without good evidence is a common pitfall, especially with qualitative analyses' for example, think carefully before using words such as *most*, *all*, *majority*, and *none*, and be sure that the justifications reflect the data. As discussed in Box 9.1, even statistical results can be interpreted in a misleading way. For example, if 8 out of 10 users preferred design A over design B, this does not mean that design A is 80 percent more attractive than design B. If you found 800 out of 1,000 users preferred design A, then you have more evidence to suggest that design A is better, but there are still other factors to consider.

ACTIVITY 9.5

Consider each of the following findings and the associated summary statement about it. For each one, comment on whether the finding supports the statement.

- Finding:** Two out of four people who filled in the questionnaire checked the box that said they prefer not to use the ring-back facility on their smartphone.
Statement: Half of the users don't use the ring-back facility.
- Finding:** One day, Joan who works in the design department was observed walking for 10 minutes to collect printouts from the high-quality color printer.
Statement: Significant time is wasted by designers who have to walk a long distance to collect printouts.
- Finding:** A data log of 1,000 hours of interaction with a website recorded during January, February, and March records eight hours spent looking at the help files.
Statement: The website's help files were used less than 1 percent of the time during the first quarter of the year.

(Continued)

Comment

1. The questionnaire didn't ask if they use the ring-back, just whether they preferred to use the ring-back facility. In addition, two users out of four is a very small number of participants, and it would be better to state the actual numbers.
2. Observing one designer on one day having to walk to get printouts does not mean that this is a general problem. There may be other reasons why this happened on this day, and other information is needed to make a clear statement.
3. This statement is justified as the log was recorded for a significant period of time and using percentages to represent this finding is appropriate as the numbers are large. ■

In-Depth Activity

The goal of this in-depth activity is to practice data analysis and presentation. Assume that you are assigned to analyze and present the findings of your data gathering in-depth activity from Chapter 8 to a group of peers, for instance, via a seminar.

1. Review the data that you gathered and identify any qualitative data and any quantitative data in the data set.
2. Is there any qualitative data that could sensibly and helpfully be translated into quantitative measures? If so, do the translation and add this data to your quantitative set.
3. Consider your quantitative data.
 - (a) Decide how best to enter it into spreadsheet software, for example, how to handle answers to close-ended questions. Then enter the data and generate some graphical representations. As the data set is likely to be small, think carefully about what, if any, graphical representations will provide meaningful summaries of the findings.
 - (b) Is there any data for which simple measures, such as percentages or averages, will be helpful? If so, calculate the three different types of average.
4. Consider your qualitative data.
 - (a) Based on your refinement of the study question "improving the product," identify some themes in the qualitative data, for example, what features of the product cause people difficulties? Did any of the participants suggest alternative designs or solutions? Refine your themes and collate extracts of data that support the theme.
 - (b) Identify any critical incidents in the data. These may arise from interviews, questionnaire responses, or observation. Describe these incidents carefully and choose one or two to analyze in more depth, focusing on the context in which they occurred.
5. Collate your findings as a presentation and deliver them to a group of peers.
6. Review the presentation and any questions from the audience. Consider how to improve the analysis and presentation.

Summary

This chapter described in detail the difference between qualitative and quantitative data and between qualitative and quantitative analysis.

Quantitative and qualitative data can be analyzed for patterns and trends using simple techniques and graphical representations. Qualitative data may be analyzed inductively or deductively using a variety of approaches. Thematic analysis (an example of inductive analysis) and data categorization (an example of deductive analysis) are common approaches. Analytical frameworks include conversation analysis, discourse analysis, content analysis, interaction analysis, grounded theory, and systems-based approaches.

It was noted that presenting the results is just as important as analyzing the data, hence it is important to make sure that any summary or claim arising from the analysis is carefully contextualized, and that it can be justified by the data.

Key Points

- The kind of data analysis that can be done depends on the data gathering techniques used.
- Qualitative and quantitative data may be collected from any of the main data gathering techniques: interviews, questionnaires, and observation.
- Quantitative data analysis for interaction design usually involves calculating percentages and averages.
- There are three different kinds of average: mean, mode, and median.
- Graphical representations of quantitative data help in identifying patterns, outliers, and the overall view of the data.
- Analysis of qualitative data analysis may be inductive, in which themes or categories are extracted from the data, or deductive, in which pre-existing concepts are used to interrogate the data.
- In practice, analysis often proceeds in iterative cycles combining inductive identification of themes and deductive application of categories and new themes.
- Which analysis approach is used is tightly coupled to the data that is collected and depends on the goals of the study.
- Several analytical frameworks exist that focus on different levels of granularity with different purposes.

Further Reading

BLANDFORD, A., FURNISS, D. and MAKRI, S. (2017) *Qualitative HCI Research: Going Behind the Scenes*. Morgan Claypool Publishers. This book in the form of a lecture discusses the practical details behind qualitative analysis in HCI. Using the analogy of making a documentary film, the authors point out that, as with movies, qualitative analysis is often presented as a finished product while the work “behind the scenes” is rarely discussed.

BRAUN, V. and CLARKE, V. (2006) Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2), pp. 77–101. This paper focuses on thematic analysis, how it relates to other qualitative analysis approaches, and how to conduct it in a rigorous fashion. It also discusses advantages and disadvantages of the approach.

CHARMAZ, K. (2014) *Constructing Grounded Theory* (2nd ed.). Sage Publications. This popular book provides a useful account of how to do grounded theory.

CORBIN, J. M. and STRAUSS, A. (2014) *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage. This presents a readable and practical account of applying the grounded theory approach. It is not tailored specifically to interaction design, and therefore it requires some interpretation. It is a good discussion of the basic approach.

HUFF, D. (1991) *How to Lie with Statistics*. Penguin. This wonderful little book illustrates the many ways in which numbers can be misrepresented. Unlike some (many) books on statistics, the text is easy to read and amusing.

ROGERS, Y. (2012) *HCI Theory: Classical, Modern, and Contemporary*. Morgan and Claypool Publishers. This short book, in the form of a lecture, charts the theoretical developments in HCI, both past and present, reflecting on how they have shaped the field. It explains how theory has been conceptualized, the different uses it has in HCI, and which theory has made the most impact.

Chapter 10

DATA AT SCALE

10.1 Introduction

10.2 Approaches for Collecting and Analyzing Data

10.3 Visualizing and Exploring Data

10.4 Ethical Design Concerns

Objectives

The main goals of the chapter are to accomplish the following:

- Provide an overview of some of the potential impacts of data at scale on society.
- Introduce key methods for collecting data at scale.
- Discuss how data at scale becomes meaningful.
- Review key methods for visualizing and exploring data at scale.
- Introduce design principles for making data at scale ethical.

10.1 Introduction

How do you start your day? How much data do you encounter when first looking at your smartphone, switching on your laptop, or turning on another device? How much do you knowingly create and how much do you create unknowingly? Upon waking up, many people routinely will ask their personal assistant, something like, “Alexa, what is the weather today?” or “Alexa, what is the news?” or “Alexa, is the S-Bahn train to Schönefeld Airport running on time?” Or, they will ask Siri, “What is my first meeting?” or “Where is the meeting?”

Having oriented themselves for the day, people will walk a few blocks to the subway entrance, dip their Metro Card in the turnstile to pay the fare, exit the station at their stop, grab their favorite morning beverage at a nearby cafe, and proceed to their office where they check in with the employee card at a security gate and take the elevator to their floor.

These are just a few of the things that many of us do to start our workdays. Each activity involves creating, searching, and storing data in some way or another. We may know that this is happening, we may suspect that it is happening, or we may be totally unaware of the data that we are generating and with which we are interacting.

There is also increasing concern about exactly what data is collected about us through personal assistants such as Amazon Echo, Google Home, Cortana, and Siri. We also know that many large cities, such as New York and London, have an enormous number of surveillance cameras (CCTV) spread around, especially in busy places such as subway stations and shopping malls. The video footage from these sources is kept for two weeks or more. Similarly, we experience being checked into an office, so we know that our movements are being tracked by security personnel. Our activities are also being tracked more surreptitiously through the technology that we use such as smartphones and credit cards.

What happens to all the data collected about us? How does it improve the services provided by society? Does it make traveling more efficient? Does it reduce traffic congestion? Does it make the streets safer? Moreover, how much of the data collected from our smartcards, smartphone Wi-Fi signals, and CCTV footage can be tracked back to us and pieced together to reveal a bigger picture of who we are and where we go? What might that data reveal about us?

Data at scale, or as it is often called *big data*, describes all kinds of data including databases of numbers, images of people, things and places, footage of conversations recorded, videos, texts, and environmentally sensed data (such as air quality). It is also being collected at an exponential rate; for example, 400 new YouTube videos are uploaded every minute, while millions of messages circulate through social media. Furthermore, sensors collect billions of bytes of scientific data.

Data at scale has huge potential for grounding and elucidating problems, and it can be collected, used, and communicated in a wide variety of ways. For example, it is increasingly being used for improving a whole range of applications in healthcare, science, education, city planning, finance, world economics, and other areas. It can also provide new insights into human behavior by analyzing data collected from people, such as their facial expressions, movements, gait, and tone of voice. These insights can be enhanced further by using machine learning and machine vision algorithms to make inferences. This includes people's emotions, their intent, and well-being, which can then be used to inform technology interventions aimed at changing or improving people's health and well-being. However, beyond societal benefits, data can also be used in potentially harmful ways.

As mentioned in Chapter 8, "Data Gathering," and Chapter 9, "Data Analysis," data can be either qualitative or quantitative. Some of the methods and tools used to collect, analyze, and communicate data can be carried out manually or using quite simple tools. What makes this chapter on data at scale different is that it considers how huge volumes of data can be analyzed, visualized, and used to inform new interventions. While having access to large volumes of data enables analysts, designers, and researchers to address large, important issues such as climate change and world economic issues, assuming that there are tools to do this, they also raise a number of user concerns. These include whether someone's privacy is being violated by the data being collected about them and whether the data corpora being used to make decisions about people, such as the provision of insurance and loans, are fair and transparent.

Furthermore, the combination of vast amounts of data from many sources and the availability of increasingly powerful data analytic tools to analyze that data is now making it possible to discover new information that is not available from any single data source. This is enabling new kinds of research to be conducted for understanding human behavior and environmental problems.

10.2 Approaches to Collecting and Analyzing Data

Collecting data has never been easier. What is challenging is knowing how best to analyze, collate, and act upon the data in ways that are socially acceptable, beneficial to society, and ethically sound. Are there certain rules or policies in place on what to reveal about people or when certain patterns, anomalies, or thresholds are reached in a data stream? For example, if people-tracking technology is used at an airport, how is that revealed to those at the airport? Is it enough only to show data that can help manage people flows and bottlenecks? For example in an airport terminal showing a public display in which one section of the terminal is detected to be much busier than another (Figure 10.1), do travelers ever stop and wonder how this data is being collected? What else is being collected about them? Do they care?



Figure 10.1 Heathrow Airport Terminal 5 Public Display in top-right corner of image showing the relative level of activity using an infographic of North vs. South Security

Source: Marc Zakian / Alamy Stock Photo

Another technique for analyzing what people are doing on websites and social media is to examine the trail of activity that they leave behind. You can see this by looking at your own Twitter feed or by looking at someone else's whom you are following, for example, a friend's, a political leader, or a celebrity. You can also examine discussions about a particular topic such as climate change, reactions to comments made by comedians like John Oliver or Stephen Colbert, or a topic that is trending on a particular day. If there are just a few posts, then it is easy to see what is going on, but often the most interesting posts are those that generate lots of comments. When examining thousands or tens of thousands of posts, analysts use automated techniques to do this (Bostock et al., 2011; Hansen et al., 2019).

10.2.1 Scraping and “Second Source” Data

One way to extract data is by “scraping” it from the web (assuming that this is allowed by the application). Once the data is scraped, it can be entered into a spreadsheet for study and analyzed using data science tools. The focus from an interaction design perspective is how one can interact with that data and the way it is displayed rather than the actual scraping process *per se*, so that it can be analyzed and sense can be made of it.

In addition, the openly available big data that Google and other companies now provide for researchers to mine offers a “second source” methodology, meaning search terms, Facebook posts, Instagram comments, and so on. Analysis of this data can indirectly reveal new insights about the users’ concerns, desires, behaviors, and habits. For example, the Google Trends tool can be used for exploring and examining the motivation behind what people ask when they type something into Google Search. Seth Stephens-Davidowitz (2018) has used it extensively to reveal what people are interested in finding out. From his analysis of Google Search data, he discovered that people type into the search box all sorts of intimate questions about their health among other topics. Moreover, he found that his analysis of search data revealed things that people would not freely admit to when asked using other research methods, such as surveys and interviews. He also makes an important assertion: to obtain new insights from big data, it requires asking the right questions of the data. Furthermore, it is not how much data can be collected or mined but what is done with the new data that has been made available. Simply mining it because there is a tool available may yield surprising results, but well-honed questions that guide and are used to interpret the data that is found will be more valuable (see Chapter 8, “Data Gathering”).

How do researchers know what are the right questions to ask of this data? This is particularly pertinent for HCI researchers to understand, especially in terms of how users will relate, trust, and confide in the next generation of technologies, including domestic robots, bots, and virtual agents.

ACTIVITY 10.1

What insights do Google Trends searches tell us about ourselves?

Go to Google Trends (<https://trends.google.com>). Then try typing into the search box statements such as “I hate my boss,” “I feel sad,” or “I eat too much.” See how many people have typed this into Google over the last month, year, and for different countries. Then type in the opposite statements: “I love my boss,” “I am happy,” or “I never eat enough.” How do the results compare? Which is asked more often? Then type in your name and see what Google returns.

Comment

It is surprising how many people confide such personal statements in Google. Some people will tell it anything. Google Trends provides a way of comparing the search data across time, country, and other topics. When you type in your name (unless you have the same name as a famous person), it often comes back with “hmm, your search doesn’t have enough data to show here.” ■

10.2.2 Collecting Personal Data

Personal data collection started becoming popular through the quantified-self (QS) movement in 2008 where monthly “show and tell” meetings were organized to enable people to come together to share and discuss their various self-tracking projects. Nowadays, many apps and wearable devices exist that people can buy off the shelf, which can collect all sorts of personal data and visualize it. These results can be matched against targets reached, and recommendations, hints, or tips can also be provided about how to act upon it. Many apps now come prebundled on a smartphone or smartwatch, including those that quantify health, screen time, and sleep. Some also allow multiple activities to be tracked, aggregated, and correlated. The most common types of apps are for physical and behavioral tracking, including mood changes, sleep patterns, sedentary behavior, time management, energy levels, mental health, exercise taken, weight, and diet. A common motivation for deciding to embark on tracking some aspect of one’s personal data over time is to see how well they are doing compared to a set threshold or level (that is, a set target, a comparison with the week before, and so on). The aggregate data may raise awareness and be revealing to the extent that someone feels compelled to act upon it (for example, changing their sleeping habits, eating more healthily, or going to the gym more regularly).

Self-tracking is also increasingly being used by people who have a condition or disease as a form of self-care, such as monitoring blood glucose levels for those who have diabetes (O’Kaine et al., 2015) and the occurrence of migraine triggers (Park and Chen, 2015). This kind of self-care monitoring has been found to help people engage in reflection when looking at their data and then learning to associate specific indicators with patterns of behavior. Making these connections can increase self-awareness and provide them with early warning signs. It can also lead them to avoid certain events or adjust their behavior accordingly. Many people are also happy to share their tracked data with others in their social networks, which has been found to enhance their social networking and motivation (Gui et al., 2017).

Quantified-self projects generate lots of data. New kinds of health data can now be collected by mobile health monitors, such as heart rate, generating masses of data per person each month, which was simply heretofore unavailable. This raises questions as to how much data should be kept and for how long? Also, how can this data be used to best effect? Should it signal to the user when their heart rate deviates from normal levels? Given that masses of data are being collected from many individuals using the same devices, it is possible to collate all of the data. Would it be useful for health clinicians and the individuals alike to have access to all of this data in order to see trends and comparisons? How can this be made to be both informative and reassuring? Translating someone’s heart rate data that is sampled many times per second along with their electroencephalogram (EEG) streamed brainwave data into an early warning sign with an appropriate form of intervention is challenging (Swan, 2013). It can easily lead to increased anxiety. Much thought needs to go into providing information into an interface that will not cause unnecessary panic. New tools should also provide flexibility in how the user might want to customize or annotate their data to meet their specific needs (Ayobi et al., 2018).

10.2.3 Crowdsourcing Data

Increasingly, people crowdsource information or work together using online technologies to collect and share data. The idea of a crowd working together has been taken one step further in *Crowd Research*, where many researchers from all over the world come together to work

on large problems, such as climate science (Vaish et al., 2018). The goal of this approach is potentially to enable hundreds of people to contribute, through collecting data, ideating, and critiquing each other's designs and research projects. Conducting research on a massive scale enables potentially hundreds or thousands of people to work on a single project, which can help address large problems, such as migration or climate change at scale.

There are also many citizen science and citizen engagement projects (see Chapter 5, "Social Interaction") that crowdsource data at scale and in doing so amass billions of different types of data (photos, sensor readings, comments, and discussion), which are collected by millions of people across the world. Most of this data is stored in the cloud as well as on local machines. Examples of large citizen science projects include iSpotNature, eBird, iNaturalist,




Figure 10.2 Abundance map for the common raven. The darkest area indicates where ravens are most abundant.

Source: <https://ebird.org/science/status-and-trends/comrav>. This link will allow you to see how abundance changes during each week of the year (purple indicates high abundance and yellow indicates low abundance).

and the Zooniverse. There are also thousands of much smaller projects that together generate huge amounts of data.

eBird.org, for example, collects data about bird sightings that is contributed by naturalists—amateurs ranging from beginning birders to highly experienced expert birders and professional scientists. The site was launched in 2002 as a collaboration effort between Cornell University’s Lab of Ornithology and the National Audubon Society. The data includes bird species data, the abundance of each species, geolocation data indicating where observations are made, profiles of the people who contribute, comments, and discussion. There are also smartphone apps and a website with links to many resources, including identification guides, data analysis tools, maps and visualizations, reports, and scientific articles. As of June 2018, there were more than 500 million bird observations recorded in a global database. eBird feeds data into aggregator sites such as the Global Biodiversity Information Facility (GBIF) so that it is available for scientists. It also provides several maps, many of which are interactive (see Figure 10.2) and other graphic representations of its data that are available for anyone to access.

Crowd projects raise a number of issues as to who owns and manages the data. This is especially pertinent when the data collected can be mined to unearth details about the people who contribute the data as well as about the endangered species, for example. For researchers and UX designers, there are interesting questions about how to balance making data available for education and research while protecting the privacy of those contributing the data and the location of endangered species in this example. Box 10.1 discusses how one citizen science project, iNaturalist.org, tries to manage this balance.

BOX 10.1

Citizen Science and UX Design for Privacy

How privacy is interpreted by citizen scientists and their desire and need for privacy regulations differs across projects (Bowser et al., 2017). Being able to share citizen science data has many advantages as well as some privacy concerns. For example, participants can see what others are observing in their area. Bird enthusiasts also often like to share first sightings, for instance, when the first swallows appear in spring or the first snow geese arrive in winter. They may also want to check identifications with each other. The downside of this community interaction is that personal profile and location data can be used to identify particular contributors and their patterns of behavior. The latter can be especially problematic, as many participants visit the same places regularly. It is, therefore, important to ask how important is privacy in citizen science compared with the benefits of community engagement? And how might UX design help protect both participants and rare species, while supporting open engagement?

Many citizen science projects and societies post privacy policies like the one shown in the following link. Other strategies involve making images and locations fuzzy so that they are not exact. This is also a good strategy for keeping the location of rare species’ observations confidential. For example, iNaturalist.org has a geoprivacy setting that can be set to “open,” “obscure,” or “private.” Obscured observations are used to hide the exact location of endangered species, as shown in Figure 10.3.

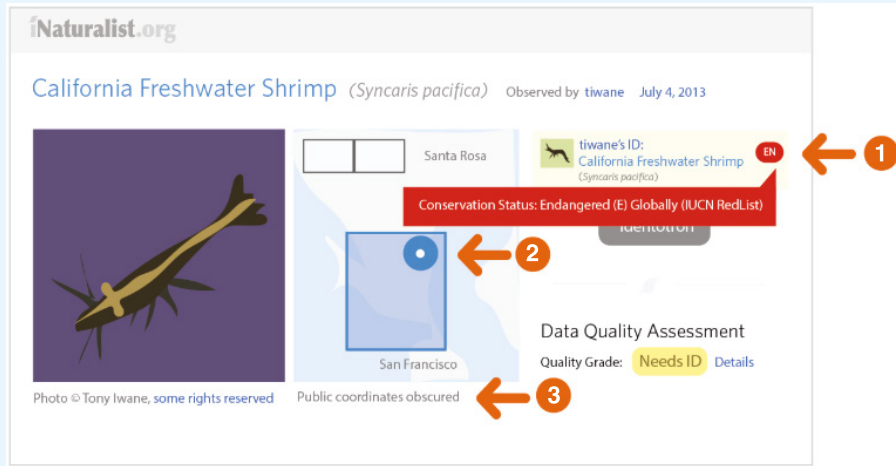


Figure 10.3 iNaturalist.org geoprivacy obscures the location of an observation.

In the above example: 1. EN indicates that the organism is endangered, so its location needs to be obscured, 2. indicates that obscuring is done by randomly placing the marker for the location within the broader area, and 3. allows the contributor to verify that this observation has been observed within iNaturalist.

Source: <https://www.inaturalist.org> ■

The European Citizen Science Organization's Data and Privacy Policy can be seen at <https://ecsa.citizen-science.net/about-us/privacy-policy>.

BOX 10.2

A Human-Data Design Approach to Sensing Data

There are a number of off-the-shelf sensor toolkits available now that can be placed in someone's home or local community to measure air quality or other aspects of the environment. One of the earliest open platforms developed was Smart Citizen (Diez and Posada, 2013). A compact device was built with a number of embedded sensors in it that could measure nitrogen dioxide (NO₂), carbon monoxide (CO), sunlight, noise pollution, temperature, and humidity levels. The data being collected from the platform was connected to a live website that could be accessed by anyone. The various data streams were presented via a dashboard using canonical types of visualisations, such as time-series graphs (see Figure 10.4a). Data streams from other Smart Citizen devices, set up throughout the world, could also be viewed via the dashboard (<https://smartcitizen.me/>) making it easy to compare data collected from

different locations. While the mass of environmental data accumulating was fascinating to data scientists and researchers, this was not the case for many of the householders who had set up a smart citizen device in their home. They found the visualizations presented via the dashboard to be difficult to understand and were unable to connect the data being sensed with what was happening in their homes (Balestrini et al., 2015). As a result, they did not find it useful and quickly stopped looking at it.

The Physikit project took this user problem as its starting point (Houben et al., 2016). A human-data design approach was adopted; the goal is to transform sensed data being collected into something that is meaningful to the general public. The goal is to provide a way of enabling users to engage with their sensed data, by giving it a physical ambient presence in the location where it is being collected. A set of colorful physical cubes were designed that could light up, move parts or vibrate, depending on how they were configured (see Figure 10.4c). Householders could easily configure a set of rules to decide which cubes to connect with which data streams and what each cube should do depending on levels or thresholds being sensed. This was intended to let them select aspects of their home that they were interested in knowing more about.

For example, one of the PhysiCubes had a rotating disk on the top that could be set to move clockwise or counterclockwise and at different speeds. One household decided to use it to measure the level of humidity in their kitchen throughout the day. They placed a basil plant on top of the cube (see Figure 10.4b) as a way of visibly showing the level of humidity in the kitchen. The rule they set up for the cube was for it to rotate only if the humidity level detected was below 60 percent. At the end of each day, they could tell how humid it had been in the room by the extent to which the plant remained upright. If it was leaning toward the window, this suggested to them that the humidity level in the kitchen had been high throughout the day because the disk had not moved the plant around for the leaves to get an even amount of light. The household, had in effect, created a naturally growing physical visualization that held historical data.



(a)



10.2.4 Sentiment Analysis

Sentiment analysis is a technique that is used to infer the effect of what a group of people or a crowd is feeling or saying. The phrases that people use when offering their opinions or views are scored as being negative, positive, or neutral. The scales used vary along a continuum from negative to positive, for example, -10 to $+10$ (where -10 is the most negative, 0 is neutral, and $+10$ is the most positive). Some sentiment systems provide more qualitative measures by identifying if the positive or negative sentiment is associated with a specific feeling, for example anger, sadness, or fear (negative feelings) or happiness, joy, or enthusiasm (positive feelings). The scores are extracted from people's tweets and texts, online reviews, and social media contributions. Their facial expressions (see Chapter 6, "Emotional Interaction") when looking at ads, movies, and other digital content and customer's voices can also be analyzed and classified using the same scales. Algorithms are then applied to the labeled data in order to identify and classify them in terms of the level of effect that has been

expressed. There are a number of online tools that can be used to do this, such as DisplayR and CrowdFlower. MonkeyLearn provides a detailed tutorial on sentiment analysis (<https://monkeylearn.com/sentiment-analysis/>).

Sentiment analysis is commonly used by marketing and advertising companies to decide on what types of ads to design and place. In addition, it is increasingly being used in research to study social science phenomena. For example, Veronikha Effendy (2018) used sentiment analysis to study people's opinions about the use of public transportation from their tweets. In particular, she was interested in determining what were the positive and negative opinions toward it, which could then be used as evidence for making a case on how to improve public transportation to increase its use in Indonesia, where there are huge traffic congestion problems.

However, sentiment analysis as a technique is not an exact science and should be viewed more as a heuristic than as an objective evaluation method. Giving a word a score from -10 to $+10$ is quite a crude way to measure. To assess how good sentiment analysis is as a method, Nicole Watson and Henry Naish (2018) compared human judgment with computer-based sentiment analysis for evaluating positive articles about the U.S. economy. They found that the computer was more often wrong than right compared with the human participants. Their analysis indicates that humans express their optimism about a topic in much richer ways. Moreover, it also shows that by focusing on emotive words in phrases, sentiment analysis misses the nuances of expression that humans understand intuitively. For example, how would sentiment analysis score the phrase written by a teen in a text to their friends that said, "I am weak"? It would probably give it a negative score. In fact, the phrase is teen slang for "That is funny," which is completely the opposite.

10.2.5 Social Network Analysis

Social network analysis (SNA) is a method based on social network theory (Wellman and Berkovitz, 1988; Hampton and Wellman, 2003) for analyzing and evaluating the strength of social ties within a network. While understanding social ties has been a strong interest of sociologists for many years, (for example, Hampton and Wellman, 2003; Putnam, 2000), as social media became increasingly successful, it also became a key interest for computer and information scientists (for example, Wasserman and Faust, 1994; Hansen et al., 2019). They want to understand the relationships that form among people and groups within and across different social media platforms, and with offline social networks too. Online, trillions of messages, tweets, pictures, and videos are posted and responded to every second of every day via Weibo, Tencent, Baidu, Facebook, Twitter, Instagram, and YouTube. Some examples include families posting pictures of their kids' birthday parties and family outings, discussion of hot political issues, and friends and colleagues chatting and keeping in contact with each other's travel experiences, hobbies, life's challenges, and successes.

Social network analysis enables these relationships to be seen more clearly. It helps to reveal who is most active in a group, who belongs to which groups, and how the groups do or do not interact and relate to each other. Analyses can also show which topics are hot and throw light on when, how, and why some topics go viral. Managers, marketers, and politicians are especially interested in how these activities can influence them, their companies, and their constituents. Many other people like to try to make their posts or YouTube videos go viral, as discussed in Chapter 5.

So, how does social network analysis work? It is a big topic, but broadly, as the name suggests, a *network* is a collection of things and their relationships to each other. A *social network* is a network of people and groups with relationships to each other. Human beings, like other primates, have formed networks for as long as our species has existed. Many other species, such as elephants, wolves, and meerkats, to name just a few, also rely on social networks for safety, for collaboratively rearing their young, and when foraging or hunting for food.

Two main entities make up a social network. *Nodes*, which are also sometimes called entities or vertices, represent people and topics. The connections between the nodes are called *edges*, which are also known as links or ties. The edges show the connections among nodes, for example, the members of a family. They can show the direction of relationships; for instance, parents may have a line with an arrow-head that points to their children, indicating the direction of the relationship between the two nodes. Similarly, an arrow in the opposite direction indicates that children have parents. These are known as *directional edges*. Edges can also indicate relationships in both directions by having arrows at each end. Edges that do not have an arrowhead are nondirectional; that is, the direction of the relationship between two nodes is not shown.

Drawing on algorithms and based in statistics, social network analysis offers a range of metrics for describing the properties of networks. One of the most important sets of metrics for visualizing networks in big data is measures of centrality. Several different measures of centrality exist based on different statistical formulae. These and other metrics are used to create visualizations like the ones shown in Figure 10.5 that show overlapping clusters. The clusters indicate voting patterns by members of the U.S. Senate in 1989, 1997, and 2013. Red represents Republicans, and blue represents Democrats. The graphs indicate how, over almost 25 years, the voting behavior of members of the two parties has become increasingly siloed, with fewer and fewer members voting with members from the other party. The nodes representing members on the far right and far left are only connected to the nodes of senators in their own party, indicating that they don't vote with members from the other party. The nodes in the middle indicate that those members sometimes vote with members of the other party. From being united on some issues back in 1989, bipartisan voting behavior has decreased over the years as indicated by the social network graphs. By 2013, few members of the two parties voted with members from the other party.

Some other topics that have been studied using social network analysis include communication during the flood in Louisiana, where Jooho Kim and Makarand Hastak (2018) examined the role of social media in flood victims' communication, both with each other and with emergency services. They found that Facebook was used particularly effectively to disseminate information. Other studies include one by Dinah Handel and her colleagues on teachers' tweets on Twitter (Handel et al, 2016), and Diane Harris Cline has used social network analysis for a number of studies to examine the relationships between historical characters (Cline, 2012). In addition, there are many other examples related to a diversity of topics, including business communication, and even the relationships and activities of characters in Shakespeare's plays (Hansen et al., 2019). Revealing as many of these social network graphs are, using the tools effectively to separate and display clusters, outliers, and other network features takes practice. However, increasing attention to the UX design and support provided by such tools enables beginners to do straightforward analyses. Two of the most well-known social network analysis tools are NodeXL (Hansen et al., 2019), which runs on Windows-based machines, and Gephi, which runs on both Windows and macOS. Many YouTube videos are available that describe how to use these tools.

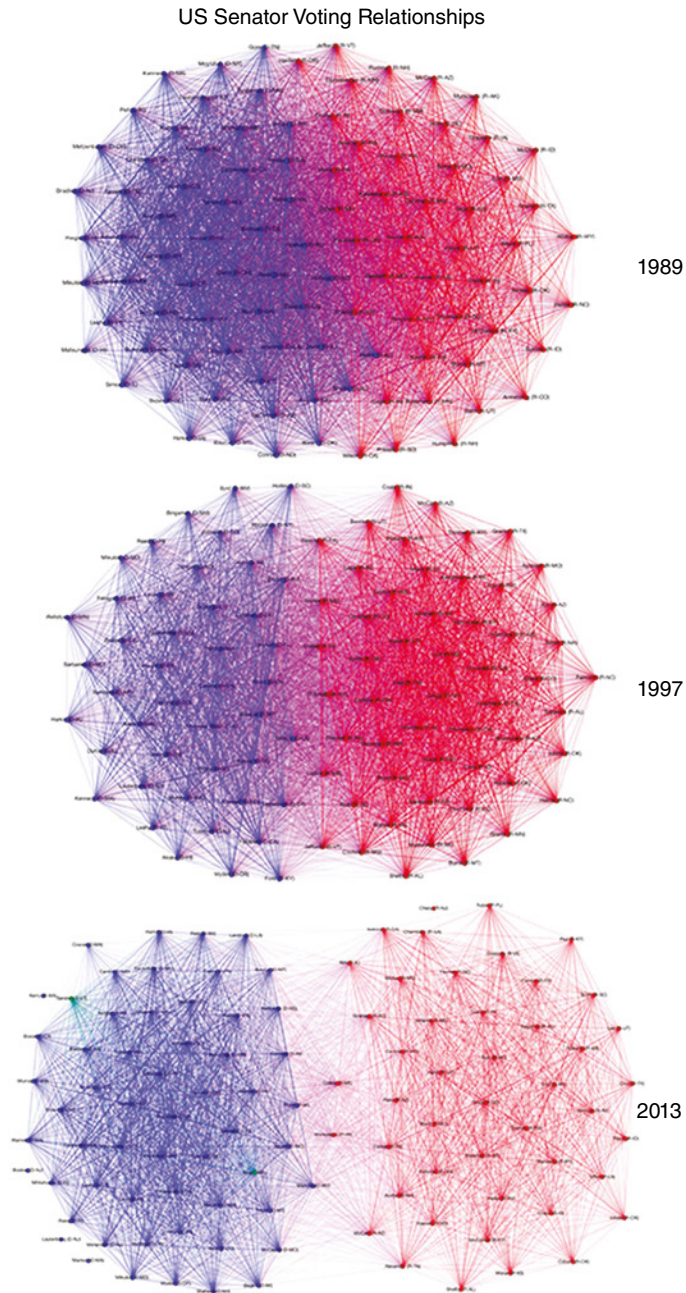


Figure 10.5 Voting behavior of U.S. Senators in 1989, 1997, and 2013. Red represents Republicans, and blue represents Democrats.

Source: Forbes Inc.

This video is an introductory tutorial about Gephi (2016) by Jen Golbeck, professor at the University of Maryland. It is one of a series, so if you continue watching at the end of the video, the next one progresses to describe more advanced features of Gephi, including how to use color to highlight particular features of interest in the network graphs: <https://www.youtube.com/watch?v=HJ4Hcq3YX4k>.

In this YouTube video, Marc Smith, a sociologist and director of the Social Media Foundation, shares with the relationship mapping workgroup how he has used NodeXL for social media network analysis and visualization: https://www.youtube.com/watch?v=Ftssu_5x7Zk.

DILEMMA

How to Probe People's Reactions to Tracking

There is often a gulf between the benefits provided to society through tracking and the level of individual privacy that is being sacrificed. It is important, therefore, to have an open debate about the costs versus the benefits of using future tracking and monitoring technologies. Ideally, this should take place before any deployment of the new technology. However, just asking people what they think about a future tracking technology may not reveal the true extent of their concerns and feelings. What other methods could be used? One approach is to use a provocative probe.

For example, a project called the Quantified Toilets did this by setting up a fake service in a public place to disrupt the accepted state of affairs. The team was interested in how a community would react to having their urine analyzed in a public toilet with the goal of improving public health. They pretended to be a commercial company called the Quantified Toilets, which had created a new urine analysis technology infrastructure and installed it in the public toilets at a convention center. Signage was placed throughout the public toilets at a convention center explaining the rationale for the initiative (see Figure 10.6). In addition, the

This facility is proud to participate in the healthy building initiative.
Behaviour at these toilets is being recorded for analysis.
Access your live data at quantifiedtoilets.com



Figure 10.6 Signage posted in the convention center Quantified Toilets

Source: Used courtesy of Quantified Toilets

team created a website (quantifiedtoilets.com) that presented fake real-time data feeds from each of the toilets in the convention center showing the results of the urine analysis, including details such as blood alcohol levels, drugs detected, pregnancy, and odor (see Figure 10.7). All sampled data were anonymized. In addition, a link to a survey was added, and the general public was invited to give their feedback.

● Recent anonymized random data feed

| Time | Toilet ID | Sex | Deposit | Odor | Blood alcohol | Drugs detected | Pregnancy | Infections |
|-------------|-----------|--------|---------|---------|---------------|----------------|-----------|------------|
| 09:39:34 AM | T205 | female | 205ml | neutral | 0.061% | no | no | none |
| 09:33:20 AM | T109 | female | 175ml | neutral | 0.054% | no | no | none |
| 09:23:07 AM | T706 | female | 185ml | nutty | 0.000% | no | no | none |
| 09:19:02 AM | T715 | female | 75ml | neutral | 0.000% | no | no | none |
| 09:18:07 AM | T704 | female | 100ml | neutral | 0.000% | no | no | none |
| 09:11:56 AM | T706 | female | 80ml | neutral | 0.000% | no | no | none |
| 09:07:09 AM | T211 | male | 150ml | neutral | 0.001% | no | no | gonorrhea |
| 09:05:30 AM | T312 | male | 250ml | neutral | 0.001% | no | no | none |
| 09:00:39 AM | T314 | female | 245ml | neutral | 0.002% | no | no | chlamydia |
| 08:57:22 AM | T107 | male | 160ml | neutral | 0.000% | no | no | none |

Figure 10.7 The real-time data provided on the quantifiedtoilets.com website

Source: Used courtesy of Quantified Toilets

The goal was to observe people’s reactions when coming across this new service on going to the toilet. Would they become upset, surprised, or outraged, or wouldn’t they mind? Would they question the reality of the situation and tell others?

So, what happened? A diverse range of responses were observed. These included disapproval (for example, “Health advice? It does not get any creepier.”); approval (“Privacy is important. But I would like to know if I was sick and this is a good way to do so.”); concern (for instance, “Imagine if your employer could find out how hard you had partied last night.”; resignation (“I am sure the government has been collecting this data for years.”; voyeurism (“I just spent the last 10 minutes watching the pee-pee logs. Can’t stop watching them.”); and even humor, where some people tried to match people entering and exiting the toilets with the data appearing on the website.

Within an hour of the project going live, #quantifiedtoilets went viral on social media, triggering a snowball of tweets and retweets. Many face-to-face discussions took place at the convention center, and articles and blogs were written, some appearing in magazines and newspapers. Some visitors were duped and tweeted how incensed they were. Arguably, this range of responses and level of discussion would never have happened if the researchers had just asked people in the street would they mind if their urine were analyzed in a public toilet.

What do you think of this type of study? Do you think it is a good way to open up debate about data tracking in society, or is it a step too far? ■

10.2.6 Combining Multiple Sources of Data

A number of researchers have started collecting data from multiple sources by combining automatic sensing and subjective reporting. The goal is to obtain a more comprehensive picture about a domain, such as a population's mental health, than if only one or two sources of data were used (for instance, interviews or surveys). One of the first comprehensive studies to do this was Studentlife (Harari et al., 2017), which was concerned with learning more about students' mental health. In particular, the research team wanted to know why some students do better than others under times of stress, why some students burn out, and still others drop out. They were also interested in the effect of stress, mood, workload, sociability, sleep, and mental health on the students' academic performance. They were especially interested in how the students' moods change in response to their workload (such as their assignments, mid-terms, finals).

During a 10-week term, the researchers collected masses of data about a cohort of 48 students studying at Dartmouth College in the United States (Harari et al., 2017). They developed an app that ran on the students' phones, without the students needing to do anything, to measure the following:

- Wake up time, bed time, and sleep duration
- The number of conversations and duration of each conversation per day
- The kind and amount of physical activity (walking, sitting, running, standing, and so on)
- Where they were located and how long they stayed there (that is, in the dorm, in class, at a party, in the gym, and so forth)
- The number of people around the student throughout the day
- Student mobility outdoors and indoors (in campus buildings)
- Their stress levels throughout the day, week, and term
- Positive affects (how good they felt about themselves)
- Eating habits (where and when they ate)
- App usage
- Their comments on campus about national events (for example, the Boston bombing)

They also used a number of pre- and post-mental health surveys and collected the students' grades. These were used as ground truth for evaluating mental health and academic performance, respectively. The researchers went to great lengths to ensure that all of the data stored was anonymized in a dataset to protect the privacy of the participants. Having achieved this, the researchers then opened up the dataset for others to examine and use to conduct further analyses (<http://studentlife.cs.dartmouth.edu/dataset.html>).

The researchers were able to mine the data that they had collected automatically from the students' smartphones and learn several new things about their behavior. In particular, they found that a number of the behavioral factors that had been tracked from their smartphones were correlated to their grades, including activity, conversational interaction, mobility, class attendance, studying, and partying.

Figure 10.8 shows a graph indicating the relationship between activity, deadlines, attendance, and sleep. It shows that students are very active at the beginning of the term and get very little sleep. This suggests that they are out partying a lot. They also have high attendance rates at the beginning of term. As the term progresses, however, their behavior changes. Toward the end of term, sleep, attendance, and activity all drop off dramatically!




Figure 10.8 Student’s activity, sleep, and attendance levels against deadlines during a term
 Source: StudentLife Study

ACTIVITY 10.2

From the two graphs shown in Figure 10.9, what can you say about the students’ activity, their stress levels, and their level of socializing in relation to deadlines over the course of the term?





Figure 10.9 Student behavioral measures over the course of a term
 Source: StudentLife Study

Comment

The top figure shows that students start the term by having long social conversations. This begins to tail off as mid-term approaches. Students resort to having fewer and shorter conversations. After the deadlines are met, students switch back to having many more and longer conversations. The bottom figure shows students started out all upbeat, having returned from vacation feeling good about themselves. They appear relaxed (high mood level) and are active (going to the gym a lot). These attributes all start going downhill as the term comes to an end—presumably as their stress levels rise because of looming deadlines. ■

10.3 Visualizing and Exploring Data

Every day, people interact with different kinds of visualizations, including road signs, maps, medical images, mathematical abstractions, tables of figures, schematic diagrams, graphs, scatter plots, and many more. These representations are intended to help us make sense of the world we live in, but for them to be useful, they have to be presented in ways that are understandable for the people who use them. Being able to take meaning from data involves being able to see it and understand the way that it is represented and its context. What kind of data is it? What is the data about? Why was it collected? Why was it analyzed and represented in a particular way? The skills needed to understand and interpret visualizations are referred to as *visual literacy*. As with any skill, different people exhibit different levels of visual literacy, depending on their experience of using visual representations (Sarıkaya et al., 2018). Figure 10.10 shows a simplified path that is followed when data is meaningful. Starting with the analyzed data, which is represented in some way, the user perceives and interprets the data representation taking into account the context of the data. The user is then able to understand and communicate what the data shows to others.

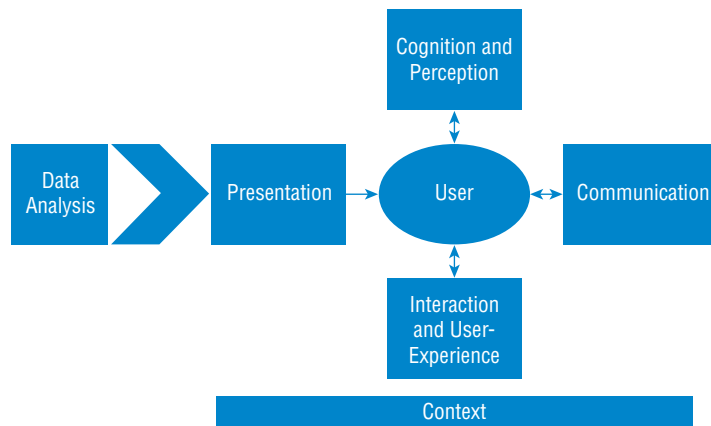


Figure 10.10 A simplified path for data to be meaningful

Source: Lugmayr et al. (2017). Used courtesy of Dr Artur Lugmayr

Even graphical representations of small amounts of data (for example 20–100 items) can be hard to interpret if the people trying to make sense of them don't understand the way that the data is being displayed. Furthermore, sometimes representations, such as bar graphs, line graphs, and scatter plots, are displayed in misleading ways. Danielle Szafir (2018), for example, asks, "How can we craft visualizations that effectively communicate the right information from our data?" She describes how data displays can mislead users when designers show axes with truncated scales, or they show data in 3D bars making it hard to read exact values from the bar because it isn't obvious which side of the 3D column is the place to read. Interactive visualizations typically include all of the canonical forms of representations (for instance, bar charts or pie charts) along with tree maps and advanced visualization techniques that enable users to interact with the data online by panning and zooming in and out of the displays. With the increased tendency to develop more complex visualizations to display increasingly large volumes of data, the question about how to craft the data representations and tools to develop and explore the data is even more relevant.

As Stu Card and his colleagues explained two decades ago, the goal of data visualization tools is to amplify human cognition so that users can see patterns, trends, correlations, and anomalies in the data that lead them to gain new insights and make new discoveries (Card et al., 1999). Many of the data visualizations and tools that have been developed since then are now being used by practitioners and researchers from fields including health and wellness, finance, business, science, educational analytics, decision-making, and personal exploration. For example, millions of people use interactive maps to find their way, benefiting from their integration into car navigation and car-sharing apps. Physicians and radiologists compare images from thousands of patients, and financiers examine trends in the stocks of hundreds of companies. Data visualization tools can help users change and manipulate variables to see what happens; for example, they can zoom in and out of the data to see an overview or to get details. Ben Shneiderman (1996) summarizes this behavior in his mantra "overview first, zoom and filter, and then details on demand."

While the early UX research on information visualizations still guides UX designers in their pursuit of designing new interactive visualizations, tools are needed for interacting with large volumes of data (Whitney, 2012). Many of these tools require expertise beyond that of most casual users in order to use them effectively. Typically, the data displays consist of many of the common techniques mentioned earlier (such as graphs and scatter plots) coupled with 3D interactive maps, time-series data, trees, heat maps, and networks (Munzner, 2014). Sometimes these visualizations were developed for uses other than those for which they are used today. For example, tree maps were originally developed to visualize file systems, enabling users to understand why they are running out of disk space on their hard drives by seeing how much space different applications and files were using (Shneiderman, 1992). However, tree maps were soon adopted by media and financial reporters for communicating changes in the stock market, and they became known as "market maps" (see Figure 10.11). Like interactive maps, tree maps have become a general-purpose tool embedded in most widely used applications, such as Microsoft's Excel (Shneiderman, 2016).

The ability to collect and store large amounts of data easily has encouraged the development of visualizations that display different types of data. For example, Figure 10.12 shows segments of sounds recorded from birds and other organisms collected by Jessie Oliver and her colleagues (2018). These researchers wanted to see how people investigated and annotated this data and in turn how this approach can be used to find and identify birds and other



Figure 10.11 A market map of the S&P 500, which is a financial index for stocks. Green indicates stocks that increased in value, and red indicates stocks that decreased in value that day.

Source: Used courtesy of FINVIZ

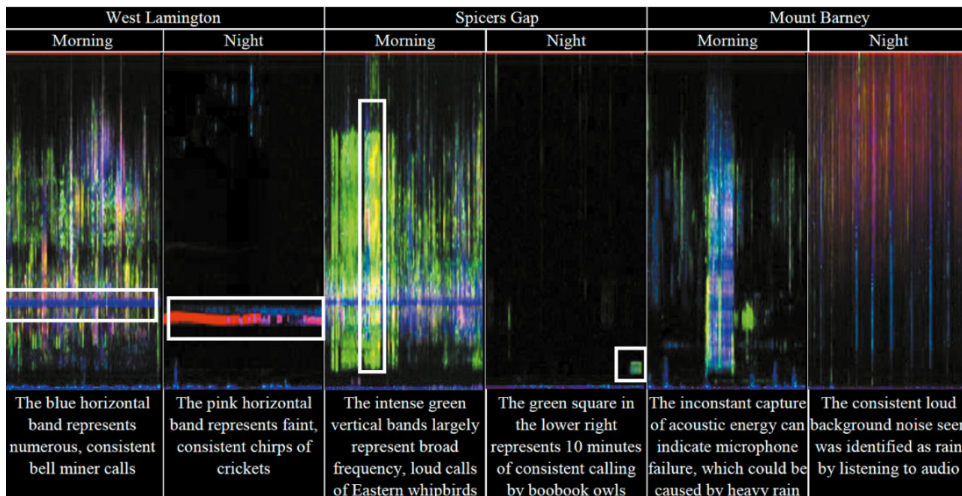


Figure 10.12 Visualization of different sounds, including birds, owls, and insects, from three areas of Australia that are displayed so they can be interpreted and compared

Source: Oliver et al. (2018). Reproduced with permission of ACM Publications

animals in the wild by recording their songs and calls. When the visualizations, known as *spectrograms*, were shown to birders, the researchers were intrigued to see how they evoked memories of hearing the birds in the field. The birders also found this data visualization to be helpful in corroborating their identifications of birds with other birders. From a UX design

perspective, Jessie Oliver and her colleagues faced the challenge of how to display long sound recordings visually. They used a technique developed by Michael Towey and his colleagues (2014) in which algorithms compress the spectrograms so that one pixel represents one minute of sound recording. The resulting spectrograms enable birders to get an overview of the recordings that in turn allows them to see patterns in the bird songs.

ACTIVITY 10.3

This video by Jeff Heer (2017) from University of Washington gives an overview of different types of data visualizations and data visualization tools: <https://www.youtube.com/watch?v=hsfWtPH2kDg>.

Watch this video and then describe (a) some of the benefits of using interactive visualizations and (b) some of the UX challenges in designing interactive visualizations.

Comment

1. By working with interactive visualizations, users can interact with data to explore aspects of interest by going deeper into particular parts of the data. This is demonstrated in the visualization of airline on-time performance in which a user can filter portions of the data to view which flights arrive late. From this exploration, the user will discover that flight delays are associated with it being late in the day. In other words, the data reveals that as time goes by, the actual arrival times of flights tend to fall further behind the scheduled arrival times. Also, by being able to filter and manipulate particular parts of the data, users can answer questions that would be difficult to investigate without data visualization tools, such as what causes flights to arrive early?
2. In the video, Jeff Heer talks about some of the human perceptual and cognitive issues about which UX designers must be aware when they create visualizations. For example, he mentions the importance of using color appropriately in a visualization of arteries. He also talks about the challenge of knowing how much detail to include in the visualization about the structure of the arteries.

In addition, Jeff mentions the power of many current tools for investigating many different variables, but he notes that using some of these tools proficiently requires programming and data analytics skills. UX visualization tool designers therefore need to find ways to support users who may not have these skills. He describes how some designers are tempted to get around this problem by automating the analyses. He points out, however, that a careful balance is needed in deciding how much automation should be provided and how much control should be left in the hands of users. Making this judgment is challenging for designers.

Jeff also mentions that there is much more to analyzing data than to visualizing it. Data has to be cleaned and prepared, a task referred to as *data wrangling*, which can take up to 80 percent of a data scientist's time. Issues of privacy also need to be considered. As a UX data visualization tool designer, Jeff suggests that all of these issues are challenges that must be considered when designing visualizations and data viz tools. ■

Powerful tools and platforms for analyzing and making predictions from large volumes of data have been designed for marketing, scientific and medical research, finance, business, and other kinds of professional use. To use these tools typically requires data analytic skills and statistical knowledge, which makes the potential benefits they offer out of reach for many people (Mittlstatd, 2012).

Many of these tools have been developed by large companies and research labs (Sakr et al., 2015). Some examples include Tableau, Qlik, Datapine, Voyager 2, Power BI, Zoho, and D3. To use these tools effectively, business managers often partner with analysts who assist them in interactive explorations that can lead to new insights. Together the analysts and managers identify widgets in the form of icons that represent their underlying functionality, from those made available in the tool, and then they create a customized “interactive dashboard” for use by the manager.

The *dashboard* is an interactive panel of control widgets that contains sliders, checkboxes, radio buttons, and coordinated multiple window displays of different kinds of graphical representations, such as bar and line graphs, heat maps, tree maps, infographics, word clouds, scatterplots, and other kinds of visualizations. Managers can then use these customized dashboards to explore the data and make informed decisions. All of the items in the dashboard are coordinated and draw from the same data selected to investigate particular questions of interest. In other words, the components of the dashboard are interactive and linked together so that they are coordinated (see Figure 10.13). This enables their users to benefit from seeing data displayed in different ways and to explore how these representations change as they manipulate sliders and other controls. The displays produced by tools like Microsoft’s Power BI, Tableau, and similar products are used by managers who can make the same dashboards available to other employees across their company. Everyone can then see, discuss, and interact with the same data.

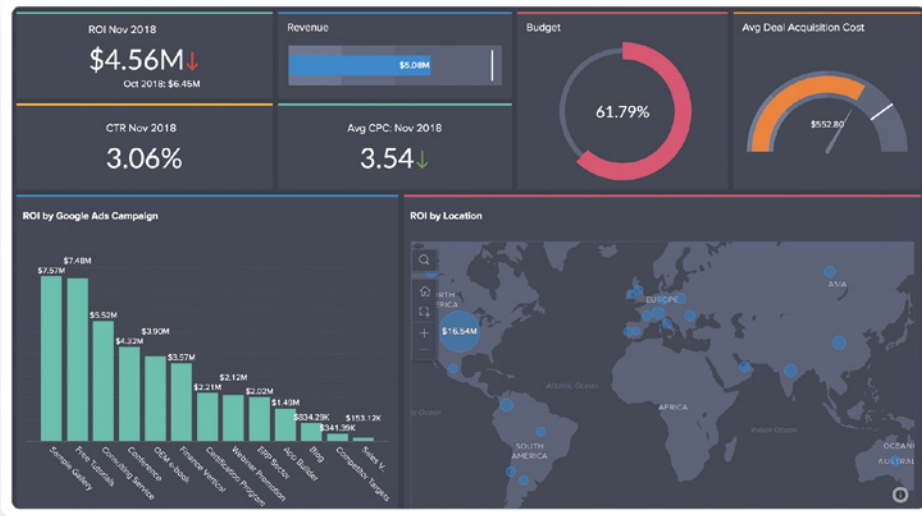


Figure 10.13 A dashboard that was created to show changes in sales information

Source: <https://www.zoho.com/analytics/tour.html>

For an overview of Tableau that shows how it is used and how Tableau dashboards are created, watch this video clip: <https://www.tableau.com/#hero-video>.

Another technique for creating interactive visualizations is *Data-Driven Documents* (D3) (Bostock et al., 2011). This tool is used to create web-based interactive displays. It is a powerful, specialist tool that extends JavaScript, and it requires programming expertise to use it effectively. It is used by journalists to create displays that appear in traditional news print and you can also interact with it online (see Figure 10.14).




Figure 10.14 An interactive graphic produced using D3 for the *New York Times*. It shows the tax rate paid by the different kinds of companies that form the S&P 500 financial index.

Source: Reproduced with permission of PARS International

Watch this *New York Times* graphic for an article entitled “Across U.S. Companies, Tax Rates Vary Greatly.” (Navigate to the following link to interact with the graphic. Try panning over the display.)

<https://archive.nytimes.com/www.nytimes.com/interactive/2013/05/25/sunday-review/corporate-taxes.html>.

A challenge is how to make powerful tools available to people who want to explore such topics as personal finance and health data but who are not trained as analysts and who do not want to employ or work with an analyst. Furthermore, some products are expensive and are unaffordable for many individuals and nonprofit organizations.

In a recent study, Alper Sarikaya and colleagues (2018) pointed out that the term *dashboard* requires more precise description and a deeper understanding of how the context of its use can impact the UX design of dashboards. They challenge UX designers to develop dashboards for different types of use cases for a wide range of users. In their study,

they analyzed a range of dashboards, first by reviewing published papers written by other researchers and then through a qualitative study in which they classified the features of different dashboards and how they are used.

They characterized the dashboards according to their design goals, levels of interaction, and the ways in which they are used. Figure 10.15 shows examples of the seven kinds of dashboards that they identified. Each type is named according to how it is used: strategic decision-making, quantified self, static operational, static organizational, operational decision-making, communication, and dashboards evolved, which was a catchall category that included features that did not fit into other categories. They state that many of these examples visually appeared as dashboards but may not fit the strictest definition of dashboard functionality.



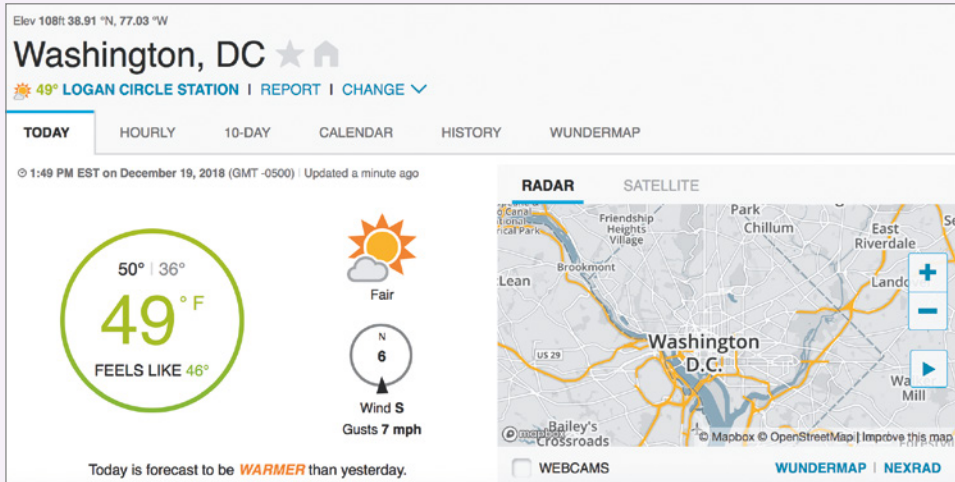
Figure 10.15 Exemplar dashboards (Sarikaya et al., 2018). Dashboard 1 and Dashboard 5 specifically target decision-making, while Dashboard 3 and Dashboard 4 target consumer awareness. Dashboard 2 represents the quantified self (such as a smart home), while Dashboard 6 represents those dashboards targeting communication. Dashboard 7 captures some novel extensions of traditional dashboards.

Source: Sarikaya et al. (2018), Graph 1. Reproduced with permission of IEEE.

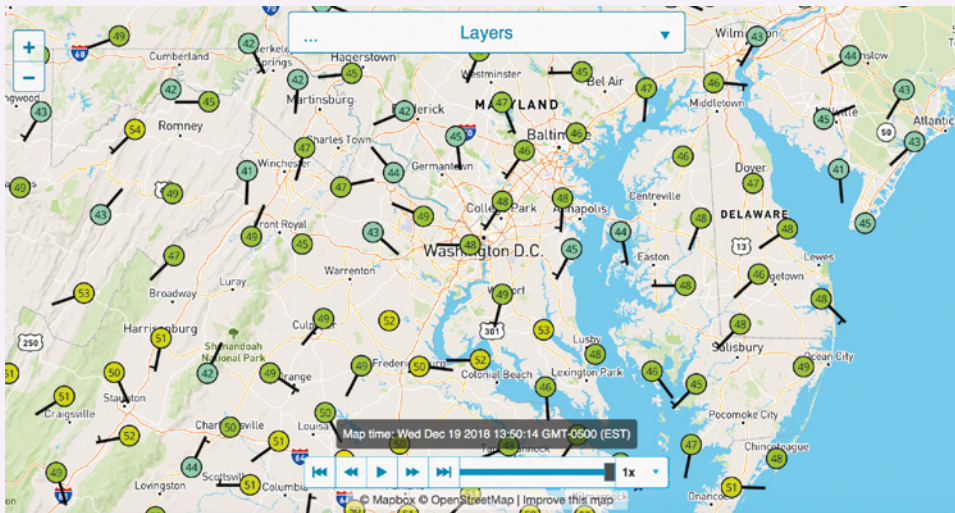
Sarikaya et al. also advocate ways of supporting users by telling stories that can help illustrate the context that the data visualizations represent. They point out the challenges that users encounter when interacting with visualizations such as enabling them to have more control over how they configure and use dashboards. A further challenge involves finding ways to support users in developing data and visual and analytic literacy. They also point out that there is a new opportunity for UX designers that involves finding ways to support users in making choices about which data and representations to use in different contexts. This includes understanding the broader social impact of dashboards.

ACTIVITY 10.4

Study Figure 10.16(a) from the weather site <https://www.wunderground.com>. It shows weather data for December 19, 2018, at Washington D.C. in the United States. Particularly take note of the temperature, precipitation, and wind data. What information do they provide? Now compare this visualization with that depicted in the “wundermap” (see Figure 10.16b). How do the two displays differ, and which do you prefer?



(a)



(b)

Figure 10.16 (a) Actual weather data and (b) a wundermap of the same area and time

Source: <https://www.wunderground.com>

Comment

The first display in Figure 10.16(a) contains representations that are fairly standard for conveying weather information. The green ring shows the maximum and minimum temperatures for now and what it feels like. A diagram of a sun indicates that it is a sunny day with some clouds, even though it is quite cold. It is also easy to see that the wind is from the south, and presumably the circle represents a compass and the pointed wedge indicates the wind direction.

The display in Figure 10.16(b) provides similar data, but it is harder to get an overview of weather in the Washington D.C. area. It uses conventional meteorological symbols to show temperature and wind. It is easier to see local effects but harder to get an overview of weather of the entire area. (If you are able to access the website, try clicking “layer” and selecting other options not shown in the figures.) Which display is preferable probably depends on how much detail you want—an overview or detail about a specific area in the Washington D.C. region—and your tolerance for clutter. ■

BOX 10.3


Visualizing the Same Sensor Data by Using Different Kinds of Representations for Environmentalists and the General Public

One of the parks in London that was used to stage the Olympics in 2012 has been transformed into a “smart park.” By this it is meant that a number of sensors were placed throughout the park to measure the health and use of the park. One type used throughout the park was bat call sensors. The goal was to ensure that the park’s bat conservation program was effective, as well as connecting visitors and residents to the wildlife around the park (<https://naturesmartcities.com/>). Monitoring bat calls is also a technique that was used to assess the general health of the park.

The data collected was primarily provided to the scientists in the form of spectrograms (see Figure 10.17b), but it was also presented in a more accessible form to the public via an interactive display (see Figure 10.17a). As part of a public kiosk, a schematic map was provided that showed where in the park the bat call data had been collected (Matej et al., 2018). A slider was provided to enable visitors to interact with the data: moving it to the left showed bat call data from the night before, while moving it to the right showed bat call data from the previous 10 nights. The LEDs on the map changed in color and intensity, representing the varying levels of bat calls. The total number was also shown in the digital display. The kiosk was deployed in the park, and many passersby stopped for a considerable length of time to learn about bats and interact with the data. The physical act of using the slider provided an engaging way of exploring the data rather than just looking at a static visualization or dashboard.



(a)



(b)

Figure 10.17 The same bat call data was made accessible (a) to the general public via an interactive visualization and (b) as a spectrogram intended for environmental scientists.

Source: (a) Used courtesy of Matej Kaninsky and (b) Used courtesy of Sarah Gallacher ■

10.4 Ethical Design Concerns

In the introduction to this chapter, we mentioned how masses of data are now regularly being collected from people for a variety of reasons, including improving public services, reducing congestion, and enhancing security measures. It is usually anonymized and sometimes aggregated to make it publicly available, for example showing the energy consumption data for a given space such as a floor of a building. Figure 10.18 shows a floor-by-floor comparison for a University of Melbourne building, where the red bar for the basement is the worst



Figure 10.18 Average daily energy consumption depicted on a public display for a building at the University of Melbourne. Green is best performer, yellow is in the middle, and red is the worst performer.

Source: Helen Sharp

performer in terms of energy usage and the green bar for Level 1 is the best performer. The idea is to provide feedback on energy consumption in the building to increase awareness among the inhabitants to encourage them to reduce their energy consumption. However, what if localized occupancy rates or energy consumption for each office were shown? It would not take much to figure out who was in that space. Would that be a step too far and an invasion of their privacy? Would people mind?

When deciding on how to analyze data and act upon data that has been automatically collected from different sensors, it is important to consider how ethical the data collection and storage processes are and how the data analysis will be used. By “ethics,” this is usually taken to mean “the standards of conduct that distinguish between right and wrong, good and bad, and so on” (Singer, 2011, p. 14). There are many codes of ethics available from official bodies that provide guidance. For example, the ACM (2018) and IEEE (2018) have both developed sets of ethics (<https://ethics.acm.org/2018-code-draft-1/>; IEEE Ethically Aligned Design, 2018). They point out that central to any ethical discussion is the importance of protecting fundamental human rights and to respect the diversity of all cultures. They also state the need to be fair, honest, trustworthy, and respectful of privacy.

To use data ethically, researchers and companies can limit the data they collect in the first place. Rather than trying to collect as much data as possible (as has often been the situation

in research—just in case it might be useful for subsequent analysis), it has been proposed that researchers and data practitioners follow an approach called *privacy by design* (Crowcroft et al., 2018). That way, they can avoid collecting excessive data that might be sensitive but not needed (see also Chapter 8 and Chapter 14). Furthermore, it may be possible to collect and analyze the data on the device itself, rather than uploading it to the cloud (Lane and Georgiev, 2015).

ACTIVITY 10.5

Watch the following TEDx talk by Jen Golbeck (the author of the *Atlantic* article on the Quantified Toilets) where she discusses why social media “likes” say more than you think. The talk was given in 2013 and since then has had more than 2.25 million views. Even though the TEDx talk is a few years old, the issues raised in it are still relevant today. In particular, she discusses how people’s behavior online enables companies to predict what they like, what they might be interested in buying, and even their political views.

https://www.ted.com/talks/jennifer_golbeck_the_curly_fry_conundrum_why_social_media_likes_say_more_than_you_might_think?language=en&utm_campaign=tedsread&utm_medium=referral&utm_source=tedcomshare

What do you think the privacy issues are here?

Comment

Jen Golbeck provides two compelling examples in her talk. The first is the well-known example of how a teenage girl’s pregnancy was predicted from her online purchases of things like vitamins. The second example was how data on liking crinkly fries coupled with a knowledge of the theory of homophily was used to predict that a group of people have above average intelligence. By understanding that the theory of homophily explains that people who are similar tend to like the same things, trust each other, and seek out each other’s company, Jen Golbeck was able to look for relationships in data about “liking” crinkly fries. The crinkly fries example indicates that even though it is absurd that liking crinkly fries is a predictor of above average intelligence, in this particular example, the person who created the post attracted “likes” from friends who were also of above average intelligence. It is an amusing example, but the main point is to illustrate that information that people contribute in social media, often unknowingly, can be used to infer all kinds of things about them, such as their ethnicity, age, gender, shopping behavior, and what they like.

The concerns highlighted in the video are prescient for politicians and others looking for ways to protect the general public by controlling what social media companies can and cannot do with personal data. For example, GDPR is a law introduced by the European Union that seeks to protect data privacy (discussed in Chapter 8, “Data Analysis,” and Chapter 14, “Introducing Evaluation”). Within the United States and in other countries, the need for controlling how personal data is used is being debated by governments and consumer protection groups, such as the Electronic Privacy Information Center (EPIC; www.epic.org), with the goal of finding ways to protect data privacy. ■

An ethical strategy that can be adopted for systems that analyze data is to have an explicit agreement in place as to how it will be used and acted upon. Such an agreement can include in what way is the analysis trustworthy. *Trustworthiness* is usually taken to mean how credible are the analyses being performed on the data (see Davis, 2012). When a decision is made on behalf of a human based on classifying their data using some machine learning algorithm, can the user be sure that the decision made is trustworthy?

Another ethical concern is whether the form of data analysis being used by a system is *socially acceptable* (Harper et al., 2008). Have clear *boundaries* been established between what is acceptable and not acceptable, especially when it is personal data that is being analysed and classified, such as health data or criminal history? Is there a clear understanding about the data that is being analyzed to provide information about a phenomenon versus that which is performed to make a decision about someone's future? Are there agreed policies in place? How much do the boundaries shift and change over time as new technology becomes more mainstream and authorities change?

ACTIVITY 10.6

Shoplifting cost U.S. retailers \$44 billion in 2014. To help combat shoplifting, DeepCam developed an intelligent system that passively monitors people coming into a store by using CCTV video footage that identifies potential suspects (see Figure 10.19). To do this, it uses AI algorithms and facial recognition software. Do you think this practice is socially acceptable? What might be the privacy concerns? To find out more about their system, check out their website at <https://deepcamai.com/>.



Figure 10.19 DeepCam's face-tracking software used in a store

Source: <https://deepcamai.com>

Comment

To address privacy concerns, the company developed its system so that it does not identify customers or link them to any sensitive information such as name, address, or date of birth. It only recognizes faces and identifies patterns of behavior that potentially are worth investigating. The video footage is indexed and structured similar to how web pages are set up for quick searching. This enables store detectives to be able to notice potential threats in real time. Many people might find this form of data analysis creepy, knowing that their faces are being matched to a database each time they enter a store. Others might find it more socially acceptable because it has the potential to reduce crime considerably. ■

The Open Data Institute (<https://www.theodi.org>) has provided a set of questions to help researchers, system developers, and data scientists to formulate a set of ethical questions to begin to address these concerns. These are grouped as sets of questions as part of a framework called the *data ethics canvas*. For example, two subsets are about the positive and negative effects that a project can have on people. The questions include “Which individuals, demographics, and organizations will be positively affected by the project?” and “How is positive impact being measured?” The negative questions include “Could the manner in which the data is collected, shared, and used cause harm?” and “Could people perceive it to be harmful?” Working through each set of questions is intended to help researchers identify potential ethical issues for a data project or activity and to encourage explicit reflection and debate within a project team as to who the project will impact and the steps needed to ensure that the project is ethical. The framework provided on the ODI website is intended to help organizations identify potential ethical issues associated with their data project.

In Chapter 1, “What Is Interaction Design?” we outlined a number of usability and UX design principles that were transformed into questions, criteria, and examples showing how to use them in the design process. Here, we introduce some other principles that relate to the ethics of collecting and using data at scale and that are often talked about in the literature on ethics, data science, HCI, and AI (see Cramer et al., 2008; Molich et al., 2001; Crowcroft et al., 2018; Chuang and Pfeil, 2018; and van den Hoven, 2015). We call these *data ethics principles* (see Box 10.4). Four principles that often appear in the reports, handbooks, and articles on ethics and interaction design are fairness, accountability, transparency, and explainability (FATE). They are also included as key principles that lie at the heart of the general data protection regime (GDPR). For example, Article 5 of the GDPR requires that personal data shall be “processed lawfully, fairly, and in a transparent manner in relation to individuals.” Within the context of HCI, Abdul et al. (2018) have proposed an agenda for how HCI researchers can help to develop more accountable intelligent systems that don’t just explain their algorithms but are also usable and useful to people.

It should be noted that the ethics principles are not mutually exclusive but interrelated as described next.

BOX 10.4

Data Ethics Principles (FATE)

Fairness *Fairness* refers to impartial and just treatment or behavior without favoritism or discrimination. While this is something to which organizations adhere in areas such as promotion and hiring, in the context of data analysis, it refers to how fair a dataset is and what the impact will be from using the results. For example, sometimes, the dataset is biased toward a particular demographic that results in unfair decisions being made. If these could be identified and revealed by the system, it would make it possible to rectify them while also developing new algorithms that can make the system fairer.

Accountability *Accountability* refers to whether an intelligent or automated system that uses AI algorithms can explain its decisions in ways that enable people to believe they are accurate and correct. This involves making clear how decisions are made from the datasets that are used. A question that arises is who is accountable for doing this? Is it the person providing the data, the company coding the algorithms, or the organization that is deploying the algorithms for its own purposes?

Transparency *Transparency* refers to the extent to which a system makes its decisions visible and how they were derived (see Maurya, 2018). There has been much debate about whether AI systems, which typically depend on large datasets when making a decision, should be designed to be more transparent (see Brkan, 2017). Examples include medical decision-making systems that can diagnose types of cancer and media service providers (for instance, Netflix) that suggest new content for you to watch based on their machine learning algorithms. Currently, many are black-box in nature; that is, they come up with solutions and decisions without any explanation as to how they were derived. Many people think this practice is unacceptable, especially as AI systems are given more responsibility to act on behalf of society, for example, deciding who goes to prison, who gets a loan, who gets the latest medical treatment, and so on. Some of the rules of the GDPR on automated decision-making are also concerned with how to ensure the transparency of decisions made by machine learning algorithms (Brkan, 2017).

Explainability *Explainability* refers to a growing expectation in HCI and AI that systems, especially those that collect data and make decisions about people, provide explanations that laypeople can understand. Research into what is a good explanation to provide has been the subject of much research since expert systems came into being in the 1980s. Following this early work, there was research into what context-aware systems should provide. For example, Brian Lim et al. (2009) conducted a study that provided different kinds of explanations for a system that made automated decisions. They found that explanations describing why a system behaved in a certain way resulted in a better understanding and stronger feelings of trust. In contrast, explanations describing why the system did not behave a certain way resulted in lower understanding. More recently, research has investigated the kinds of explanations that are appropriate and helpful for users of automated systems (see Binnes et al., 2018).

The FATE framework suggests that the design of future systems, which use AI algorithms in combination with personal or societal data, should ensure that they are fair, transparent, accountable, and explainable. Achieving this goal is complex, and it involves being aware of the potential for bias, discrimination in big data and algorithms, ethics in big data, legal and policy implications, data privacy, and transparency (Abdul et al., 2018).

Achieving this objective is inevitably difficult. For example, as pointed out by Cynthia Dwork at a panel on big data and transparency (transcribed by Maurya, 2018), it is difficult to know what a good explanation of a decision might be for human beings. She uses the example of what should a system say when a user asks, “Why was I turned down for the loan?” to illustrate this. The system might be able to reply, “There is a classifier, we feed your data into it, and the outcome was that you were turned down.” However, that is of little help to a user, and it is likely to be more annoying than not having any explanation.

Reuben Binnes et al. (2018) conducted an experiment to determine what kinds of explanations users found to be fair, accountable, and transparent for an automated system. In particular, they compared four different styles of explanation, ranging from being largely numerical scores to more comprehensive ones that provided a breakdown of the statistics used for certain demographic categories, including age, gender, income level, or occupation. The different styles were presented for scenarios in which a decision had been made about individuals automatically, such as applying for a personal financial loan and where passengers on over-booked airline flights were selected for rerouting. The results of their experiment showed that some of the participants found that they engaged with the explanations to assess the fairness of the decisions being made, but at times they found them impersonal, and even dehumanizing. What constitutes a fair explanation may need to be more than providing an account of the processes used by the algorithms. From an interaction design perspective, it might help if the explanations were interactive, enabling the user to interrogate and negotiate with the system, especially if a decision that has been made is contrary to what they expected or had hoped.

Jure Leskovec (2018) comments on how the consequences of a system making a decision on behalf of a human can vary. This will determine whether an explanation is needed to support a decision made by a system and what it should include. For example, if a decision is made to pop up an ad for slippers in a user’s browser, based on an analysis of their tracked online app usage (a common practice used in targeted advertising), it might be mildly annoying, but it is unlikely to upset them. However, if it means a person is going to be denied a loan or a visa based on the outcome of an automated algorithm, it may have more dire consequences for someone’s life, and they would want to know why the particular decision was made. Jure suggests that humans and algorithms need to work together for system decisions that implicate more important societal concerns.

Another reason why ethics and data have become a big concern is that automated systems that rely on existing datasets can sometimes make decisions that are erroneous or biased toward a particular set of criteria. In doing so, they end up being unfair. An example that caused a public outcry was the misidentification of people with dark skin. Traditional AI systems have been found to have much higher error rates for this demographic. In particular, 35 percent of darker-skinned females were misidentified compared with 7 percent of lighter-skinned females (Buolamwini and Gebru, 2018). This difference was exacerbated by the error rate found for lighter-skinned males, which was less than one percent. One of the main reasons for this large discrepancy in misidentification is thought to be due to the make-up of the images in the datasets used. One widely used collection of images was estimated to have more than 80 percent white images of which most were male.

(Continued)

This bias is clearly unacceptable. The challenge facing companies that want to use or provide this data corpora is to develop fairer, more transparent, and accountable facial analysis algorithms that can classify people more accurately regardless of demographics, such as skin color or gender. A number of AI researchers have begun addressing this problem. Some have started developing 3D facial algorithms that continuously learn multiracial characteristics from 2D pictures. Others have introduced new face datasets that are more balanced (see Buolamwini and Gebru, 2018). ■

BOX 10.5

The Living Room of the Future: An Ethical Approach to Using Personal Data

There are now more than 250 *smart cities* projects throughout the world in nearly 200 cities. Each has different aspirations, but a major goal is to make cities more energy efficient, to make them safer, and to improve the quality of life. Most work in partnership with tech companies, central councils, and local communities to realize the benefits of new economic opportunities. IoT technologies and big data are often a central concern. One aspect has been to develop new approaches and toolkits to empower local individuals with the tools to measure and act upon what they want to find out about their city. For example, a citizen-sensing approach, adopted in the city of Bristol, brought together local citizens, community workers, and volunteers who together developed an innovative and affordable DIY sensing tool that people could use to record and collect data about the level of dampness in their homes (www.bristolapproach.org/).

In addition to the many smart cities projects, there are others that focus on how people adopt, accept, and approach new sensing technologies in a particular building or home. For example, the Living Room of the Future project is investigating how people will live in the future in a home that has been embedded with a range of IoT devices (<https://www.bbc.co.uk/rd/projects/living-room-of-the-future>). The project is researching how to make personal data transparent and trustworthy, while at the same time respecting people's privacy. It is also concerned with designing methods that can offer explicit awareness and transparency of those individuals' data.

A particular challenge that they are addressing is how to enable people to be in control of their own personal data while at the same time letting it be used by the home system to adapt their experiences, for example, choosing what media to play and from which device(s). The devices, like sensors and everyday RFID-enabled objects, have the potential of being able to collect a diversity of personal data, including music preferences, history of where everyone is sitting, what they are doing, and how they are feeling. To ensure privacy, the data is collected and stored, not in the cloud but on a home-fenced data server called a *databox*. The box collates, curates, and mediates access to each person's data. It only uses verified and audited

third-party services to ensure that it cannot be accessed by anyone else. As such, the data never leaves the living room and ensures that the IoT services provided can be deemed trustworthy. It also provides a platform that allows personalization to occur without needing to ask the users if they are OK with any suggested changes.

The home of the future project is also investigating how data can be collected and how it is being used in the house, for example, to control home heating and lighting. ■

In-Depth Activity

Go to labinthewild.com, and select the test “What is your privacy profile?” This test has been designed to tell you what you think about data privacy and how different you are compared to what others think about this topic. It should take about 10–15 minutes to complete. At the end of the test, it will provide you with your results and classify you in terms of whether you are not concerned, somewhat concerned, or very concerned.

1. Do you consider this to be an accurate reflection of how you view privacy?
2. Did you think the video shown was effective at raising potential problems of what data is collected in a smart building? If not, what other scenario could be used in a video to ask people to consider privacy concerns?
3. What impact do you think the context chosen for the scenario might have on your reactions? For example, if the scenario involved a doctor’s surgery, might you have reacted differently and if so, why?
4. What do you think of labinthewild.com as a platform for conducting large-scale online experiments from volunteers?
5. Did you find any other information on the website interesting?

Summary

This chapter described how data at scale involves bringing together large volumes of data from different sources that is then analyzed to address new questions and provide insights that could not be gained by analyzing data from a single source. The chapter explains techniques and tools for collecting and analyzing large volumes of data. It also raises some concerns about how data at scale is used, particularly as to the need for personal data privacy. UX designers are encouraged to consider the impact of their designs on how data is used and to ensure that it is used ethically. Four core principles are advocated for ethical design: fairness, accountability, transparency, and explainability (FATE).

Key Points

- Data at scale concerns very large volumes of data, which is also known as *big data*.
- A defining feature of data at scale is that it includes different types of data collected from different sources that are analyzed to address particular questions.
- Data at scale can be quantitative and qualitative data; it consists of social media messages, sentiment and facial recognition data, documents, sensor, sound and sonic data, and video surveillance data.
- Analyzing data from different sources is powerful because it provides different perspectives on people's behavior.
- Analyzing data at scale can have positive outcomes, such as understanding people's health problems, but there are also dangers if personal data is revealed and then misused.
- Data at scale is collected and analyzed in many different ways including data scraping, monitoring oneself and others, crowdsourcing, and sentiment and social network analysis.
- Data visualization provides tools and techniques for representing, understanding, and exploring data.
- Ethical design principles suggest ways that UX designers can create designs and interaction processes that make clear how data is being used.
- Ensuring that artificial intelligence systems are transparent is a particularly important ethical design principle. ■

Further Reading

HANSEN, D., SHNEIDERMAN, B., SMITH, M. A. AND HIMELBOIN, I., (2019) *Analyzing Social Media Networks with NodeXL. Insights from a Connected World* (2nd ed.). Morgan Kaufmann. This book provides an introduction to social network analysis. It focuses on NodeXL, but much of the discussion is helpful when using any social network analysis tool.

SCHRAEFEL, M.C. GOMER, R. ALAN, A. GERDING, E., AND MAPLE, C. (2017) The Internet of Things: Interaction Challenges to Meaningful Consent at Scale. *Interactions*, 24, 6 (October 2017), 26–33. This short article discusses how HCI researchers can be involved in helping users manage their privacy and personal data especially in view of IoT.

SZAFIR, D. (2018) The good, the bad, and the biases: Five Ways Visualizations Can Mislead and How to Fix Them. *Interactions*. xxv.4. As the title suggests, this article discusses some of the well-known problems and design flaws with visualizations and suggests ways to fix them.

SARIKAYA, A., CORELL, M., BARTRAM, L, TOREY, M, FISHER, D. (2018) What Do We Talk About When We Talk About Dashboards? *IEEE Trans Vis Comput Graph*. This paper characterizes dashboards, and it reviews and critiques their design and how they are used.

SHILTON, K (2018) Values and Ethics in Human-Computer Interaction. *Foundations and Trends in Human-Computer Interaction*: Vol. 12, No. 2, 107–171. This article provides a good overview of issues being debated in HCI about ethics, data, and HCI.

Chapter 11

DISCOVERING REQUIREMENTS

11.1 Introduction

11.2 What, How, and Why?

11.3 What Are Requirements?

11.4 Data Gathering for Requirements

11.5 Bringing Requirements to Life: Personas and Scenarios

11.6 Capturing Interaction with Use Cases

Objectives

The main goals of the chapter are to accomplish the following:

- Describe different kinds of requirements.
- Allow you to identify different kinds of requirements from a simple description.
- Explain additional data gathering techniques and how they may be used to discover requirements.
- Enable you to develop a persona and a scenario from a simple description.
- Describe use cases as a way to capture interaction in detail.

11.1 Introduction

Discovering requirements focuses on exploring the problem space and defining what will be developed. In the case of interaction design, this includes: understanding the target users and their capabilities; how a new product might support users in their daily lives; users' current tasks, goals, and contexts; constraints on the product's performance; and so on. This understanding forms the basis of the product's requirements and underpins design and construction.

It may seem artificial to distinguish between requirements, design, and evaluation activities because they are so closely related, especially in an iterative development cycle like the one used for interaction design. In practice, they are all intertwined, with some design taking place while requirements are being discovered and the design evolving through a series of evaluation—redesign cycles. With short, iterative development cycles, it's easy to confuse the

purpose of different activities. However, each of them has a different emphasis and specific goals, and each of them is necessary to produce a quality product.

This chapter describes the requirements activity in more detail, and it introduces some techniques specifically used to explore the problem space, define what to build, and characterize the target audience.

11.2 What, How, and Why?

This section briefly considers what is the purpose of the requirements activity, how to capture requirements, and why bother at all.

11.2.1 What Is the Purpose of the Requirements Activity?

The *requirements activity* sits in the first two phases of the double diamond of design, introduced in Chapter 2, “The Process of Interaction Design.” These two phases involve exploring the problem space to gain insights about the problem and establishing a description of what will be developed. The techniques described in this chapter support these activities, and they capture the outcomes in terms of requirements for the product plus any supporting artifacts.

Requirements may be discovered through targeted activities, or tangentially during product evaluation, prototyping, design, and construction. Along with the wider interaction design lifecycle, requirements discovery is iterative, and the iterative cycles ensure that the lessons learned from any of these activities feed into each other. In practice, requirements evolve and develop as the stakeholders interact with designs and learn what is possible and how features can be used. And, as shown in the interaction design lifecycle model in Chapter 2, the activity itself will be repeatedly revisited.

11.2.2 How to Capture Requirements Once They Are Discovered?


Requirements may be captured in several different forms. For some products, such as an exercise monitoring app, it may be appropriate to capture requirements implicitly through a prototype or operational product. For others, such as process control software in a factory, a more detailed understanding of the required behavior is needed before prototyping or construction begins, and a structured or rigorous notation may be used to investigate the product’s requirements. In all cases, capturing requirements explicitly is beneficial in order to make sure that key requirements aren’t lost through the iterations. Interactive products span a wide range of domains with differing constraints and user expectations. Although it may be disappointing if a new app to alert shoppers about offers on their favorite purchases turns out to be unusable or slightly inaccurate, if the same happens to an air traffic control system, the consequences are far more significant and could threaten lives.

As we discuss in this section, there are different kinds of requirements, and each can be emphasized or de-emphasized by different notations because notations emphasize different characteristics. For example, requirements for a product that relies on processing large amounts of data will be captured using a notation that emphasizes data characteristics. This means that a range of representations is used including prototypes, stories, diagrams, and photographs, as appropriate for the product under development.

11.2.3 Why Bother? Avoiding Miscommunication

One of the goals of interaction design is to produce usable products that support the way that people communicate and interact in their everyday and working lives. Discovering and communicating requirements helps to advance this goal, because defining what needs to be built supports technical developers and allows users to contribute more effectively. If the product turns out to be unusable or inappropriate, then everyone will be disappointed.

User-centered design with repeated iteration and evaluation along with user involvement mitigates against this from happening. The following cartoon illustrates the consequences of misunderstanding or miscommunication. The goal of an iterative user-centered approach is to involve different perspectives and make sure that there is agreement. Miscommunication is more likely if requirements are not clearly articulated.



11.3 What Are Requirements?

A *requirement* is a statement about an intended product that specifies what it is expected to do or how it will perform. For example, a requirement for a smartwatch GPS app might be that the time to load a map is less than half a second. Another, less precise requirement might be for teenagers to find the smartwatch appealing. In the latter example, the requirements

activity would involve exploring in more detail exactly what would make such a watch appealing to teenagers.

One of the goals of the requirements activity is to identify, clarify, and capture the requirements. The process of discovering requirements is iterative, allowing requirements and their understanding to evolve. In addition to capturing the requirements themselves, this activity also involves specifying criteria that can be used to show when the requirements have been fulfilled. For example, usability and user experience criteria can be used in this way.

Requirements come in different forms and at different levels of abstraction. The example requirement shown in Figure 11.1(a) is expressed using a generic requirements structure called an *atomic requirements shell* (Robertson and Robertson, 2013); Figure 11.1(b) describes the shell and its fields. Note the inclusion of a “fit criterion,” which can be used to assess when the solution meets the requirement, and also note the indications of “customer satisfaction,” “dissatisfaction,” and “priority.” This shell indicates the information about a requirement that needs to be identified in order to understand it. The shell is from a range of resources, collectively called *Volere* (<http://www.volere.co.uk>), which is a generic requirements framework. Although not specifically designed for interaction design, *Volere* is widely used in many different domains and has been extended to include UX analytics (Porter et al, 2014).

An alternative way to capture what a product is intended to do is via user stories. *User stories* communicate requirements between team members. Each one represents a unit of customer-visible functionality and serves as a starting point for a conversation to extend and clarify requirements. User stories may also be used to capture usability and user experience goals. Originally, user stories were normally written on physical cards that deliberately constrained the amount of information that could be captured in order to prompt conversations between stakeholders. While these conversations are still highly valued, the use of digital support tools such as Jira (<https://www.atlassian.com/software/jira>) has meant that additional information to elaborate the requirement is often stored with user stories. As an example, this additional information might be detailed diagrams or screenshots.

A user story represents a small chunk of value that can be delivered during a sprint (a short timebox of development activity, often about two weeks’ long), and a common and simple structure for user stories is as follows:

- As a <role>, I want <behavior> so that <benefit>.

Example user stories for a travel organizer might be:

- As a <traveler>, I want <to save my favorite airline for all my flights> so that <I will be able to collect air miles>.
- As a <travel agent>, I want <my special discount rates to be displayed to me> so that <I can offer my clients competitive rates>.

User stories are most prevalent when using an agile approach to product development. User stories form the basis of planning for a sprint and are the building blocks from which the product is constructed. Once completed and ready for development, a story consists of a description, an estimate of the time it will take to develop, and an acceptance test that determines how to measure when the requirement has been fulfilled. It is common for a user story such as the earlier ones to be decomposed further into smaller stories, often called *tasks*.

| | | |
|---|-----------------------------|---------------------|
| Requirement #: 75 | Requirement Type: 9 | Event/use case #: 6 |
| Description: The product shall issue an alert if a weather station fails to transmit readings. | | |
| Rationale: Failure to transmit readings might indicate that the weather station is faulty and needs maintenance, and that the data used to predict freezing roads may be incomplete. | | |
| Source: Road Engineers | | |
| Fit Criterion: For each weather station the product shall communicate to the user when the recorded number of each type of reading per hour is not within the manufacturer's specified range of the expected number of readings per hour. | | |
| Customer Satisfaction: 3 | Customer Dissatisfaction: 5 | |
| Dependencies: None | Conflicts: None | |
| Supporting Materials: Specification of Rosa Weather Station | | |
| History: Raised by GBG, 28 July | | |

Volere
Copyright © Atlantic Systems Guild

(a)

| | | |
|--|---|--|
| Requirement #: Unique id | Requirement Type: <i>The type from the template</i> | Event/use case #: <i>List of events / use cases that need this requirement</i> |
| Description: A one sentence statement of the intention of the requirement | | |
| Rationale: A justification of the requirement | | |
| Originator: The person who raised this requirement | | |
| Fit Criterion: A measurement of the requirement such that it is possible to test if the solution matches the original requirement | | |
| Customer Satisfaction: | Customer Dissatisfaction: | Other requirements that cannot be implemented if this one is |
| Priority: A rating of the customer value | Conflicts: | |
| Supporting Materials: Pointer to documents that illustrate and explain this requirement | Volere Copyright © Atlantic Systems Guild | |
| History: Creation, changes, deletions, etc. | | |

Degree of stakeholder happiness if this requirement is successfully implemented.
Scale from 1 = uninterested to 5 = extremely pleased.

Measure of stakeholder unhappiness if this requirement is not part of the final product.
Scale from 1 = hardly matters to 5 = extremely displeased.

(b)

Figure 11.1 (a) An example requirement expressed using an atomic requirements shell from Volere (b) the structure of an atomic requirements shell

Source: Atlantic Systems Guild

During the early stages of development, requirements may emerge in the form of epics. An *epic* is a user story that may take weeks or months to implement. Epics will be broken down into smaller chunks of effort (user stories), before they are pulled into a sprint. Example epics for a travel organizer app might be the following:

- As a <group traveler>, I want <to choose from a range of potential vacations that suit the group's preferences> so that <the whole group can have a good time>.
- As a <group traveler>, I want <to know the visa restrictions for everyone in the group> so that <visas can be arranged for everyone in the group in plenty of time>.
- As a <group traveler>, I want <to know the vaccinations required to visit the chosen destination> so that <vaccinations can be arranged for everyone in the group in plenty of time>.
- As a <travel agent>, I want <up-to-date information displayed> so that <my clients receive accurate information>.

11.3.1 Different Kinds of Requirements

Requirements come from several sources: from the user community, from the business community, or as a result of the technology to be applied. Two different kinds of requirements have traditionally been identified: *functional requirements*, which describe what the product will do, and *nonfunctional requirements*, which describe the characteristics (sometimes called *constraints*) of the product. For example, a functional requirement for a new video game might be that it will be challenging for a range of user abilities. This requirement might then be decomposed into more specific requirements detailing the structure of challenges in the game, for instance, levels of mastery, hidden tips and tricks, magical objects, and so on. A *nonfunctional requirement* for this same game might be that it can run on a variety of platforms, such as the Microsoft Xbox, Sony PlayStation, and Nintendo Switch game systems. Interaction design involves understanding both functional and nonfunctional requirements.

There are many more different types of requirements, however. Suzanne and James Robertson (2013) suggest a comprehensive categorized set of requirements types (see Table 11.1), while Ellen Gottesdiener and Mary Gorman (2012) suggest seven product dimensions (see Figure 11.2).

| | |
|-------------------------|--|
| Project Drivers | 1. The Purpose of the Product |
| | 2. The Stakeholders |
| Project Constraints | 3. Mandated Constraints |
| | 4. Naming Conventions and Terminology |
| | 5. Relevant Facts and Assumptions |
| Functional Requirements | 6. The Scope of the Work |
| | 7. Business Data Model and Data Dictionary |
| | 8. The Scope of the Product |
| | 9. Functional Requirements |

| | |
|----------------------------|--|
| Nonfunctional Requirements | 10. Look and Feel Requirements |
| | 11. Usability and Humanity Requirements |
| | 12. Performance Requirements |
| | 13. Operational and Environmental Requirements |
| | 14. Maintainability and Support Requirements |
| | 15. Security Requirements |
| | 16. Cultural Requirements |
| Project Issues | 17. Compliance Requirements |
| | 18. Open Issues |
| | 19. Off-the-Shelf Solutions |
| | 20. New Problems |
| | 21. Tasks |
| | 22. Migration to the New Product |
| | 23. Risks |
| | 24. Costs |
| | 25. User Documentation and Training |
| | 26. Waiting Room |
| | 27. Ideas for Solutions |

Table 11.1 A comprehensive categorization of requirements types

Source: Atlantic Systems Guild, Volere Requirements Specification Template, Edition 18 (2017), <http://www.volere.co.uk/template.htm>

The 7 Product Dimensions


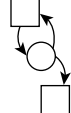

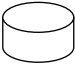
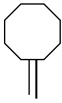


| | | | | | | |
|---|---|---|---|---|---|---|
|  |  |  |  |  |  |  |
| User | Interface | Action | Data | Control | Environment | Quality Attribute |
| Users interact with the product | The product connects to users, systems, and devices | The product provides capabilities for users | The product includes a repository of data and useful information | The product enforces constraints | The product conforms to physical properties and technology platforms | The product has certain properties that qualify its operation and development |

Figure 11.2 The seven product dimensions

Source: Gottesdiener and Gorman (2012), p. 58. Used courtesy of Ellen Gottesdiener

To see how the seven product dimensions can be used to discover requirements, see <https://www.youtube.com/watch?v=x9olpZaXTDs>.

In this section, six of the most common types of requirements are discussed: functional, data, environment, user, usability, and user experience.

Functional requirements capture what the product will do. For example, a functional requirement for a robot working in a car assembly plant might be that it is able to place and weld together the correct pieces of metal accurately. Understanding the functional requirements for an interactive product is fundamental.

Data requirements capture the type, volatility, size/amount, persistence, accuracy, and value of the required data. All interactive products have to handle some data. For example, if an application for buying and selling stocks and shares is being developed, then the data must be up-to-date and accurate, and it is likely to change many times a day. In the personal banking domain, data must be accurate and persist over many months and probably years, and there will be plenty of it.

Environmental requirements, or context of use, refer to the circumstances in which the interactive product will operate. Four aspects of the environment lead to different types of requirements. First is the *physical environment*, such as how much lighting, noise, movement, and dust is expected in the operational environment. Will users need to wear protective clothing, such as large gloves or headgear that might affect the choice of interface type? How crowded is the environment? For example, an ATM operates in a very public physical environment, thus using a speech interface is likely to be problematic.

The second aspect of the environment is the *social environment*. Issues regarding the social aspects of interaction design, such as collaboration and coordination, were raised in Chapter 5, “Social Interaction.” For example, will data need to be shared? If so, does the sharing have to be synchronous (for instance, viewing the data at once) or asynchronous (for example, two people authoring a report taking turns to edit it)? Other factors include the physical location of fellow team members, such as collaborators communicating across great distances.

The third aspect is the *organizational environment*, for example, how good is user support likely to be, how easily can it be obtained, and are there facilities or resources for training, how efficient or stable is the communications infrastructure, and so on?

Finally, the *technical environment* will need to be established. For example, what technologies will the product run on or need to be compatible with, and what technological limitations might be relevant?

User characteristics capture the key attributes of the intended user group, such as the users’ abilities and skills, and depending on the product, also their educational background, preferences, personal circumstances, physical or mental disabilities, and so on. In addition, a user may be a novice, an expert, a casual user, or a frequent user. This affects the ways in which interaction is designed. For example, a novice user may prefer step-by-step guidance. An expert, on the other hand, may prefer a flexible interaction with more wide-ranging powers of control. The collection of characteristics for a typical user is called a *user profile*. Any one product may have several different user profiles.

Usability goals and *user experience goals* are another kind of requirement, and they should be captured together with appropriate measures. Chapter 2 briefly introduced

usability engineering, an approach in which specific measures for the usability goals of the product are agreed upon early in the development process and are used to track progress as development proceeds. This both ensures that usability is given due priority and facilitates progress tracking. The same is true for user experience goals. Although it is harder to identify quantifiable measures that track these qualities, an understanding of their importance is needed during the requirements activity.

Different interactive products will be associated with different requirements. For example, a telecare system designed to monitor an elderly person's movements and alert relevant care staff will be constrained by the type and size of sensors that can be easily worn by the users as they go about their normal activities. Wearable interfaces need to be light, small, fashionable, preferably hidden, and not get in the way. A desirable characteristic of both an online shopping site and a robotic companion is that they are trustworthy, but this attribute leads to different nonfunctional requirements—in the former, security of information would be a priority, while in the latter behavioral norms would indicate trustworthiness. A key requirement in many systems nowadays is that they be secure, but one of the challenges is to provide security that does not detract from the user experience. Box 11.1 introduces usable security.

BOX 11.1

Usable Security

Security is one requirement that most users and designers will agree is important, to some degree or another, for most products. The wide range of security breaches, in particular of individuals' private data, that have occurred in recent years has heightened everyone's awareness of the need to be secure. But what does this mean for interaction design, and how can security measures be suitably robust, while not detracting from the user experience? As long ago as 1999, Anne Adams and Angela Sasse (1999) discussed the need to investigate the usability of security mechanisms and to take a user-centered approach to security. This included informing users about how to choose a secure password, but it also highlighted that ignoring a user-centered perspective regarding security will result in users circumventing security mechanisms.

Many years later, usable security and the role of users in maintaining secure practices is still being discussed. Users are now bombarded with advice about how to choose a password, but most adults interact with so many systems and have to maintain a wide variety of login details and passwords that this can be overwhelming. Instead of improving security, this can lead to users developing coping strategies to manage their passwords, which may end up compromising rather than strengthening security. In their study, Elizabeth Stobert and Robert Biddle (2018) identify a password lifecycle that shows how passwords are developed, reused, adapted, discarded, and forgotten. Users are not necessarily ignoring password advice when they create weak passwords or write them down, but instead they are carefully managing their resources and expending more effort to protect the most valued accounts. Chapter 4, "Cognitive Aspects," highlighted issues around memory and passwords and the move toward using biometrics instead of passwords. The need to identify usable security requirements, however, will still exist even with biometrics. ■

ACTIVITY 11.1

Suggest some key requirements in each category (functional, data, environmental, user characteristics, usability goals, and user experience goals) for each of the following situations:

1. An interactive product for navigating around a shopping center.
2. A wearable interactive product to measure glucose levels for an individual with diabetes.

Comment

You may come up with alternative suggestions. These are merely indicative answers.

1. Interactive product for navigating around a shopping center.

Functional The product will locate places in the shopping center and provide routes for the user to reach their destination.

Data The product needs access to GPS location data for the user, maps of the shopping center, and locations of all the places in the center. It also requires knowledge about the terrain and pathways for people with different needs.

Environmental The product design needs to take into account several environmental aspects. Users may be in a rush, or they may be more relaxed and wandering about. The physical environment will be noisy and busy, and users may be talking with friends and colleagues while using the product. Support or help with using the product may not be readily available, but the user can probably ask a passerby for directions if the app fails to work.

User Characteristics Potential users are members of the population who have their own mobile device and for whom the center is accessible. This suggests quite a wide variety of users with different abilities and skills, a range of educational backgrounds and personal preferences, and different age groups.

Usability Goals The product needs to be easy to learn so that new users can use it immediately, and it should be memorable for more frequent users. Users won't want to wait around for the product to display fancy maps or provide unnecessary detail, so it needs to be efficient and safe to use; that is, it needs to be able to deal easily with user errors.

User Experience Goals Of the user experience goals listed in Chapter 1, "What Is Interaction Design?" those most likely to be relevant here are satisfying, helpful, and enhancing sociability. While some of the other goals may be appropriate, it is not essential for this product to, for example, be cognitively stimulating.

2. A wearable interactive product to measure glucose levels for an individual with diabetes.

Functional The product will be able to take small blood samples and measure glucose readings from them.

Data The product will need to measure and display the glucose reading—but possibly not store it permanently—and it may not need other data about the individual. These questions would be explored during the requirements activity.

Environmental The physical environment could be anywhere the individual may be—at home, in hospital, visiting the park, and so on. The product needs to be able to cope with a wide range of conditions and situations and to be suitable for wearing.

User Characteristics Users could be of any age, nationality, ability, and so forth, and may be novice or expert, depending on how long they have had diabetes. Most users will move rapidly from being a novice to becoming a regular user.

Usability Goals The product needs to exhibit all of the usability goals. You wouldn't want a medical product being anything other than effective, efficient, safe, easy to learn and remember how to use, and with good utility. For example, outputs from the product, especially any warning signals and displays, must be clear and unambiguous.

User Experience Goals User experience goals that are relevant here include the device being comfortable, while being aesthetically pleasing or enjoyable may help encourage continued use of the product. Making the product surprising, provocative, or challenging is to be avoided, however. ■

11.4 Data Gathering for Requirements

Data gathering for requirements covers a wide spectrum of issues, including who are the intended users, the activities in which they are currently engaged and their associated goals, the context in which the activities are performed, and the rationale for the current situation. The goals for the data gathering sessions will be to discover all of the types of requirements relevant for the product. The three data gathering techniques introduced in Chapter 8, “Data Gathering” (interviews, observation, and questionnaires), are commonly used throughout the interaction design lifecycle. In addition to these techniques, several other approaches are used to discover requirements.

For example, documentation, such as manuals, standards, or activity logs, are a good source of data about prescribed steps involved in an activity, any regulations governing a task, or where records of activity are already kept for audit or safety-related purposes. Studying documentation can also be good for gaining background information, and it doesn't involve stakeholder time. Researching other products can also help identify requirements. For example, Jens Bornschein and Gerhard Weber (2017) analyzed existing nonvisual drawing support packages to identify requirements for a digital drawing tool for blind users. Xiangping Chen et al. (2018) propose a recommender system for exploring existing app stores and extracting common user interface features to identify requirements for new systems.

It is usual for more than one data gathering technique to be used in order to provide different perspectives. Examples are observation to understand the context of the activity, interviews to target specific user groups, questionnaires to reach a wider population, and focus groups to build a consensus view. Many different combinations are used in practice, and Box 11.2 includes some examples. Note that the example from Orit Shaer et al. (2012) also illustrates the development of an interactive product for a specialist domain, where users join the development team to help them understand the domain complexities.

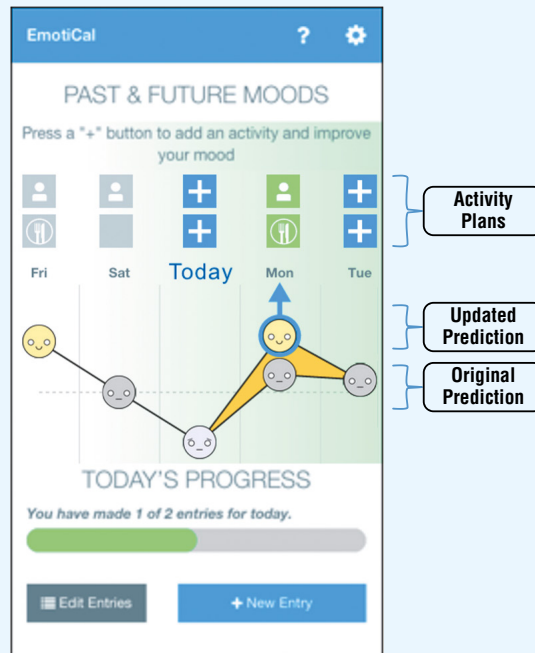
BOX 11.2

Combining Data Gathering in Requirements Activities

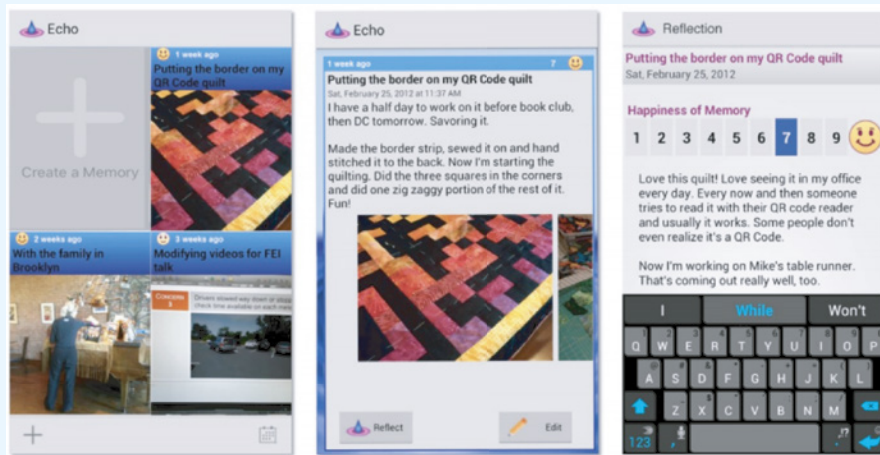
The following describes some examples where different data gathering techniques have been combined to progress requirements activities.

Direct Observation in the Field, Indirect Observation Through Log Files, Interviews, Diaries, and Surveys

Victoria Hollis et al. (2017) performed a study to inform the design of reflective systems that promote emotional well-being. Specifically, they wanted to explore the basis for future recommendations to improve a person's well-being and the effects of reflecting on negative versus positive past events. They performed two direct observation studies in the field with 165 participants. In both studies, surveys were administered before and after the field study period to assess emotional well-being, behaviors, and self-awareness. The first study also performed interviews. In the first study (60 participants, over 3 weeks), they investigated the relationship between past mood data, emotional profiles, and different types of recommendations to improve future well-being. In the second study (105 participants, over 28 days), using a smartphone diary application, they investigated the effects of reflection and analysis of past negative and positive events on well-being. Together, these studies provided insights into requirements for systems to support the promotion of emotional well-being. Figure 11.3 shows the visualization displayed to emotion-forecasting participants in week 3 of the first study. The leftmost two points in the line graph indicate average mood ratings on previous days, and the center point is the average rating for the immediate day. The two rightmost points indicate predicted mood for upcoming days.



(a)



(b)

Figure 11.3 (a) The image shown to participants in the first field study (b) the smartphone diary app for the second study

In Figure 11.3 (b) The left panel shows the home screen: Participants record a new experience by clicking the large plus sign (+) in the upper left. The center panel shows a completed event record, which consists of a header, textual entry, emotion rating, and image. The right panel shows participant reflection by rating their current emotional reaction to the initial record and providing a new textual reappraisal.

Source: (a) Hollis et al. (2017), Figure 1. Used courtesy of Taylor & Francis and (b) Hollis et al. (2017), Figure 9. Used courtesy of Taylor & Francis

Diaries and Interviews

Tero Jokela et al. (2015) studied how people currently combine multiple information devices in their everyday lives to inform the design of future interfaces, technologies, and applications that better support multidevice use. For the purpose of this study, an *information device* is any device that can be used to create or consume digital information, including personal computers, smartphones, tablets, televisions, game consoles, cameras, music players, navigation devices, and smartwatches. They collected diaries over a one-week period and interviews from 14 participants. The study indicates that requirements for the technical environment needed to improve the user experience of multiple devices, including being able to access any content with any device and improved reliability and performance for cloud storage.

Interview, Think-Aloud Evaluation of Wireframe Mock-Up, Questionnaire, and Evaluation of Working Prototype

Carole Chang et al. (2018) developed a memory aid application for traumatic brain injury (TBI) sufferers. They initially conducted interviews with 21 participants to explore memory impairments after TBI. From these, they identified common themes in the use of external memory aids. They also learned that TBI sufferers do not want just another reminder system but something that helps them to remember and hence can also train their memory, and that their technology requirements were for something simple, customizable, and discreet.

(Continued)

Studying Documentation, Evaluating Other Systems, User Observations, and Group Interviews
Nicole Costa et al. (2017) describe their ethnographic study of the design team for the user interface of a ship's maneuvering system (called a *conning display*). The design team started by studying the accident and incident reports to identify requirements for things to avoid, such as mixing up turn-rate meter with rudder indicator. They used Nielsen's heuristics to evaluate other existing systems, and specifically how to represent the vessel on the display. Once a suitable set of requirements had been discovered, sketching, prototyping, and evaluating with the help of users was used to produce the final design.

Ethnographic Study, Interviews, Usability Tests, and User Participation

Orit Shaer et al. (2012) report on the design of a multitouch tabletop user interface for collaborative exploration of genomic data. In-depth interviews were conducted with 38 molecular and computational biologists to understand the current work practices, needs, and workflow of small research groups. A small team of nine researchers investigating gene interaction in tuberculosis was studied for eight weeks using an ethnographic approach, and other labs were also observed. Because the application area was specialized, the design team needed to be comfortable with the domain concepts. To achieve this, biologists were integrated into the development team, and other members of the design team regularly visited biology research group partners, attended courses to teach them relevant domain concepts, and conducted frequent usability tests with users. ■

11.4.1 Using Probes to Engage with Users

Probes come in many forms and are an imaginative approach to data gathering. They are designed to prompt participants into action, specifically by interacting with the probe in some way, so that the researchers can learn more about users and their contexts. Probes rely on some form of logging to gather the data—either automatically in the case of technology probes or manually in the case of diaries or design probes.

The idea of a probe was developed during the Presence Project (Gaver et al., 1999), which was investigating novel interaction techniques to increase the presence of elderly people in their local community. They wanted to avoid more traditional approaches, such as questionnaires, interviews, or ethnographic studies, and developed a technique called *cultural probes*. These probes consisted of a wallet containing eight to ten postcards, about seven maps, a disposable camera, a photo album, and a media diary. Recipients were asked to answer questions associated with certain items in the wallet and then return them directly to the researchers. For example, on a map of the world, they were asked to mark places where they had been. Participants were also asked to use the camera to take photos of their home, what they were wearing today, the first person they saw that day, something desirable, and something boring.

Inspired by this original cultural probe idea, different forms of probes have been adapted and adopted for a range of purposes (Boehner et al., 2007). For example, *design probes* are objects whose form relates specifically to a particular question and context. They are intended to gently encourage users to engage with and answer the question in their own context. Figure 11.4 illustrates a Top Trumps probe; the participant was given six cards and asked to describe objects which were powerful to them and to rate the object's powers using numerical values out of 100 (Wallace et al., 2013).




Figure 11.4 Top Trumps probe

Source: Wallace et al. (2013), Figure 6. Reproduced with permission of ACM Publications

Other types of probes include *technology probes* (Hutchinson et al., 2003) and *provocative probes* (Sethu-Jones et al., 2017). Examples of technology probes include toolkits, such as the SenseBoard for developing IoT applications (Richards and Woodthorpe, 2009), mobile phone apps such as Pocketsong, a mobile music listening app (Kirk et al., 2016), and M-Kulinda, a device that uses sensors to monitor movement that was deployed in rural Kenya (Chidziwisano and Wyche, 2018). The last of these, M-Kulinda, worked with the participant's mobile phones to alert participants of unexpected movement in their homes. By doing this, the researchers hoped to provide insights into how sensor-based technology may be used in rural households and to learn about domestic security measures in rural Kenya.

Provocative probes are technology probes designed to challenge existing norms and attitudes in order to provoke discussion. For example, Dimitros Raptis et al. (2017) designed a provocation called “The Box” to challenge domestic laundry practices. The intention was to learn about users' laundry practices and also to provoke users across three dimensions: conceptually, functionally, and aesthetically. Conceptual provocation challenged the assumption that electricity is always available and that its source is not relevant. Functional provocation was provided through an emergency override button that could be pressed if the electricity had been cut off, but its size and color implied that doing so was somehow wrong. Aesthetic provocation was achieved by designing a separate physical box, rather than a mobile phone app, and by designing it to be bulky and utilitarian (see Figure 11.5). This kind of provocation was found to increase participants' engagement.



a) electricity status – 12 hour forecast,
b) savings account, c) override button presses,
d) override button, and e) electricity status at the moment

Figure 11.5 The Box provocative probe

Source: Raptis et al. (2017). Reproduced with permission of ACM Publications

11.4.2 Contextual Inquiry

Contextual inquiry was originally developed in the 1990s (Holtzblatt and Jones, 1993) and has been adapted over time to suit different technologies and the different ways in which technology fits into daily life. Contextual inquiry is the core field research process for Contextual Design (Holtzblatt and Beyer, 2017), which is a user-centered design approach that explicitly defines how to gather, interpret, and model data about how people live in order to drive design ideation. However, contextual inquiry is also used on its own to discover requirements. For example, Hyunyoung Kim et al. (2018) used contextual inquiry to learn about unresolved usability problems related to devices for continuous parameter controls, such as the knobs and sliders used by sound engineers or aircraft pilots. From their study, they identified six needs: fast interaction, precise interaction, eyes-free interaction, mobile interaction, and retro-compatibility (the need to use their existing expertise with interfaces).

One-on-one field interviews (called *contextual interviews*) are undertaken by every member of the design team, each lasting about one-and-a-half to two hours. These interviews focus on matters of daily life (work and home) that are relevant for the project scope. Contextual inquiry uses a model of master/apprentice to structure data gathering, based on the idea that the interviewer (apprentice) is immersed in the world of the user (master), creating an attitude of sharing and learning on either side. This approach shifts the perceived “power” relationship that can exist in a more traditional interviewer–interviewee relationship. Users talk as they “do,” and the apprentice learns by being part of the activity while also observing it, which has all of the advantages of observation and ethnography. Hidden and specific details that people don’t make explicit, and don’t necessarily realize themselves, emerge this way and can be shared and learned. While observing and learning, the apprentice focuses on why, not just what.

Four principles guide the contextual interview, each of which defines an aspect of the interaction and enhances the basic *apprenticeship model*. These principles are context, partnership, interpretation, and focus.

The *context principle* emphasizes the importance of going to the user, wherever they are, and seeing what they do as they do it. The benefits of this are exposure to ongoing experience instead of summary data, concrete details rather than abstract data, and experienced motives rather than reports. The *partnership principle* creates a collaborative context in which the user and interviewer can explore the user’s life together, on an equal footing. In a traditional interview or workshop situation, the interviewer or workshop leader is in control, but in contextual inquiry, the spirit of partnership means that understanding is developed together.

Interpretation turns the observations into a form that can be the basis of a design hypothesis or idea. These interpretations are developed collaboratively by the user and the design team member to make sure that they are sound. For example, imagine that during a contextual interview for an exercise monitor, the user repeatedly checks the data, specifically looking at the heart rate display. One interpretation of this is that the user is very worried about their heart rate. Another interpretation is that the user is concerned that the device is not measuring the heart rate effectively. Yet another interpretation might be that the device has failed to upload data recently, and the user wants to make sure that the data is saved regularly. The only way to make sure that the chosen interpretation is correct is to ask the user and see their reaction. It may be that, in fact, they don’t realize that they are doing this and that it has simply become a distracting habit.

The fourth principle, *focus*, is established to guide the interview setup and tells the interviewer what they need to pay attention to among all of the detail that will be unearthed. While the apprenticeship model means that the master (user) will choose what to share or teach, it is also the apprentice's responsibility to capture information relevant to the project. In addition, the interviewer will have their own interests and perspectives, and this allows different aspects of the activity to surface when all members of the team conduct interviews around the project focus. This leads to a richer collection of data.

Together with the principles that shape how the session will run, the contextual interview is also guided by a set of “cool concepts.” *Cool concepts* are an addition to the original contextual inquiry idea, and they are derived from a field study that investigated what it is about technologies that users find “cool” (Holtzblatt and Beyer, 2017, p10). Seven cool concepts emerged from this study, and they are divided into two groups: four concepts that enhance the joy of life and three concepts that enhance the joy of use.

The *joy of life concepts* capture how products make our lives richer and more fulfilling. These concepts are *accomplish* (empower users), *connection* (enhance real relationships), *identity* (support users' sense of self), and *sensation* (pleasurable moments).

The *joy of use concepts* describes the impact of using the product itself; they are: *direct in action* (provide fulfillment of intent), *the hassle factor* (remove all glitches and inconveniences), and *the learning delta* (reduce the time to learn). During a contextual interview, cool concepts are identified as the user does their activity, although often they only emerge retrospectively when reflecting on the session.

The contextual interview has four parts: getting an overview, the transition, main interview, and wrap-up. The first part can be performed like a traditional interview, introducing each other and setting the context of the project. The second part is where the interaction changes as both parties get to know each other, and the nature of the contextual interview engagement is set up. The third part is the core data gathering session when the user continues with their activities and the interviewer observes them and learns. Finally, the wrap-up involves sharing some of the patterns and observations the interviewer has made.

During the interview, data is collected in the form of notes and initial Contextual Design models and perhaps audio and video recordings. Following each contextual interview, the team holds an interpretation session that allows the whole team to talk about the user and hence establish a shared understanding based on the data. During this session, specific contextual design models are also generated or consolidated. There are 10 models suggested by Contextual Design, and the team can choose the most relevant for the project. Five of these models are linked to the cool concepts: the *day-in-the-life model* (representing accomplishment), the *relationship* and *collaboration models* (representing connection), the *identity model*, and the *sensation board*. Five others provide a complete view of the users' tasks, but they are used only for some projects: the *flow model*, the *decision point model*, the *physical model*, the *sequence model*, and the *artifact model*. The affinity diagram, described in Chapter 2, is produced after several interpretation sessions have taken place. The contextual design method follows this up with an immersion exercise called the *Wall Walk*, in which all of the generated models are hung up on the walls of a large conference room for stakeholders to read and suggest design ideas. For more detail about these models and how to generate them, see Holtzblatt and Beyer (2017).

ACTIVITY 11.2

How does contextual inquiry compare with the data gathering techniques introduced in Chapter 8, specifically ethnography and interviews?

Comment

Ethnography involves observing a situation without any *a priori* structure or framework; it may include other data gathering techniques such as interviews. Contextual inquiry also involves observation with interviews, but it provides more structure, support, and guidance in the form of the apprenticeship model, the principles to which one must adhere, the cool concepts to look out for, and a set of models to shape and present the data. Contextual inquiry also explicitly states that it is a team effort, and that all members of the design team conduct contextual interviews.

Structured, unstructured, and semi-structured interviews were introduced in Chapter 8. Contextual inquiry could be viewed as a form of unstructured interview with props, but it has other characteristics as discussed earlier, which gives it added structure and focus. A contextual inquiry interview requires the interviewee to be going about their daily activities, which may also mean that the interview moves around—something very unusual for a standard interview. ■

11.4.3 Brainstorming for Innovation

Requirements may emerge directly from the data gathered, but they may also involve innovation. *Brainstorming* is a generic technique used to generate, refine, and develop ideas. It is widely used in interaction design specifically for generating alternative designs or for suggesting new and better ideas to support users.

Various rules have been suggested for making a brainstorming session successful, some of which are listed next. In the context of the requirements activity, two key success factors are that the participants know the users and that no ideas are criticized or debated. Other suggestions for successful requirements brainstorming sessions are as follows (Robertson and Robertson, 2013; Kelley with Littman, 2016):

1. *Include participants from a wide range of disciplines with a broad range of experience.*
2. *Don't ban silly stuff.* Wild ideas often turn into really useful requirements.
3. *Use catalysts for further inspiration.* Build one idea on top of another, jump back to an earlier idea, or consider alternative interpretations when energy levels start to flag. If you get stuck, use a word pulled randomly from a dictionary to prompt ideas related to the product.
4. *Keep records.* Capture every idea, without censoring. One suggestion is to number ideas so that they can be referred to more easily at a later stage. Cover the walls and tables in paper, and encourage participants to sketch, mind-map, and diagram ideas, including keeping the flow of ideas, as spatial memory is very strong, and this can facilitate recall. Sticky notes, each with one idea, are useful for re-arranging and grouping ideas.

5. *Sharpen the focus.* Start the brainstorm with a well-honed problem. This will get the brainstorm off to a good start, and it makes it easier to pull people back to the main topic if the session wanders.
6. *Use warm-up exercises and make the session fun.* The group will require warming up if they haven't worked together before, most of the group doesn't brainstorm regularly, or the group is distracted by other pressures. Warm-up exercises might take the form of word games or the exploration of physical items related or unrelated to the problem at hand, such as the TechBox in Chapter 2.

11.5 Bringing Requirements to Life: Personas and Scenarios

Using a format such as those shown in Figure 11.1 or user stories captures the essence of a requirement, but neither of them is sufficient on their own to express and communicate the product's purpose and vision. Both can be augmented with prototypes, working systems, screenshots, conversations, acceptance criteria, diagrams, documentation, and so on. Which of these augmentations is required, and how much, will be determined by the kind of system under development. In some cases, capturing different aspects of the intended product in more formal or structured representations is appropriate. For example, when developing safety-critical devices, the functionality, user interface, and interaction of the system need to be specified unambiguously and precisely. Sapna Jaidka et al. (2017) suggest using the Z formal notation (a mathematically based specification language) and petri nets (a notation for modeling distributed systems based on directed graphs) to model the interaction behavior of medical infusion pumps. Harold Thimbleby (2015) points out that using a formal expression of requirements for number entry user interfaces such as calculators, spreadsheets, and medical devices could avoid bugs and inconsistencies.

Two techniques that are commonly used to augment the basic requirements information and to bring requirements to life are personas and scenarios. Often used together, they complement each other in order to bring realistic detail that allows the developer to explore the user's current activities, future use of new products, and futuristic visions of new technologies. They can also guide development throughout the product lifecycle.

11.5.1 Personas

Personas (Cooper, 1999) are rich descriptions of typical users of the product under development on which the designers can focus and for which they can design products. They don't describe specific people, but rather they are realistic, and not idealized. Any one persona represents a synthesis of a number of real users who have been involved in data gathering, and it is based on a set of user profiles. Each persona is characterized by a unique set of goals relating to the particular product under development, rather than a job description or a simple demographic. This is because goals often differ among people within the same job role or the same demographic.

In addition to their goals, a persona will include a description of the user's behavior, attitudes, activities, and environment. These items are all specified in some detail. For instance, instead of describing someone simply as a competent sailor, the persona includes

that they have completed a Day Skipper qualification, have more than 100 hours of sailing experience in and around European waters, and get irritated by other sailors who don't follow the navigation rules. Each persona has a name, often a photograph, and some personal details such as what they do as a hobby. It is the addition of precise, credible details that helps designers to see the personas as real potential users, and hence as people for whom they can design.

A product will usually require a small set of personas rather than just one. It may be helpful to choose a few, or maybe only one, *primary personas* who represent a large section of the intended user group. Personas are used widely, and they have proved to be a powerful way to communicate user characteristics and goals to designers and developers.

A good persona helps the designer understand whether a particular design decision will help or hinder their users. To this end, a persona has two goals (Caddick and Cable, 2011):

- To help the designer make design decisions
- To remind the team that real people will be using the product

A good persona is one that supports the kind of reasoning that says, “What would Bill (persona 1) do in this situation with the product?” and “How would Clara (persona 2) respond if the product behaved this way?” But good personas can be challenging to develop. The kind of information they include needs to be pertinent to the product being developed. For example, personas for a shared travel organizer would focus on travel-related behavior and attitudes rather than the newspapers the personas read or where they buy their clothes. On the other hand, personas for a shopping center navigation system might consider these aspects.

Steven Kerr et al. (2014) conducted a series of studies to identify user needs and goals in the kitchen as a way to improve the design of technology to assist with cooking. They conducted observations and interviews with three members each of three user groups, identified through background research: beginners, older experts, and families (specifically parents). The interview focused on topics such as cooking experience, meal planning, and grocery shopping. Two researchers attended each home visit. Notes, video, audio, and photographs were used to capture the data, including a think-aloud session during some of the activities.

Personas were developed following both inductive and deductive analysis (see Chapter 9, “Data Analysis, Interpretation, and Presentation”), looking for patterns in the data and commonalities that could be grouped into one persona. Three primary and three secondary personas were developed from the data. Figure 11.6 shows one primary (beginner) persona and one secondary (older expert) persona that they derived from this work. Note that the two types (primary and secondary) in this case have different formats, with Ben being more detailed than Olive.

They also conducted a survey online to validate the personas and to create a list of requirements for new technology to support cooking.

The style of personas varies widely, but commonly they include a name and photograph, plus key goals, user quotes, behaviors, and some background information. The examples in Figure 11.7(a) illustrate the persona format developed and used by the company in Box 11.3, which has been tailored for their purposes. Figure 11.7(b) shows the user journey associated with their personas.

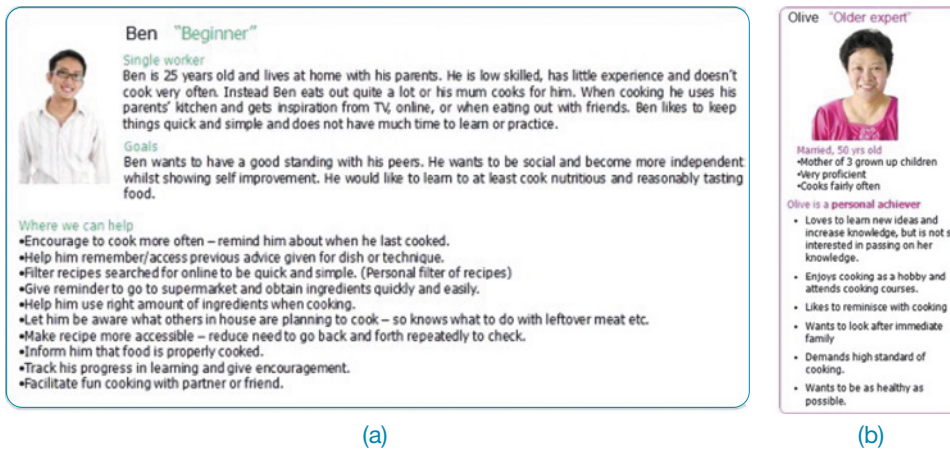


Figure 11.6 (a) One primary (beginner) persona and (b) one secondary (older expert) persona for cooking in Singapore

Source: Kerr et al. (2014). Used courtesy of Elsevier

BOX 11.3

Persona-Driven Development in London

Caplin Systems is based in the City of London and provides a framework to investment banks that enables them to build, or enhance, their single-dealer offering (a platform that integrates services and information for trading in capital markets) quickly or to create a single-dealer platform for the first time.

The company was drawn to use personas to increase the customer focus of their products by better understanding for whom they were developing their system. Personas were seen as a way to provide a unified view of their users and to start building more customer-focused products.

The first step was to run a workshop for the whole company, introduce personas, show how other companies were using them, and have employees experience the benefits of using personas firsthand through some simple team exercises. The following proposition was then put forward:

Should we adopt personas and persona-driven development?

The response was a resounding “yes!” This was a good thing to do. Gaining this “buy in” was fundamentally important in ensuring that everyone was behind the use of personas and committed to the change.

Everyone got excited, and work began to define the way forward. Further workshops were run to refine the first persona, though in hindsight the Caplin team believes that too long a time was spent trying to get the first persona perfect. Now they are much more agile about persona creation.

Eighteen months after the persona breakthrough workshop, the main persona for Caplin Trader, Jack, and his “pain points” were the focus of development, design decisions, and team discussions. Ongoing persona development focused on end users of the software built with Caplin’s technology, and Narrative Journey Maps captured their interactions and helped to define goals/motivations and pain points (see Figure 11.7b).



(a)



(b)

Figure 11.7 (a) Example personas, (b) the narrative journey maps—sad faces show pain points for the persona

Source: Caplin Systems ■

ACTIVITY 11.3

Develop two personas for a group travel organizer app that supports a group of people, perhaps a family, who are exploring vacation possibilities together. Use the common persona structure of a photo, name, plus key goals, user quotes, behaviors, and some background information. Personas are based on real people, so choose friends or relatives that you know well to construct them.

These can be drawn by hand, or they can be developed in PowerPoint, for example. There are also several tailorable persona templates available on the Internet that can be used instead.

Comment

The personas shown in Figure 11.8 were developed for a father and his daughter using templates from <https://xtensio.com/templates/>.

Family traveler



"I want a travel organiser that will offer me a range of potential vacations that suit our needs"

Age: 35
 Work: Plumber
 Family: Married, two children

Personality



- Organised
- Practical
- Expects high standard

Goals

- To book comprehensive travel quickly
- To find a trip that meets the needs of the whole family
- To feel supported and guided from the beginning of the booking experience right to the end.

Frustrations

- Wasting time filling in forms
- Too much irrelevant information
- Existing systems tend to be too diverse and complicated


Bio

Will loves to take his family on adventure holidays to explore new challenges. His children, Sky (8) and Eamonn (15) are old enough to take part in several sporting activities and he wants to make the most of this before they no longer want to go on trips with him and his wife, Claire. He likes the fact that choosing travel options is so much easier than it used to be, but is frustrated by the many different sources and disjointed options that this can result in. He wants a travel organiser that can provide clear support for family holidays while offering as wide a choice as possible.


Motivation



Favourite destinations



Young traveler



"I want a travel organiser that will allow all of us to choose the vacation together"

Age: 8
 Work: Schoolgirl
 Family: Mum Dad and Eamonn (15)

Personality



- Energetic
- Inquisitive
- Likes reading

Goals

- To find a good vacation without any fuss
- To find a destination with other children her age
- To make sure that the travel time is short


Frustrations

- Sitting around discussing things for too long
- Not getting clear answers to her questions
- Feeling that everything is organised for adults and not children her age

Bio

Sky likes having adventures. She is very energetic and takes part in lots of sporting activities at school, such as gymnastics and swimming. She enjoys playing games with her older brother, Eamonn. Sky is keen to make new friends, but is also happy sitting reading a book, painting or making a model. She likes going to visit new places but expects to see something familiar, such as playground or food that she recognises! The most important thing for her is that she can go on vacation with her family where there will be something for everyone to do - but especially for her and Eamonn.

Motivation



Favourite destinations




Figure 11.8 Two personas for the group travel organizer ■

This article by Jared Spool explains why personas on their own are not enough and why scenarios also need to be developed:

<https://medium.com/user-interface-22/when-it-comes-to-personas-the-real-value-is-in-the-scenarios-4405722dd55c>

11.5.2 Scenarios

A *scenario* is an “informal narrative description” (Carroll, 2000). It describes human activities or tasks in a story that allows exploration and discussion of contexts, needs, and requirements. It does not necessarily describe the use of software or other technological support used to achieve a goal. Using the vocabulary and phrasing of users means that scenarios can be understood by stakeholders, and they are able to participate fully in development.

Imagine that you have been asked to investigate how a design team working on a large building project shares information. This kind of team includes several roles, such as an architect, mechanical engineer, client, quantity surveyor, and electrical engineer. On arrival, you are greeted by Daniel, the architect, who starts by saying something like the following:

Every member of the design team needs to understand the overall purpose, but we each take a different perspective on the design decisions that have to be made. For example, the quantity surveyor will keep an eye on how much things cost, the mechanical engineer will want to make sure that the design accounts for ventilation systems, and so on. When the architect presents a design concept, such as a spiral staircase, each of us will view that concept from our own discipline and assess whether it will work as envisioned in the given location. This means that we need to share information about the project goals, the reason for decisions, and the overall budget, as well as drawing on our own discipline expertise to advise the client on options and consequences.

Telling stories is a natural way for people to explain what they are doing, and stakeholders can easily relate to them. The focus of such stories is also naturally likely to be about what the users are trying to achieve, that is, their goals. Understanding why people do things as they do and what they are trying to achieve in the process focuses the study on human activity rather than interaction with technology. Starting with current behavior allows the designer to identify the stakeholders and artifacts involved in an activity. Repeated reference to a particular app, drawing, behavior, or location indicates that it is somehow central to the activity being performed and that it deserves close attention to uncover the role it plays. Understanding current behavior also allows the designer to explore the constraints, contexts, irritations, facilitators, and so on, under which people operate. Steven Kerr et al. (2014), who devised the cooking persona Ben introduced in the previous section, also produced a scenario for Ben, which is shown in Figure 11.9. Reflect on the persona and read the scenario alongside to see how the two complement each other to give a fuller picture of the activities related to cooking in which a user such as Ben would engage.

Ben is out with friends, catching up over a meal. A friend asks if he has ever made the dish they are eating. This reminds Ben that he has not cooked for a while.

Later in the week, Ben sees a TV programme on cooking and again he is reminded that he has not cooked for a while. Thus he decides to cook later that week. Ben goes online and Google's "mee soto". He looks through the various sites for pictures that not only look good, but also have few steps and a short time to cook. He spends some time browsing through the sites for other ideas to use in the future.

Next day at work, he remembers to go to the supermarket on his way home. As he is only getting a few ingredients (for a simple recipe), he just remembers what he needs to get. He buys ingredients in whatever size is readily available and is not too concerned about the freshness. Whilst in the shop, he also makes some spontaneous purchases.

On the day he is cooking, Ben checks out a YouTube video on preparing chicken pieces. When he gets home he prints out a recipe and asks his mum for some last minute advice. He had asked her previously about a similar question, but was so long ago that he has forgotten her advice. His mum looks at the recipe and suggests some alterations to it and writes notes on the recipe.

Ben places the recipe near the stove and follows it closely in terms of steps (one process at a time, no preparation) though haphazardly gauges the amount of ingredients to put in. He tries to gauge one portion's worth of ingredients (the recipe is for 3 people), only using some of the chicken, so puts rest back in fridge (he thinks his mum might use it later). He will estimate the time the mee soto needs for cooking, tasting it when he thinks it is ready. He is initially worried about hot oil splashing on his face, so he is hesitant when handling the hot pan.

When ready, he serves the dish and finds out he has made a bit too much. Depending on how much leftover is available, he will either put it in the fridge for later or it will just be wasted.

Afterwards he loads up pictures of the dish on Facebook. He is greatly encouraged when people 'like' his link or leave a comment and it makes him feel good.

Figure 11.9 Scenario developed for Ben, the Beginner persona, for cooking in Singapore

Source: Kerr et al. (2014), Table 1. Used courtesy of Elsevier

The previous scenario describes existing behavior, but scenarios can also be used to describe behavior with a potential new technology. For example, a scenario that might be generated by potential users of a new navigation app for a large shopping center is given here:

Charlie wants to take his elderly mother, Freia, to his favorite home products store, ComfortAtHome. He knows that the store has moved within the shopping center, but he doesn't know where. He also needs to find a route that is suitable for his mother who uses a walker but doesn't like elevators. He opens the navigation app on his smartphone and enters the name of the store in the search feature. Two different branches of the store are listed, and Charlie asks for directions to the one nearest to their current location. A map of the shopping center is displayed, showing their current location, the location of the nearest store, and the suggested route. This route, however, includes a series of steps that are unsuitable for his mother. So, he asks for an alternative route that uses only ramps, which the app displays. They set off, following the new route provided.

Note the following in this limited scenario: the provision of a search facility to find the store's location (but what if the user doesn't know the store name or wants to locate all home products shops), the facility to display a map, and the importance of different options for the navigation routes to accommodate different users. These are all indicators of potential design choices for the new system. The scenario also describes one use of the app—to find a route to the nearest branch of a specific store.

During the requirements activity, scenarios emphasize the context, the usability and user experience goals, and the activities in which the user is engaged. Scenarios are often generated during workshop, interview, or brainstorming sessions to help explain or discuss some aspect of the user's goals. They capture only one perspective, perhaps a single use of the product, or one example of how a goal might be achieved.

The following scenario for the group travel organizer introduced in Activity 11.3 describes how one function of the system might work—to identify potential vacation options. This scenario incorporates information about the two personas shown in Figure 11.8. This is the kind of story that you might glean from a requirements interview:

The Thomson family enjoys outdoor activities and wants to try their hand at sailing this year. There are four family members: Sky (8 years old), Eamonn (15), Claire (32), and Will (35).

One evening after dinner, they decide to start exploring the possibilities. They want to discuss the options together, but Claire has to visit her elderly mother so she will be joining the conversation from her mother's house down the road. As a starting point, Will raises an idea they had been discussing over dinner—a sailing trip for four novices in the Mediterranean.

The system allows users to log in from different locations using different devices so that all members of the family can interact easily and comfortably with it wherever they are. The system's initial suggestion is a flotilla, where several crews (with various levels of experience) sail together on separate boats.

Sky and Eamonn aren't very happy at the idea of going on vacation with a group of other people, even though the Thomsons would have their own boat. The travel organizer shows them descriptions of flotilla experiences from other children their ages, and they are all very positive, so eventually, everyone agrees to explore flotilla opportunities.

Will confirms this recommendation and asks for detailed options. As it's getting late, he asks for the details to be saved so that everyone can consider them tomorrow. The travel organizer emails them a summary of the different options available.

Developing this type of scenario, which focuses on how a new product may be used, helps to uncover implicit assumptions, expectations, and situations in which the users might find themselves, such as the need to plan travel when participants are situated in different locations. These in turn translate into requirements, in this case an environment requirement, which may be expressed as follows:

As a <group traveler>, I want <to be able to share vacation discussions when I am not co-located> so that <the whole group can discuss their choices together under a wide range of circumstances>.

Futuristic scenarios describe an envisioned situation in the future, perhaps with a new technology and new world view. Different kinds of future visions were discussed in Chapter 3, “Conceptualizing Interaction,” and one approach that is an extension of the scenario idea and that can be used to discover requirements is design fiction (see Box 11.4).

BOX 11.4

Design Fiction

Design fiction is a way to communicate a vision about the world in which a future technology is situated. It has become popular in interaction design as a way to explore envisioned technologies and their uses without having to grapple with pragmatic challenges. In a fictional world, ethics, emotions, and context can be explored in detail and in depth without worrying about concrete constraints or implementations. The term was first coined by Bruce Sterling in 2005, and its use has gradually increased as different ways of using it have emerged.

For example, Richmond Wong et al. (2017) took a design fiction approach to engage in issues of privacy and surveillance around futuristic sensing technologies. Their design fictions were inspired by a near-future science-fiction novel, *The Circle* by David Eggers (2013). Their design fictions are visual, and they take the form of a workbook containing conceptual designs. They draw on three technologies in the novel, such as *SeeChange*, a small camera about the size of a lollipop that wirelessly records and broadcasts live video. They also include a new prototyped technology designed to detect a user's breathing pattern and heart rate (Adib et al., 2015).

The design fictions go through three rounds. The first round adapted the technologies from the novel, for example, by adding concrete interfaces. As there were no photos in *The Circle*, the authors were able to design interfaces for these technologies based only on the textual descriptions. The second round considered privacy concerns and placed the technologies in an extended world from the novel. The third round considered privacy concerns in situations that went beyond the novel and the designs they had created up to that point in time. They suggested that their design fictions could help broaden the design space for people designing sensing technologies or be used as interview probes in further research. They also reflect that an existing fictional world is a good starting point from which to develop design fictions and this helps to explore futures that might otherwise go unnoticed.

Other examples of design fiction include Eric Baumer et al.'s (2018) consideration of how design fiction can support the exploration of ethics and Steve North's (2017) approach to embrace the perspective of a nonhuman user—in this case a horse.

What's the difference between scenarios and design fiction? Mark Blythe (2017) uses the "basic plots" of literature to suggest that scenarios employ the plot of "overcoming the monster," where the monster is some problem to be solved, while design fiction more frequently takes the form of a "voyage and return" or a "quest." ■

ACTIVITY 11.4

This activity illustrates how a scenario of an existing activity can help identify requirements for a future application to support the same user goal.

Write a scenario of how you would go about choosing a new hybrid car. This should be a new car, not a secondhand car. Having written it, think about the important aspects of the task, your priorities and preferences. Then imagine a new interactive product that supports this goal and takes account of these issues. Write a futuristic scenario showing how this product would support you.

(Continued)

Comment

The following example is a fairly generic view of this process. Your scenario will be different, but you may have identified similar concerns and priorities.

The first thing I would do is to observe cars on the road, specifically hybrid ones, and identify those whose looks I find appealing. This may take several weeks. I would also try to identify any consumer reports that include an assessment of hybrid cars' performance. Hopefully, these initial activities will result in identifying a likely car for me to buy.

The next stage will be to visit a car showroom and see at first hand what the car looks like and how comfortable it is to sit in. If I still feel positive about the car, then I'll ask for a test drive. Even a short test drive helps me to understand how well the car handles, if the engine is noisy, how smooth are the gear changes, and so on. Once I've driven the car myself, I can usually tell whether I would like to own it or not.

From this scenario, it seems that there are broadly two stages involved in the task: researching the different cars available and gaining firsthand experience of potential purchases. In the former, observing cars on the road and getting expert evaluations of them are highlighted. In the latter, the test drive has quite a lot of significance.

For many people who are in the process of buying a new car, the smell and touch of the car's exterior and interior and the driving experience itself are the most influential factors in choosing a particular model. Other attributes such as fuel consumption, interior roominess, colors available, and price may rule out certain makes and models, but at the end of the day, cars are often chosen according to how easy they are to handle and how comfortable they are inside. This makes the test-drive a vital part of the process of choosing a new car.

Taking these comments into account, we've come up with the following scenario describing how an innovative "one-stop shop" for new cars might operate. This product makes use of immersive virtual reality technology that is already in use by other applications, such as designing buildings and training bomb disposal experts.

I want to buy a hybrid car, so I go down the street to the local "one-stop car shop." The shop has a number of booths in it, and when I enter, I'm directed to an empty booth. Inside, there's a large seat that reminds me of a racing car seat, and in front of that there's a large display screen.

As I sit down, the display screen jumps to life. It offers me the option of browsing through video clips of new cars that have been released in the last two years, or of searching through video clips of cars by make, model, or year. I can choose as many of these as I like. I also have the option of searching through consumer reports that have been produced about the cars in which I'm interested.

I spend about an hour looking through materials and deciding that I'd like to experience the up-to-date models of a couple of vehicles that look promising. Of course, I can go away and come back later, but I'd like to have a go right now at some of the cars I've found. By flicking a switch in my armrest, I can call up the options for virtual reality simulations for any of the cars in which

I'm interested. These are really great, as they allow me to take the car for a test drive, simulating everything about the driving experience in this car—from road holding to windshield display and foot pedal pressure to dashboard layout. It even re-creates the atmosphere of being inside the car.

Note that the product includes support for the two research activities mentioned in the original scenario, as well as the important test-drive facility. This would be only a first-cut scenario, which would then be refined through discussion and further investigation. ■

Scenarios may be constructed as textual descriptions, like those shown previously, but they can also use audio or video. For example, Alen Keirnan et al. (2015) used animated scenarios to present early user research findings about wearable emergency alarms for older people. They found that the animated scenarios helped participants to describe problems associated with the use of emergency alarm technology and to discuss and evaluate key emotions and themes (see Figure 11.10 and the following links to videos).



Figure 11.10 Screen captures of the animated scenarios used to explore early user research insights regarding emergency alarm technology for elderly people: I forgot, Dress Code, and Cow Bell

Source: Keirnan et al. (2015), Figure 1. Reproduced with permission of ACM Publications

Here are links to three animated scenarios: “I forgot,” “Dresscode,” and “Cowbell.”

<https://vimeo.com/126443388>

<https://vimeo.com/136821334>

<https://vimeo.com/123466330>

BOX 11.5

Scenarios and Personas

Writing personas and scenarios can be difficult at first, leading to a set of narratives that conflate details of the person with details of the scenario. A scenario describes one use of a product or one example of achieving a goal, while a persona characterizes a typical user of the product. Figure 11.11 captures this difference graphically.

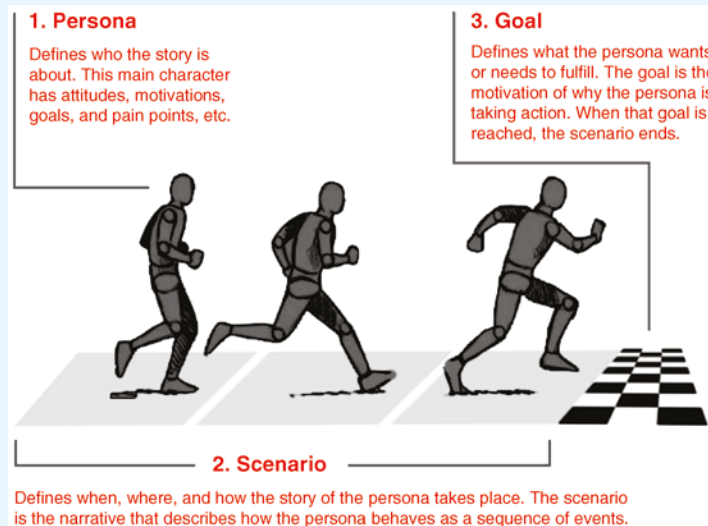


Figure 11.11 The relationship between a scenario and its associated persona

Source: <http://www.smashingmagazine.com/2014/08/06/a-closer-look-at-personas-part-1/>. Reproduced with permission of Smashing Magazine

This image also shows that the scenarios and personas are tightly linked. Each scenario represents a single experience of using the product from the perspective of one persona. Note that this figure introduces the notion of a scenario goal. Thinking about the persona’s goal for the scenario helps to scope the scenario to one use of the product. ■

11.6 Capturing Interaction with Use Cases

Use cases focus on functional requirements and capture interaction. Because they focus on the interaction between a user and the product, they may be used in design to think about the new interaction being designed, but they may also be used to capture requirements—to think through details about what the user needs to see, to know about, or to react to. Use cases define a specific process because they are a step-by-step description. This is in contrast to a *user story*, which focuses on outcomes and user goals. Nonetheless, capturing the detail of this interaction in terms of steps is useful as a way to enhance the basic requirement statement. The style of use cases varies. Two styles are shown in this section.

The first style focuses on the division of tasks between the product and the user. For example, Figure 11.12 illustrates an example of this kind of use case, focusing on the visa requirements element of the group travel application. Note that nothing is said about how the user and product might interact, but instead it focuses on user intentions and product responsibilities. For example, the second user intention simply states that the user supplies the required information, which could be achieved in a variety of ways including scanning a passport, accessing a database of personal information based on fingerprint recognition, and so on. This style of use case has been called *essential use cases*, and they were developed by Constantine and Lockwood (1999).

| retrieveVisa | |
|--|--|
| USER INTENTION | SYSTEM RESPONSIBILITY |
| find visa requirements | request destination and nationality |
| supply required information | obtain appropriate visa information |
| obtain a personal copy of visa information | offer information in different formats |
| choose suitable format | provide information in chosen format |

Figure 11.12 An “essential use case for “retrieve visa” that focuses on how the task is split between the product and the user

The second style of use cases is more detailed, and it captures the user’s goal when interacting with the product. In this technique, the main use case describes the *normal course*, that is, the set of actions most commonly performed. Other possible sequences, called *alternative courses*, are then captured at the bottom of the use case. A use case for retrieving the visa requirements for the group travel organizer with the normal course being that information about the visa requirements is available, might be as follows:

1. The product asks for the name of the destination country.
2. The user provides the country’s name.
3. The product checks that the country is valid.
4. The product asks the user for their nationality.
5. The user provides their nationality.

6. The product checks the visa requirements of that country for a passport holder of the user's nationality.
7. The product provides the visa requirements.
8. The product asks whether the user wants to share the visa requirements on social media.
9. The user provides appropriate social media information.

Alternative Courses

4. If the country name is invalid:
 - 4.1 The product provides an error message.
 - 4.2 The product returns to step 1.
6. If the nationality is invalid:
 - 6.1 The product provides an error message.
 - 6.2 The product returns to step 4.
7. If no information about visa requirements is found:
 - 7.1 The product provides a suitable message.
 - 7.2 The product returns to step 1.

Note that the number associated with the alternative course indicates the step in the normal course that is replaced by this action or set of actions. Also note how specific the use case is about how the user and the product will interact compared with the first style.

In-Depth Activity

This activity is the first of five assignments that together go through the complete development lifecycle for an interactive product.

The goal is to design and evaluate an interactive product for booking tickets for events such as concerts, music festivals, plays, and sporting events. Most venues and events have booking websites or apps already, and there are many ticket agencies that also provide reduced tickets and exclusive options, so there are plenty of existing products to research first. Carry out the following activities to discover requirements for this product:

1. Identify and capture some user requirements for this product. This could be done in a number of ways. For example, observing friends or family using ticket agents, thinking about your own experience of purchasing tickets, studying websites for booking tickets, interviewing friends and family about their experiences, and so on.
2. Based on the information you glean about potential users, choose two different user profiles and produce one persona and one main scenario for each, capturing how the user is expected to interact with the product.
3. Using the data gathered in part 1 and your subsequent analysis, identify different kinds of requirements for the product, according to the headings introduced in section 11.3. Write up the requirements using a format similar to the atomic requirements shell shown in Figure 11.1 or in the style of user stories.

Summary

This chapter examined the requirements activity in greater detail. The data gathering techniques introduced in Chapter 8 are used here in various combinations in the requirements activity. In addition, contextual inquiry, studying documentation, and researching similar products are commonly used techniques. Personas and scenarios help to bring data and requirements to life, and in combination they can be used to explore the user experience and product functionality. Use cases and essential use cases are helpful techniques for documenting the findings from data gathering sessions.

Key Points

- A *requirement* is a statement about an intended product that specifies what it is expected to do or how it will perform.
- Articulating requirements and defining what needs to be built avoids miscommunication and supports technical developers and allows users to contribute more effectively.
- There are different kinds of requirements: functional, data, environmental (context of use), user characteristics, usability goals, and user experience goals.
- Scenarios provide a story-based narrative to explore existing behavior, potential use of new products under development, and futuristic visions of technology use.
- Personas capture characteristics of users that are relevant to the product under development, synthesized from data gathering sessions.
- Scenarios and personas together can be used throughout the product lifecycle.
- Use cases capture details about an existing or imagined interaction between users and the product.


Further Reading

HOLTZBLATT, K. and BEYER, H. (2017) *Contextual Design (second edition) Design for life*. Morgan Kaufmann. This book provides a comprehensive treatment of contextual design—design for life, cool concepts, and all of the models, techniques, principles, and underpinnings for the complete contextual design method.

COHN, M. (2004) *User Stories Applied*. Addison-Wesley. This is a practical guide to writing good user stories.

PRUITT, J. and ADLIN, T. (2006) *The Persona Lifecycle: Keeping people in mind throughout product design*. Morgan Kaufmann. This book explains how to use personas in practice—how to integrate them into a product lifecycle, stories from the field, and bright ideas, as well as many example personas. It also includes five guest chapters that place personas in the context of other product design concerns. *See also* Adlin and Pruitt (2010).

ROBERTSON, S. and ROBERTSON, J. (2013) *Mastering the Requirements Process (3rd edn)*. Pearson Education. In this book, Suzanne Robertson and James Robertson provide a very practical and useful framework for software requirements work.



INTERVIEW with Ellen Gottesdiener

Ellen is an Agile Product Coach and CEO of EBG Consulting focused on helping product and development communities produce valuable outcomes through product agility. She is the author of three acclaimed books on product discovery and requirements including *Discover to Deliver: Agile Product Planning and Analysis*. Ellen is a frequent speaker and works with clients globally. She is Producer of Boston's Agile Product Open and Director of Agile Alliance's Agile Product Management initiative. Visit these websites for resources: www.ebgconsulting.com and www.DiscoverToDeliver.com.

What are requirements?

Product requirements are needs that *must* be satisfied to achieve a goal, solve a problem, or take advantage of an opportunity. The word “requirement” literally means something that is absolutely, positively, without question, necessary. Product requirements need to be defined in sufficient detail for planning and development. But before going to that effort and expense, are you sure they are not only must-haves but also the right and relevant requirements?

To arrive at this level of certainty, stakeholders ideally start by exploring the product's *options*. An option represents a potential characteristic, facet, or quality of the product. Stakeholders, who I like to refer to as product partners, use expansive thinking to surface a range of options that

could fulfill the vision. Then they collaboratively analyze the options and collectively select options, based on value.

Every product has multiple dimensions, seven in fact. Discovering options for each of the seven product dimensions yields a comprehensive, realistic view of the product (see Figure 11.2).

You want to engage diverse stakeholders across the full product lifecycle, from birth to retirement and demise.

So how do you know who are the stakeholders?

Successful teams work hand in hand with their stakeholders as *product partners*, defining value and then actively discovering—and delivering—high-value solutions. This goes beyond feature requests and requirements documents—beyond user stories and product backlogs—beyond the push-pull of competing interests. It's a partnership where the ideas, perspectives, and experiences of three different stakeholder groups converge. The result? Product partners who collaborate to discover and deliver value.

A product partnership includes people from three realms: customer, business, and technology. Each offers a unique perspective and has its own ideas of what is valuable.

The *customer* partners represent users, buyers, and advisers—people or systems that interface with the product, choose to buy it, or influence others to buy it.

They tend to value improved productivity, heightened efficiency, greater speed, entertainment, and similar benefits.

Business partners represent the people in your organization who authorize, champion, or support the product or who provide subject-matter expertise. They find value in improving market position, complying with regulations, achieving a business case, reducing overhead costs, enhancing internal performance, and so on.

Technology partners (your delivery team, internal or third parties) design, deliver, test, and support the product or advise those who do. They may value building a high-quality product, offering smooth, continual delivery, adopting a stable architecture, and the like.

This mix of partners and perspectives is essential, no matter what kind of delivery method you adopt (agile, traditional, hybrid, or another approach). For the partnership to work, these three disparate groups must collaborate to reach their shared goal: discover and deliver value.

How do you go about identifying requirements?

Requirements discovery is highly proactive, interactive, and, well, sometimes hyperactive! You are engaged in eliciting, analyzing, specifying, prototyping, and testing. And in the best practices we've been involved in, you are constantly discovering (aka identifying) product needs. It's not a "one-and-done" activity.

Elicitation includes interviews, existing documentation study, exploratory prototypes, facilitated workshops, focus groups, observation (including apprenticing, contextual inquiry, and ethnography), surveys (and other research-based techniques), and user task analysis (including storyboarding and scenario analysis). There

are a number of specific techniques within each of these general categories, and some techniques overlap. Analyzing involves using lightweight models, often combined with specifications, which are often in the form of acceptance tests or prototypes or both.

It's not enough to get the right people together and ask the right questions. To communicate efficiently and effectively about how to deliver, product partners need a focused way to communicate and make decisions together.

What we've found in our work is that the most efficient and effective discovery mechanism is a collaborative approach called the "structured conversation." In a structured conversation, the product partners first explore possible requirements (options) for their next increment. They do this within and across the seven product dimensions. This enables product partners to collaboratively and creatively explore a range of possibilities. This expansive thinking opens up product innovation, experimentation, and mutual learning.

They then evaluate these many options in terms of value. This means having shared understanding of what value really means at that point in time. Once they have narrowed the list of options through the evaluation process, they confirm how they will verify and validate these candidate solutions with unambiguous acceptance criteria. The validation includes how to test that they delivered the right requirements, and that they achieved the anticipated value from each delivery.

How do you know when you have collected enough requirements to go on to the next step?

I often get asked by clients how I know when I have a complete set of requirements.

(Continued)

I think it's more important to ask whether you are going after the right requirements.

I characterize a “right requirement” as one that is:

1. **Just in time, just enough.** It is essential for achieving the business objectives in this time period.
2. **Realistic.** It is capable of being delivered with the available resources.
3. **Clearly and unambiguously defined.** Acceptance criteria exist that all partners understand and will use to verify and validate the product.
4. **Valuable.** It is indispensable for achieving the anticipated outcomes for the next delivery cycle.

What's the hardest thing about establishing requirements?

People. Seriously. We humans are nonlinear creatures. We are unpredictable, fickle, and (as adults) often inflexible. As requirements seekers, we swim in a stew of complex, ever-evolving human systems that interoperate as we do our requirements work.

To top that off, most products' requirements are fraught with complexity and interdependency; there are truly wicked problems, whereby the problem space overlaps with solution space. As Frederick Brooks said [in his essay *No Silver Bullet*], “the hardest single part of building a software system is deciding precisely what to build.”

You can't make those decisions without trust. And trust is not an easy thing to build.

Do you have any other tips for establishing requirements?

Employ small, tightly wound cycles of requirements-build-release. Use interactive and incremental (aka agile) practices to get feedback early and often on the smallest viable releases.

For successful requirements discovery, you need to keep the focus on value—the *why* behind the product and the value considerations of the product partners. During discovery work, some people view a specific option as a “requirement” for the next delivery cycle, whereas others consider it a “wish list” item for a future release.

Such was the case in a recent release planning workshop. The team wrestled with a particular option, questioning if it could deliver enough value to justify the cost to develop it. The product champion explained why the option was a requirement—without it the organization was at risk for regulatory noncompliance. Once the others understood this, they all agreed it would be included in the release.

In the end, requirements work is human-centric and central to successful product delivery. At the same time, the subject matter and content of product requirements is complex. Thus, requirements work is the hardest part of software and will always be.

To be successful with requirements, engineer collaboration into requirements work.

Personally, I'm excited and grateful for the growing recognition of the value of collaboration and the explosion of interest in collaborative practices in the product and software development community—because collaboration works!

Chapter 12

DESIGN, PROTOTYPING, AND CONSTRUCTION

12.1 Introduction

12.2 Prototyping

12.3 Conceptual Design

12.4 Concrete Design

12.5 Generating Prototypes

12.6 Construction

Objectives

The main goals of this chapter are to accomplish the following:

- Describe prototyping and the different types of prototyping activities.
- Enable you to produce simple prototypes from the models developed during the requirements activity.
- Enable you to produce a conceptual model for a product and justify your choices.
- Explain the use of scenarios and prototypes in design.
- Introduce both physical computing kits and software development kits and their role in construction.

12.1 Introduction

Design, prototyping, and construction fall within the Develop phase of the double diamond of design, introduced in Chapter 2, “The Process of Interaction Design,” in which solutions or concepts are created, prototyped, tested, and iterated. The final product emerges iteratively through repeated design-evaluation-redesign cycles involving users, and prototypes facilitate this process. There are two aspects to design: the conceptual part, which focuses on the idea of a product, and the concrete aspect, which focuses on the details of the design. The former involves developing a conceptual model that captures what the product will do and how it will behave, while the latter is concerned with the details of the design, such as menu

types, haptic feedback, physical widgets, and graphics. The two are intertwined, and concrete design issues will require some consideration in order to prototype ideas, and prototyping ideas will lead to an evolution of the concept.

For users to evaluate the design of an interactive product effectively, designers prototype their ideas. In the early stages of development, these prototypes may be made of paper and cardboard, or ready-made components pulled together to allow evaluation, while as the design progresses, they become more polished, tailored, and robust so that they resemble the final product.

This chapter presents the activities involved in progressing a set of requirements through the cycles of prototyping and construction. The next section explains the role and techniques of prototyping and then explores how prototypes may be used in the design process. The chapter ends by discussing physical computing and software development kits (SDKs), which provide a basis for construction.

12.2 Prototyping

It is often said that users can't tell you what they want, but when they see something and get to use it, they soon know what they don't want. *Prototyping* provides a concrete manifestation of an idea—whether it is a new product or a modification of an existing one—which allows designers to communicate their ideas and users to try them out.

12.2.1 What Is a Prototype?

A *prototype* is one manifestation of a design that allows stakeholders to interact with it and to explore its suitability. It is limited in that a prototype will usually emphasize one set of product characteristics and de-emphasize others (see Box 12.1). Prototypes take many forms, for example, a scale model of a building or a bridge, or a piece of software that crashes every few minutes. A prototype can also be a paper-based outline of a display, a collection of wires and ready-made components, a digital picture, a video simulation, a complex piece of software and hardware, or a three-dimensional mockup of a workstation.


In fact, a prototype can be anything from a paper-based storyboard to a complex piece of software, and from a cardboard mockup to a molded or pressed piece of metal. For example, when the idea for the PalmPilot (a line of palmtop computers introduced in 1992) was being developed, Jeff Hawkin (founder of the company) carved up a piece of wood about the size and shape of the device he had imagined (see Figure 12.1).

Jeff Hawkin used to carry this piece of wood around with him and pretend to enter information into it, just to see what it would be like to own such a device (Bergman and Haitani, 2000). This is an example of a simple (some might even say bizarre) prototype, but it served its purpose of simulating scenarios of use. Advances in 3D printer technologies, coupled with reduced prices, have increased their use in design. It is now common practice to take a 3D model from a software package and print a prototype. Soft toys, chocolate, dresses, and whole houses may be “printed” in this way (see Figure 12.2 and the following links).

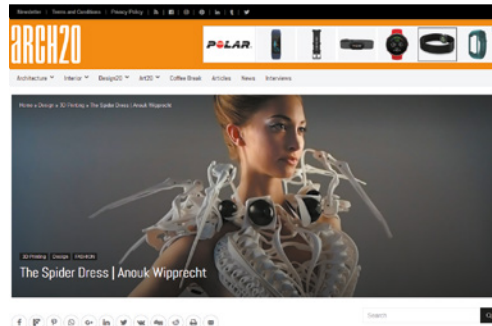


Figure 12.1 The PalmPilot wooden prototype


Source: <https://www.computerhistory.org/revolution/mobile-computing/18/321/1648>. © Mark Richards



(a)



(b)



(c)

Figure 12.2 Examples of 3D printing (a) model jet engine, (b) Spider Dress 2.0 by Anouk Wipprecht: embedded with sensors, the arms of the ‘spider’ will extend to defend the wearer if her breath becomes heavier, and (c) a teddy bear “printed” from a wireframe design

Source: (a) <https://www.thingiverse.com/thing:392115>. Licensed under CC-BY-3.0, (b) <http://www.arch2o.com>, and (c) Used courtesy of Scott Hudson

A video that shows how soft interactive objects can be printed is available at <https://www.youtube.com/watch?v=8jErWRddFYs>.

To see how 3D printing is facilitating fashion and interactive wearables, see the following articles:

<https://interestingengineering.com/high-fashion-meets-3d-printing-9-3d-printed-dresses-for-the-future>

<https://medium.com/@scientific/designing-interactive-3d-printed-things-with-tinkercad-circuit-assemblies-518ee516adb6>

12.2.2 Why Prototype?

Prototypes are useful when discussing or evaluating ideas with stakeholders; they are a communication device among team members and an effective way for designers to explore design ideas. The activity of building prototypes encourages reflection in design, as described by Donald Schön (1983), and it is recognized by designers from many disciplines as an important aspect of design.

Prototypes answer questions and support designers in choosing between alternatives. Hence, they serve a variety of purposes, for example, to test the technical feasibility of an idea, to clarify some vague requirements, to do some user testing and evaluation, or to check that a certain design direction is compatible with the rest of product development. The purpose of a prototype will influence the kind of prototype that is appropriate to build. So, for example, to clarify how users might perform a set of tasks and whether the proposed design would support them in doing this, a paper-based mockup might be produced. Figure 12.3 shows a paper-based prototype of a handheld device to help an autistic child communicate. This prototype shows the intended functions and buttons, their positioning and labeling, and the overall shape of the device, but none of the buttons actually works. This kind of prototype is sufficient to investigate scenarios of use and to decide, for example, whether the button images and labels are appropriate and the functions sufficient, but not to test whether the speech is loud enough or the response fast enough. Another example is the development of Halo, a new air quality monitor that can detect 10 different allergens and connects to an air purifier that will remove them (see the following reference). The design of Halo used a range of prototypes including many sketches (paper-based and electronic) and working prototypes.

Dan Saffer (2010) distinguishes between a product prototype and a service prototype, where the latter involves role-playing and people as an integral part of the prototype as well as the product itself. Service prototypes are sometimes captured as video scenarios and used in a similar way to the scenarios introduced in Chapter 11, “Discovering Requirements.”

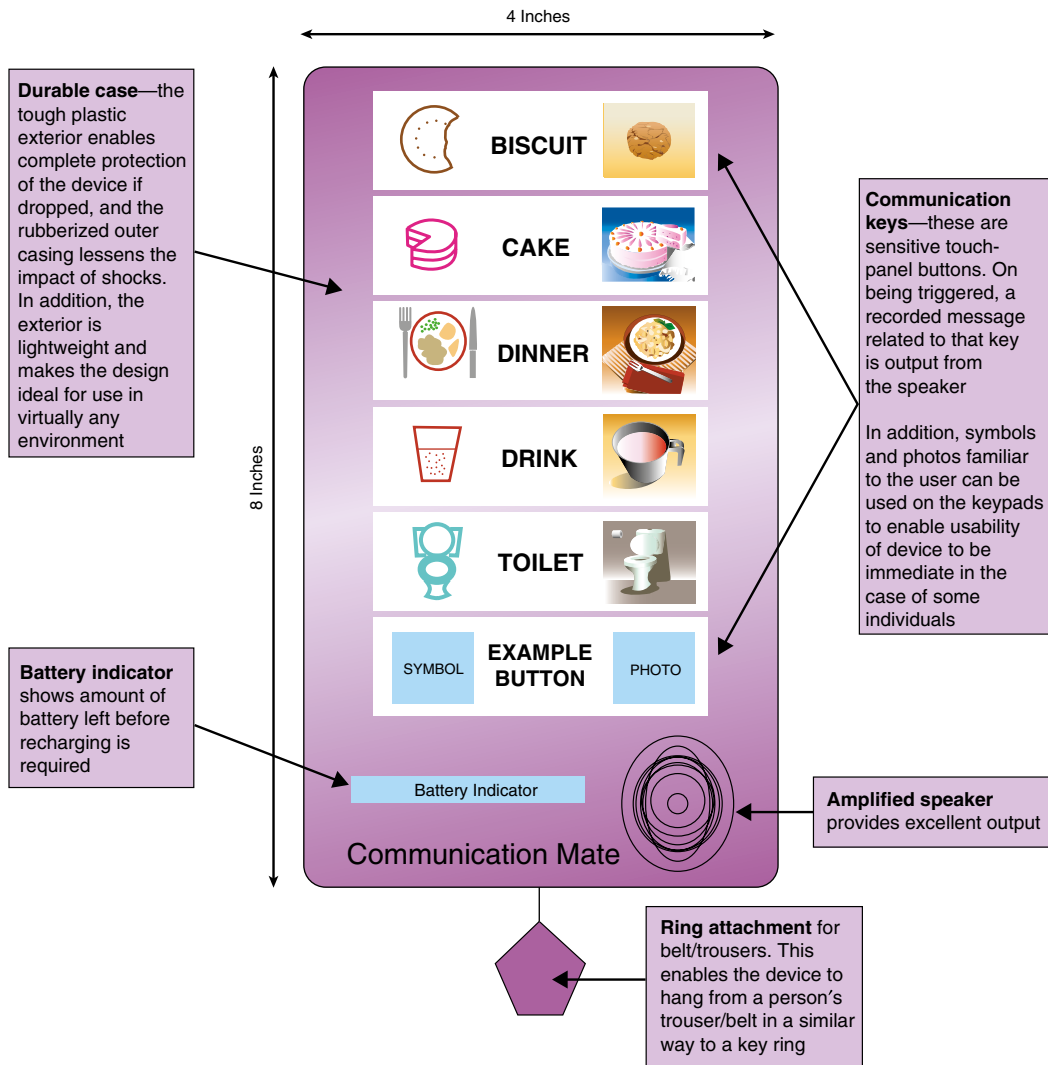


Figure 12.3 A paper-based prototype of a handheld device to support an autistic child

Source: Used courtesy of Sigil Khwaja

To see more about the development of the Wynd Halo and Wynd Home Purifier, see this website:

<https://www.kickstarter.com/projects/882633450/wynd-halo-home-purifier-keep-your-homes-air-health?ref=discovery>

12.2.3 Low-Fidelity Prototyping

A *low-fidelity prototype* does not look very much like the final product, nor does it provide the same functionality. For example, it may use very different materials, such as paper and cardboard rather than electronic screens and metal, it may perform only a limited set of functions, or it may only represent the functions and not perform any of them. The lump of wood used to prototype the PalmPilot described earlier is a low-fidelity prototype.

Low-fidelity prototypes are useful because they tend to be simple, cheap, and quick to produce. This also means that they are simple, cheap, and quick to modify so that they support the exploration of alternative designs and ideas. This is particularly important in the early stages of development, during conceptual design for example, because prototypes that are used for exploring ideas should be flexible and encourage exploration and modification. Low-fidelity prototypes are not meant to be kept and integrated into the final product.

Low-fidelity prototyping has other uses, for example in education. Seobkin Kang et al. (2018) use low-fidelity prototyping to help children represent creative ideas when designing and experimenting with complex systems. Their system, called *Rainbow*, is composed of a collection of low-fidelity materials such as paper, scissors, and markers that can be used to create prototypes, a top-down camera that can recognize the prototype, and a monitor to display augmented reality visualizations.

Storyboarding

Storyboarding is one example of low-fidelity prototyping that is often used in conjunction with scenarios, as described in Chapter 11, “Discovering Requirements.” A *storyboard* consists of a series of sketches showing how a user might progress through a task using the product under development. It can be a series of screen sketches or a series of scenes showing how a user can perform a task using an interactive device. When used in conjunction with a scenario, the storyboard provides more detail and offers stakeholders a chance to role-play with a prototype, interacting with it by stepping through the scenario. The example storyboard shown in Figure 12.4 depicts a person (Christina) using a new mobile device for exploring historical

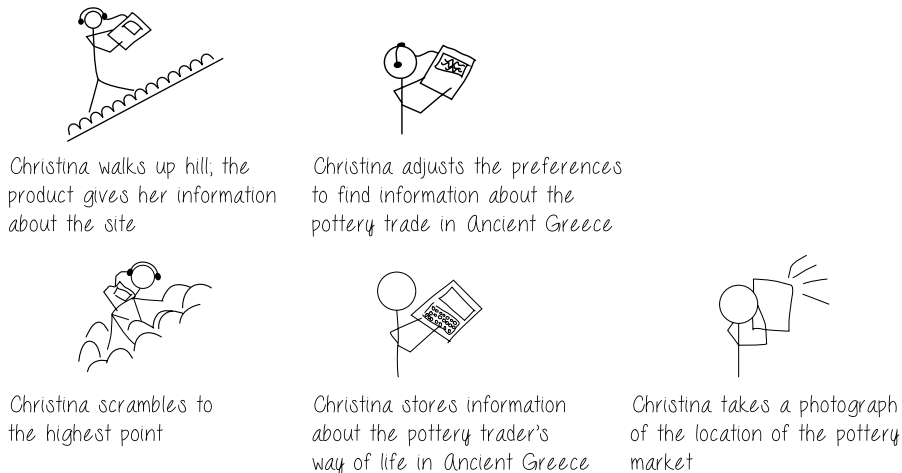


Figure 12.4 An example storyboard for a mobile device to explore ancient sites such as The Acropolis

sites. This example shows the context of use for this device and how it might support Christina in her quest for information about the pottery trade at The Acropolis in Ancient Greece.

Sketching

Low-fidelity prototyping often relies on hand-drawn sketches. Many people find it difficult to engage in *sketching* because they are inhibited by the quality of their drawing. As Saul Greenberg et al. (2012) put it, however, “Sketching is not about drawing. Rather, it is about design” (p. 7). You can get over any inhibition by devising your own symbols and icons and practicing them—referred to by Saul Greenberg et al. as a *sketching vocabulary* (p. 85). They don’t have to be anything more than simple boxes, stick figures, and stars. Elements that might be required in a storyboard sketch, for example, include digital devices, people, emotions, tables, books, and so forth, and actions such as give, find, transfer, and write. If you are sketching an interface design, then you might need to draw various icons, dialog boxes, and so on. Some simple examples are shown in Figure 12.5. The next activity requires other sketching symbols, but they can still be drawn quite simply. Mark Baskinger (2008) provides further tips for those new to sketching.




Figure 12.5 Some simple sketches for low-fidelity prototyping

Prototyping with Index Cards

Using *index cards* (small pieces of cardboard about 3 × 5 inches) is a successful and simple way to prototype an interaction, and it is used for developing a range of interactive products including websites and smartphone apps. Each card represents one element of the interaction, perhaps a screen or just an icon, menu, or dialog exchange. In user evaluations, the user can step through the cards, pretending to perform the task while interacting with the cards. A more detailed example of this kind of prototyping is provided in Section 12.5.2.

ACTIVITY 12.1

Produce a storyboard that depicts how to fill a car with fuel.

Comment

Our attempt is shown in Figure 12.6.




Figure 12.6 A storyboard depicting how to fill a car with fuel ■

Wizard of Oz

Another low-fidelity prototyping method called *Wizard of Oz* assumes that you have a software-based prototype. With this technique, the user interacts with the software as though interacting with the product. In fact, however, a human operator simulates the software's response to the user. The method takes its name from the classic story of the little girl who is swept away in a storm and finds herself in the Land of Oz (Baum and Denslow, 1900). The Wizard of Oz is a small shy man who operates a large artificial image of himself from behind a screen where no one can see him. The Wizard of Oz style of prototyping has been used successfully for various applications, including analyzing gestural behavior (Henschke et al., 2015) and when studying dialogues between children and virtual agents (Fialho and Coheur, 2015). The Wizard of Oz technique is often used in human-robot interaction studies. One such example is *Marionette*, a Wizard of Oz system for performing studies on the road with autonomous vehicles (Wang et al., 2017). Prototyping AI systems also draws on this style of prototyping, where the designer sketches the AI for themselves, and as the design matures, implementations of the AI can take its place (van Allen, 2018).

12.2.4 High-Fidelity Prototyping

A *high-fidelity prototype* looks more like the final product and usually provides more functionality than a low-fidelity prototype. For example, a prototype of a software system developed in Python or other executable language is higher fidelity than a paper-based mock-up; a

molded piece of plastic with a dummy keyboard would be a higher-fidelity prototype of the PalmPilot than the lump of wood. There is a continuum between low- and high-fidelity, and prototypes used “in the wild,” for example, will have enough fidelity to be able to answer their design questions and to learn about interaction or technological constraints or contextual factors. It is common for prototypes to evolve through various stages of fidelity, within the design-evaluate-redesign cycles. Boban Blazeovski and Jean Haslwanter (2017) describe their successful trials of two fully working prototypes for an assembly line mobile worker assistance system. They developed a smartphone app and a tablet-based app, both of which were integrated into the production system so that suitable instructions could be provided. Workers used the two versions for five days, and this allowed them to be evaluated *in situ*. The paper concludes that although producing a working prototype takes more effort, being able to try the prototype in real contexts provided valuable feedback for this kind of environment.

High-fidelity prototypes can be developed by modifying and integrating existing components—both hardware and software—which are widely available through various developer kits and open source software, for example. In robotics, this approach has been called *tinkering* (Hendriks-Jansen, 1996), while in software development it has been referred to as *Opportunistic System Development* (Ncube et al, 2008). For example, Ali Al-Humairi et al. (2018) used existing hardware (Arduino) and open source software to build a prototype to test their idea of robotically playing musical instruments automatically from a mobile phone.

12.2.5 Compromises in Prototyping

By their very nature, prototypes involve compromises: the intention is to produce something quickly to test an aspect of the product. Youn-Kiung Lim et al. (2008) suggest an anatomy of prototyping that structures the different aspects of a prototype and what it aims to achieve. Their ideas are expanded in Box 12.1. The kind of questions that any one prototype can answer is limited, and the prototype must be built with the key issues in mind. In low-fidelity prototyping, it is fairly clear that compromises have been made. For example, with a paper-based prototype, an obvious compromise is that the device doesn’t actually work. For physical prototypes or software prototypes, some of the compromises will still be fairly clear. For example, the casing may not be very robust, the response speed may be slow, the look and feel may not be finalized, or only a limited amount of functionality may be available. Box 12.2 discusses when to use low- or high-fidelity prototyping.

BOX 12.1

The Anatomy of Prototyping: Filters and Manifestations

Youn-Kyung Lim et al. (2008) propose a view of prototypes that focuses on their role as filters, for example to emphasize specific aspects of a product being explored by the prototype, and as manifestations of designs, for instance, as tools to help designers develop their design ideas through external representations.

They suggest three key principles in their view of the anatomy of prototypes:

1. *Fundamental prototyping principle*: Prototyping is an activity with the purpose of creating a manifestation that, in its simplest form, filters the qualities in which designers are interested without distorting the understanding of the whole.

(Continued)

2. *Economic principle of prototyping*: The best prototype is one that, in the simplest and the most efficient way, makes the possibilities and limitations of a design idea visible and measurable.
3. *Anatomy of prototypes*: Prototypes are filters that traverse a design space and are manifestations of design ideas that concretize and externalize conceptual ideas.

Youn-Kyung Lim et al. identify several dimensions of filtering and of manifestation that may be considered when developing a prototype, although they point out that these dimensions are not complete but provide a useful starting point for consideration of prototype development. These are shown in Table 12.1 and Table 12.2.

| Filtering dimension | Example variables |
|---------------------|---|
| Appearance | size; color; shape; margin; form; weight; texture; proportion; hardness; transparency; gradation; haptic; sound |
| Data | data size; data type (for example, number; string; media); data use; privacy type; hierarchy; organization |
| Functionality | system function; users' functionality needs |
| Interactivity | input behavior; output behavior; feedback behavior; information behavior |
| Spatial structure | arrangement of interface or information elements; relationship among interface or information elements, which can be either two- or three-dimensional, intangible or tangible, or mixed |

Table 12.1 Example variables of each filtering dimension

| Manifestation dimension | Definition | Example variables |
|-------------------------|---|---|
| Material | Medium (either visible or invisible) used to form a prototype | Physical media, for example, paper, wood, and plastic; tools for manipulating physical matters, such as a knife, scissors, pen, and sandpaper; computational prototyping tools, for instance, Python; physical computing tools, such as, Phidgets and Basic Stamps; available existing artifacts, such as a beeper to simulate a heart attack |
| Resolution | Level of detail or sophistication of what is manifested (corresponding to fidelity) | Accuracy of performance, for instance, feedback time responding to an input by a user (giving user feedback in a paper prototype is slower than in a computer-based one); appearance details; interactivity details; realistic versus faked data |
| Scope | Range of what is covered to be manifested | Level of contextualization, for example, website color scheme testing with only color scheme charts or color schemes placed in a website layout structure; book search navigation usability testing with only the book search related interface or the whole navigation interface |

Table 12.2 The definition and variables of each manifestation dimension ■

BOX 12.2

High-Fidelity and Low-Fidelity Prototypes

Table 12.3 summarizes proclaimed advantages and disadvantages of high- and low-fidelity prototyping.

| Type | Advantages | Disadvantages |
|-------------------------|---|--|
| Low-fidelity prototype | <ul style="list-style-type: none"> • Quick revision possible • More time can be spent on improving the design before starting development • Evaluates multiple design concepts • Useful communication device • Proof of concept | <ul style="list-style-type: none"> • Limited error checking • Poor detailed specification for development • Facilitator-driven • Limited usefulness for usability tests • Navigational and flow limitations |
| High-fidelity prototype | <ul style="list-style-type: none"> • (Almost) complete functionality • Fully interactive • User-driven • Clearly defines navigational scheme • Use for exploration and test • Look and feel of intended product • Serves as a “living” or evolving specification • Marketing and sales tool | <ul style="list-style-type: none"> • More resource-intensive to develop • Time-consuming to modify • Inefficient for proof-of-concept designs • Potential of being mistaken for the final product • Potential of setting inappropriate expectations |

Table 12.3 Advantages and disadvantages of low- and high-fidelity prototypes

Component kits and pattern libraries for interface components (see section 12.7 and Chapter 13, “Interaction Design in Practice”) make it quite easy to develop polished functional prototypes quickly, but there is a strong case for the value of low-fidelity prototypes, such as paper-based sketches, sticky note designs, and storyboarding to explore initial ideas. Paper prototyping, for example, is used in game design (Gibson, 2014), website development, and product design (Case Study 12.1). Both high- and low-fidelity prototypes can provide useful feedback during evaluation and design iterations, but how do you know which to choose? The advantages and disadvantages listed earlier will help, but there are other factors too. Beant Dhillon et al. (2011) found that a low-fidelity video prototype elicited comparable user feedback as a high-fidelity one, but it was quicker and cheaper to produce. Gavin Sim et al. (2016) investigated the effect of using a paper booklet versus iPads with children who were rating a game concept. They found that the children’s rating of the game was unaffected by the form factor but that they rated the paper version significantly higher for aesthetics. ■

This article covers the benefits of high- and low-fidelity prototyping and how to produce them:

<https://www.nngroup.com/articles/ux-prototype-hi-lo-fidelity/?Im=aesthetic-usability-effect&pt=article>

CASE STUDY 12.1

Paper Prototyping as a Core Tool in the Design of Telephone User Interfaces

Paper prototyping is being used by telephone and tablet companies as a core part of their design process (see Figure 12.7). Mobile devices are feature-rich. They include megapixel cameras, music players, media galleries, downloaded applications, and more. This requires designing interactions that are complex but that are clear to learn and use. Paper prototyping offers a rapid way to work through every detail of the interaction design across multiple applications.



Figure 12.7 Prototype developed for cell phone user interface

Mobile device projects involve a range of disciplines—all with their own perspective on what the product should be. A typical project may include programmers, project managers, marketing experts, commercial managers, handset manufacturers, user experience specialists, visual designers, content managers, and network specialists. Paper prototyping provides a vehicle for everyone involved to be part of the design process—considering the design from multiple angles in a collaborative way.

The case study on the id-book.com website describes the benefits of using paper prototyping from the designer's viewpoint, while considering the bigger picture of its impact across the entire project lifecycle. It starts by explaining the problem space and how paper prototyping is used as an integrated part of user interface design projects for European and U.S.-based mobile operator companies. The case study uses project examples to illustrate the approach and explains step by step how the method can be used to include a range of stakeholders in the design process—regardless of their skill set or background. The case study offers exercises so that you can experiment with the approach yourself. ■

Two common properties that are often traded off against each other are breadth of functionality versus depth. These two kinds of prototyping are called *horizontal prototyping* (providing a wide range of functions but with little detail) and *vertical prototyping* (providing a lot of detail for only a few functions). Another common compromise is level of robustness versus degree of changeability. Making a prototype robust may lead to it being harder to change. This compromise may not be visible to a user of the product until something goes wrong. For example, the internal structure of a piece of code may not have been carefully designed, or the connections between electronic components may be delicate.

One of the consequences of high-fidelity prototypes is that the prototype can appear to be good enough to be the final product, and users may be less prepared to critique something if they perceive it as a finished product. Another consequence is that fewer alternatives are considered because the prototype works and users like it.



"THEN IN HERE WE DO A CLAY MOCK-UP
OF THE COMPUTER MODEL"

Source: Reproduced with permission of Penwil Cartoons

Although prototypes will have undergone extensive user evaluation, they will not necessarily have been built with good engineering principles or subjected to rigorous quality testing for other characteristics such as security and error-free operation. Building a product to be used by thousands or millions of people running on various platforms and under a wide range of circumstances requires a different construction and testing regime than producing a quick prototype to answer specific questions.

The next "Dilemma" box discusses two different development philosophies. In *evolutionary prototyping*, a prototype evolves into the final product and is built with these engineering principles in mind. *Throwaway prototyping* uses the prototypes as stepping stones toward the final design. In this case, the prototypes are thrown away, and the final product is built from scratch. In an evolutionary prototyping approach, each stage will be subjected to rigorous testing; for throwaway prototyping, such testing is not necessary.

DILEMMA

Prototyping vs. Engineering

The compromises made when developing low-fidelity prototypes are evident, but compromises in high-fidelity prototypes are not so obvious. When a project is under pressure, it can become tempting to integrate a set of existing high-fidelity prototypes together to form the final product. Many hours will have been spent developing them, and evaluation with users has gone well. So, why throw it all away? Generating the final product this way will simply store up testing and maintenance problems for later (see Box 13.1 on technical debt). In short, this is likely to compromise the quality of the product, unless the prototypes have been built with sound engineering principles from the start.

On the other hand, if the device is an innovation, then being first to market with a “good enough” product may be more important for securing market position than having a very high-quality product that reaches the market two months after a competitor’s product.

The dilemma arises in deciding how to treat high-fidelity prototypes—engineer them from the start or accept that they will be thrown away. ■

12.3 Conceptual Design

Conceptual design is concerned with developing a conceptual model; conceptual models were introduced in Chapter 3, “Conceptualizing Interaction.” The idea of a conceptual model can be difficult to grasp because these models take many different forms and there isn’t a definitive detailed characterization of one. Instead, conceptual design is best understood by exploring and experiencing different approaches to it, and the purpose of this section is to provide some concrete suggestions about how to go about doing this.

A *conceptual model* is an outline of what people can do with a product and which concepts are needed for the user to understand how to interact with it. The former will emerge from an understanding of the problem space and the current functional requirements. Which concepts are needed to understand how to interact with the product depends on a variety of issues such as who the user will be, what kind of interaction will be used, what kind of interface will be used, terminology, metaphors, application domain, and so on? The first step in developing a conceptual model is to steep yourself in the data about the users and their goals and try to empathize with them.

There are different ways to achieve empathy with users. For example, Karen Holtzblatt and Hugh Beyer’s (2017) contextual interviews, interpretation sessions and the Wall Walk, introduced in Chapter 11, supports this. These three activities together ensure that different people’s perspectives on the data and what they observed are captured, help to deepen understanding and to expose the whole team to different aspects of the problem space, and immerse the team in the users’ world. This stimulates ideas based on an extended understanding of the user and their context. Once captured, ideas are tested against other data and scenarios, discussed with

other design team members, and prototyped for testing with users. *Contextual Design* describes further activities that cover the whole design process. Trying to empathize with users may not be the right approach, however, as discussed in the following “Dilemma” box.

Using different creativity and brainstorming techniques to explore ideas with other members of the team can help build a picture of the users and their goals. Gradually, a picture of the desired users’ experience will emerge and become more concrete. This process is helped by considering the issues in this section and by using scenarios and prototypes to capture and experiment with ideas. The availability of ready-made components increases the ease with which ideas can be prototyped, which also helps to explore different conceptual models and design ideas. Mood boards (traditionally used in fashion and interior design) may be used to capture the desired feel of a new product (see Figure 12.8). This is informed by any data gathering or evaluation activities and considered in the context of technological feasibility.

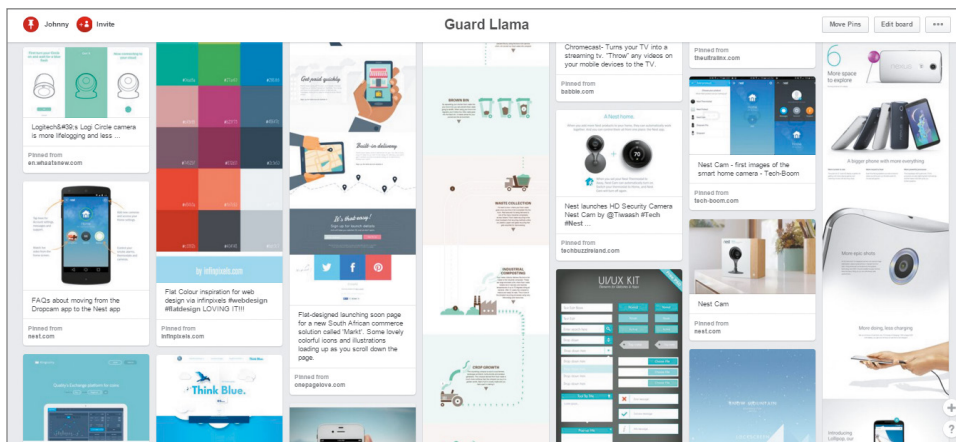


Figure 12.8 An example mood board developed for a personal safety product called Guard Llama

Source: <http://johnnyhuang.design/guardllama.html>

Read about how to create mood boards for UX projects here:

<https://uxplanet.org/creating-better-moodboards-for-ux-projects-381d4d6daf70>

Invision offers a tool to help with this. See the following web page:

<https://www.invisionapp.com/inside-design/boards-share-design-inspiration-assets/>

Developing a range of scenarios, as described in Chapter 11, can also help with conceptual design (Bødker, 2000) and to think through the consequences of different ideas. Suzanne Bødker (2000) also proposes the notion of plus and minus scenarios. These attempt to capture the most positive and the most negative consequences of a particular proposed design solution, thereby helping designers to gain a more comprehensive view of the proposal.

This idea has been extended by Mancini et al. (2010) who used positive and negative video scenarios to explore futuristic technology. Their approach used video to represent positive and negative consequences of a new product to help with diet and well-being, which was designed to explore privacy concerns and attitudes. The two videos (each with six scenes) focus on Peter, an overweight businessman who has been advised by his doctor to use a new product DietMon to help him lose weight. The product consists of glasses with a hidden camera, a microchip in the wrist, a central data store, and a text messaging system to send messages to Peter's mobile phone telling him the calorific value of the food he is looking at and warning him when he is close to his daily limit (Price et al, 2010). Figure 12.9 shows the content of two scenes from the videos and the positive and negative reactions; Figure 12.10 is a snapshot from the negative video.

| Scene 2: breakfast at home | |
|--|---|
| <p>Peter starts preparing his breakfast with his new glasses on. His wife notices them and he <i>keenly</i> gives her a demonstration of what they are and how they work, and tells her about the microchip. She seems <i>impressed</i> and leaves the room to get ready for work. Peter opens the fridge to put away the butter and sees a pastry. He looks at it and gets a DietMon message telling him the calorie content of the pastry. He shows that to his wife, who is entering the kitchen and looks at him with a <i>smile</i>.</p> | <p>Peter prepares breakfast with his new glasses on. His wife notices them. While looking at his toast, he gets a text. His wife enquires what that is. He says it's nothing and he does not feel like having toast after all. When she questions why he becomes <i>tense</i> and <i>reluctantly</i> tells her about DietMon. <i>Skeptical</i>, she leaves the room with a sarcastic comment. Peter opens the fridge and sees a pastry. As he gives in and takes a bite, he is caught by his wife, who is entering the kitchen and looks at him with a <i>grin</i>.</p> |
| Scene 3: birthday party at the office | |
| <p>Peter is working away at his desk when some colleagues invite him to a small birthday celebration. He tries to refuse but they insist. As he joins them, wearing his glasses, he greets the birthday-lady. His colleague Chris serves him a slice of cake. Peter looks at it and takes out his mobile. He gets a text, checks it and says the slice is too big, and asks Chris to cut it in a half. Chris is intrigued and asks for an explanation, so Peter gives his colleagues a <i>keen</i> demonstration of how the technology works. His audience is <i>impressed</i>, gathered around him.</p> | <p>Peter is working away at his desk when some colleagues invite him to a small birthday celebration. He tries to refuse but they insist. As he joins them, wearing his glasses, his colleague Chris gives him a slice of cake. He takes the plate and greets the birthday-lady. He gets a text and, <i>pretending</i> it's an important phone call, moves <i>away</i> from the others with the cake. Turned away from them, he <i>throws</i> the cake in a bin and goes back pretending to have already finished it. Chris comments on how fast he ate. Peter excuses himself, saying he has a deadline to meet, and leaves.</p> |

Figure 12.9 How two scenes from the videos differ in terms of positive and negative reactions to the system. The positive version is on the left and the negative on the right

Source: Mancini et al. (2010)



Figure 12.10 Peter being caught eating the pastry out of the fridge at breakfast (scene 2, negative reaction)

Source: Price et al. (2010)

DILEMMA

Is Trying to Empathize with Users the Right Approach?

Empathizing with users who live in a very different context than the designers is not easy, no matter how much data is collected. Interaction designers have tried several ways to understand situations that are outside their experience, two of which are experience prototyping and the Third Age suit.

(Continued)

Experience prototyping was introduced by Marion Buchenau and Jane Fulton Suri (2000) who describe a team designing a chest-implanted automatic defibrillator. A defibrillator is used with victims of cardiac arrest to deliver a strong electric shock to the heart that is intended to restore the heart muscle to its regular rhythm. This kind of event is completely outside most people's experience. To simulate some critical aspects of the experience, one of which is the random occurrence of a defibrillating shock, each team member was sent text messages at random times over one weekend. Each message simulated the occurrence of a defibrillating shock, and team members were asked to record where they were, who they were with, what they were doing, and what they thought and felt knowing that this represented a shock. Example insights ranged from anxiety around everyday happenings, such as holding a child and operating power tools, to being in social situations and at a loss how to communicate to onlookers what was happening. This first-hand experience brought new insights to the design effort.

The second example is the Third Age suit, designed so that car designers can experience what it is like for people with some loss of mobility or declining sensory perception to drive their cars. The suit restricts movement in the neck, arms, legs, and ankles. Originally developed by Ford Motor Company and Loughborough University (see Figure 12.11), it has been used to raise awareness within groups of car designers, architects, and other product designers.

Using these techniques appears to have brought new insights to the design process, but how deep are those insights and how accurate are they? According to Michelle Nario-Redmond



Figure 12.11 The Third Age suit helps designers experience the loss of mobility and sensory perception.

Source: Ford Motor Co.

et al. (2017), who conducted experiments to investigate the impact of disability simulations, they have unexpected negative consequences. They found that these simulations can result in feelings of fear, apprehension, and pity toward those with disabilities, rather than any sense of empathy. In addition, experiencing the disability for only a short time does not take into account the various coping strategies and innovative techniques that individuals develop. They suggest that a better way to design for these circumstances is to involve people with those disabilities and to understand their experiences more holistically.

To see the Third Age suit in action, watch this video:
<https://www.youtube.com/watch?v=Yb0aqr0rzrs>

To read an overview of the disability simulation experiment results, see this article:
<https://blog.prototypr.io/why-i-wont-try-on-disability-to-build-empathy-in-the-design-process-and-you-should-think-twice-7086ed6202aa>

12.3.1 Developing an Initial Conceptual Model

The core components of the conceptual model are metaphor and analogies, the concepts to which users are exposed, the relationship between those concepts, and the mappings between the concepts and user experience being supported (Chapter 3). Some of these will derive from the product's requirements, such as the concepts involved in a task and their relationships, such as through scenarios and use cases. Others such as suitable metaphors and analogies will be informed by immersion in the data and attempting to understand the users' perspectives.

This section introduces approaches that help to produce an initial conceptual model. In particular, it considers the following:

- How to choose interface metaphors that will help users understand the product?
- Which interaction type(s) would best support the users' activities?
- Do different interface types suggest alternative design insights or options?

All of these approaches provide different ways of thinking about the product and help generate potential conceptual models.

Interface Metaphors

Interface metaphors combine familiar knowledge with new knowledge in a way that will help users understand the product. Choosing suitable metaphors and combining new and familiar concepts requires a balance between utility and relevance, and it is based on an understanding of the users and their context. For example, consider an educational system to

teach 6-year-olds mathematics. One possible metaphor is a classroom with a teacher standing at the front. But considering the users of the product and what is likely to engage them, a metaphor that reminds them of something enjoyable is more likely to keep them engaged, such as a ball game, the circus, a playroom, and so on.

Different approaches to identifying and choosing an interface metaphor have been tried. For example, Dietrich Kammer et al. (2013) combined creativity methods to explore everyday objects, paper prototypes, and toolkits to support groups of students designing novel interface metaphors and gestures for mobile devices. They found that developing metaphors for both tablets and smartphones resulted in flexible metaphors. On the other hand, Marco Speicher et al. (2018) decided on an apartment metaphor for a VR online shopping experience by considering the limitations of systems that try to mimic physical stores.

Tom Erickson (1990) suggests a three-step process for choosing a good interface metaphor. Although this work is quite old, the approach is remarkably useful with current technologies. The first step is to understand what the system will do, that is, to identify functional requirements. Developing partial conceptual models and trying them may be part of the process. The second step is to understand which bits of the product are likely to cause users problems, that is, which tasks or subtasks cause problems, are complicated, or are critical. A metaphor is only a partial mapping between the product and the real thing upon which the metaphor is based. Understanding areas in which users are likely to have difficulties means that the metaphor can be chosen to support those aspects. The third step is to generate metaphors. Looking for metaphors in the users' description of relevant activities, or identifying metaphors used in the application domain, is a good starting point.

When suitable metaphors have been generated, they need to be evaluated. Erickson (1990) suggests five questions to ask:

- How much structure does the metaphor provide? A good metaphor will provide structure—preferably familiar structure.
- How much of the metaphor is relevant to the problem? One of the difficulties of using metaphors is that users may think they understand more than they do and start applying inappropriate elements of the metaphor to the product, leading to confusion or false expectations.
- Is the interface metaphor easy to represent? A good metaphor will be associated with particular physical, visual, and audio elements, as well as words.
- Will your audience understand the metaphor?
- How extensible is the metaphor? Does it have extra aspects that may be useful later?

For the group travel organizer introduced in Chapter 11, one potential metaphor that was prompted by the quote from Sky in her persona is a family restaurant. This seems appropriate because the family is all together, and each can choose what they want. Evaluating this metaphor using the previous five questions listed prompted the following thoughts:

- Does it supply structure? Yes, it supplies structure from the users' point of view, based on the familiar restaurant environment. Restaurants can be very different in their interior and the food they offer, but the structure includes having tables and a menu and people to serve the food. The experience of going to a restaurant involves arriving, sitting at a table, ordering food, being served the food, eating it, and then paying before leaving. From a different point of view, there is also structure around food preparation and how the kitchens are run.

- How much of the metaphor is relevant? Choosing a vacation involves seeing what is being offered and deciding what is most attractive, based on the preferences of everyone in the group. This is similar to choosing a meal in a restaurant. For example, a restaurant will have a menu, and visitors to the restaurant will sit together and choose individual meals, but they all sit in the same restaurant and enjoy the environment. For a group vacation, it may be that some members of the group want to do different activities and come together for some of the time, so this is similar. Information about the food such as allergens is available from the server or in the menu, reviews of restaurants are available, and photos or models of the food available are common. All of these characteristics are relevant to the group travel organizer. One of the characteristics of a restaurant that is not so applicable to a vacation is the need to pay at the end of the meal rather than before you get there.
- Is the metaphor easy to represent? There are several options in this regard, but the basic structure of a restaurant can be represented. The key aspect of this conceptual model will be to identify potential vacations that suit everyone and choose one. In a restaurant this process involves looking at menus, talking to the server, and ordering the food. Vacation information including photos and videos could be presented in a menu—maybe as one menu for adults and one for children. So, the main elements of the metaphor seem straightforward to represent.
- Will your audience understand the metaphor? For this example, the user group has not yet been investigated in detail, but eating in a restaurant is common.
- How extensible is the metaphor? There are several different types of restaurant experiences—à la carte, fixed menu, serve yourself, all you can eat, and food courts, for example. Elements from these different types of restaurants may be used to expand initial ideas.

ACTIVITY 12.2

One of the disadvantages of the restaurant metaphor is the need to have a shared experience when members of the group are in different locations. Another possible interface metaphor for the group travel organizer is the travel consultant. A travel consultant discusses the requirements with the traveler(s) and tailors the vacation accordingly, offering maybe two or three alternatives, but making most of the decisions on the travelers' behalf. Ask the earlier five questions about this metaphor.

Comment

1. Does the travel consultant metaphor supply structure?

Yes. The key characteristic of this metaphor is that the travelers specify what they want, and the consultant researches the options. It relies on the travelers giving the consultant sufficient information to search within a suitable range rather than leaving them to make key decisions.

2. How much of the metaphor is relevant?

The idea of handing over responsibility to someone else to search for suitable vacations may be appealing to some users, but it might feel uncomfortable to others. The level of

(Continued)

responsibility given to the consultant can be adjusted, though, depending on user preferences. It is common for individuals to put together vacations themselves based on web searches, but this can be time-consuming and diminish the excitement of planning a vacation. It would be attractive to some users if the initial searching and sifting is done for them.

3. Is the metaphor easy to represent?

Yes, it could be represented by a software agent or by having a sophisticated database entry and search facility. But the question is, would users like this approach?

4. Will your audience understand the metaphor?

Yes.

5. How extensible is the metaphor?

The wonderful thing about people is that they are flexible; hence, the metaphor of the travel consultant is also pretty flexible. For example, the consultant could be asked to refine their vacation recommendations according to as many different criteria as the travelers require. ■

Interaction Types

Chapter 3 introduced five different types of interaction: instructing, conversing, manipulating, exploring, and responding. Which type of interaction is best suited to the current design depends on the application domain and the kind of product being developed. For example, a computer game is most likely to suit a manipulating style, while a software application for drawing or drafting has aspects of instructing and conversing.

Most conceptual models will include a combination of interaction types, and different parts of the interaction will be associated with different types. For example, in the group travel organizer, one of the user tasks is to find out the visa regulations for a particular destination. This will require an instructing approach to interaction as no dialog is necessary for the system to show the regulations. The user simply has to enter a predefined set of information, for instance, the country issuing the passport and destination. On the other hand, trying to identify a vacation for a group of people may be conducted more like a conversation. For example, the user may begin by selecting some characteristics of the destination and some time constraints and preferences. Then the organizer will respond with several options, and the user will provide more information or preferences and so on. Alternatively, for users who don't have any clear requirements yet, they might prefer to explore availability before asking for specific options. Responding could be used when the user chooses an option that has additional restrictions and the system asks if the user meets them.

Interface Types

Considering different interfaces at this stage may seem premature, but it has both a design and a practical purpose. When thinking about the conceptual model for a product, it is important not to be unduly influenced by a predetermined interface type. Different interface types prompt and support different perspectives on potential user experiences and possible behaviors, hence prompting alternative design ideas.

In practical terms, prototyping the product will require an interface type, or at least candidate alternative interface types. Which ones to choose depends on the product constraints

that arise from the requirements. For example, input and output modes will be influenced by user and environmental requirements. Therefore, considering interfaces at this point also takes one step toward producing practical prototypes.

To illustrate this, we consider a subset of the interfaces introduced in Chapter 7, “Interfaces,” and the different perspectives they bring to the group travel organizer.

- **Shareable Interface** The travel organizer has to be shareable, as it is intended to be used by a group of people and it should be exciting and fun. The design issues for *shareable interfaces*, which were introduced in Chapter 7, will need to be considered for this system. For example, how best (whether) to use the individuals’ own devices such as smartphones in conjunction with a shared interface. Allowing group members to interact at a distance suggests the need for multiple devices, so a combination of form factors is required.
- **Tangible Interface** *Tangible interfaces* are a kind of sensor-based interaction, where blocks or other physical objects are moved around. Thinking about a travel organizer in this way conjures up an interesting image of people collaborating, maybe with the physical objects representing themselves while traveling, but there are practical problems of having this kind of interface, as the objects may be lost or damaged.
- **Virtual Reality** The travel organizer seems to be an ideal product for making use of a *virtual reality* interface, as it would allow individuals to experience the destination and maybe some of the activities available. Virtual reality would not be needed for the whole product, just for the elements where users want to experience the destination.

ACTIVITY 12.3

Consider the new navigation app for a large shopping center introduced in Chapter 11.

1. Identify tasks associated with this product that would best be supported by each of the interaction types instructing, conversing, manipulating, exploring, and responding.
2. Pick out two interface types from Chapter 7 that might provide different perspectives on the design.

Comment

1. Here are some suggestions. You may have identified others.
 - **Instructing** The user wants to see the location of a specific shop.
 - **Conversing** The user wants to find one particular branch out of several; the app might ask them to pick one from a list. Or, the user might want to find a particular kind of shop, and the app will display a list from which to choose.
 - **Manipulating** The chosen route could be modified by dragging the path to encompass other shops or specific walkways.
 - **Exploring** The user might be able to walk around the shopping center virtually to see what shops are available.
 - **Responding** The app asks whether the user wants to visit their favorite snack bar on the way to the chosen shop.

(Continued)

2. Navigation apps tend to be smartphone-based, so it is worth exploring other styles to see what insights they may bring. We had the following thoughts, but you may have had others. The navigation app needs to be mobile so that the user can move around to find the relevant shop. Using voice or gesture interfaces is one option, but this could still be delivered through a mobile device. Thinking more broadly, perhaps a haptic interface that guides the user to the required location might suffice. Smart interfaces, such as one built into the environment is an alternative, but privacy issues may arise if an individual's data is displayed for all to see. ■

12.3.2 Expanding the Initial Conceptual Model

The previous elements represent the core of the conceptual model. For prototyping or testing with users, these ideas need some expansion. Examples include which functions the product will perform and which the user will perform, how those functions are related, and what information is required to support them. Some of this will have been considered during the requirements activity and will evolve after prototyping and evaluation.

What Functions Will the Product Perform?

This question is about whether the product or the user takes responsibility for different parts of the overall goal. For example, the travel organizer is intended to suggest specific vacation options for a group of people, but is that all it should do? What about if it automatically reserved the bookings? Or does it wait until it is given a preferred choice? In the case of visa requirements, will the travel organizer simply provide the information, or link to visa services? Deciding what the system will do and the user will do is sometimes called *task allocation*. This trade-off has cognitive implications (see Chapter 4, “Cognitive Aspects”) and affects social aspects of collaboration (see Chapter 5, “Social Interaction”). If the cognitive load is too high for the user, then the device may be too stressful to use. On the other hand, if the product has too much control and is too inflexible, then it may not be used at all.

Another decision is which functions to hard-wire into the product and which to leave under software control, thereby indirectly in the control of the human user.

How Are the Functions Related to Each Other?

Functions may be related temporally; for example, one must be performed before another, or two can be performed in parallel. They may also be related through any number of possible categorizations, for instance, all functions relating to privacy on a smartphone or all options for viewing photographs on a social networking site. The relationships between tasks may constrain use or may indicate suitable task structures within the product. For example, if one task depends on another, the order in which tasks can be completed may need to be restricted. If use cases or other detailed analysis of the tasks have been generated, these will help. Different styles of requirements (for example, stories or atomic requirements shell) provide different levels of detail, so some of this information will be available, and some will evolve as the design team explores and discusses the product.

What Information Is Needed?

What data is required to perform a task? How is this data to be transformed by the system? Data is one of the categories of requirements identified and captured through the requirements activity. During conceptual design, these requirements are considered to ensure that the model provides the information needed to perform the task. Detailed issues of structure and display, such as whether to use an analog display or a digital display, will more likely be dealt with during the concrete design activity, but implications arising from the type of data to be displayed may impact conceptual design issues.

For example, identifying potential vacations for a group of people using the travel organizer requires the following: what kind of vacation is required; available budget; preferred destinations (if any); preferred dates and duration (if any); how many people it is for; and are there any special requirements (such as disability) within the group? To perform the function, the system needs this information and must have access to detailed vacation and destination descriptions, booking availability, facilities, restrictions, and so on.

Initial conceptual models may be captured in *wireframes*—a set of documents that show structure, content, and controls. Wireframes may be constructed at varying levels of abstraction, and they may show part of the product or a complete overview. Chapter 13, “Interaction Design in Practice,” includes more information and some examples.

12.4 Concrete Design

Conceptual design and concrete design are closely related. The difference between them is rather a matter of changing emphasis: during design, conceptual issues will sometimes be highlighted, and at other times, concrete detail will be stressed. Producing a prototype inevitably means making some concrete decisions, albeit tentatively, and since interaction design is iterative, some detailed issues will come up during conceptual design, and vice versa.

Designers need to balance the range of environmental, user, data, usability, and user experience requirements with functional requirements. These are sometimes in conflict. For example, the functionality of a wearable interactive product will be restricted by the activities the user wants to perform while wearing it; a computer game may need to be learnable but also challenging.

There are many aspects to the concrete design of interactive products: visual appearance such as color and graphics, icon design, button design, interface layout, choice of interaction devices, and so on. Chapter 7 introduces several interface types, together with their associated design considerations, guidelines, principles, and rules, which help designers ensure that their products meet usability and user experience goals. These represent the kinds of decision that are made during *concrete design*.

Concrete design also deals with issues related to user characteristics and context. Two aspects that have drawn particular attention for concrete design are discussed in this section: accessibility and inclusiveness; and designing for different cultures. Accessibility and inclusiveness were introduced in Section 1.6. *Accessibility* refers to the extent to which a product is accessible to as many people as possible, while *inclusiveness* means being fair, open, and equal to everyone. The aim of inclusive design is to empower users in their everyday and working lives (Rogers and Marsden, 2013).

A range of input and output modes is available for interaction design. Apart from standard keyboard, mouse, and touchscreen, there are also different pointing devices and

keyboards, screen readers, refreshable Braille, and eye-tracking, among others. Regardless of which alternate input or output modalities are used, interactive interfaces must be flexible enough to work with these various devices. This is particularly important for users with disabilities, who may be unable to use pointing devices or standard keyboards and may instead interact using a head or mouth stick, voice recognition, video with captions, audio transcripts, and so on.

Making an interface accessible involves engaging users with disabilities in the development process to better understand their needs and using the Web Content Accessibility Guidelines (WCAG), which applies to all interfaces, not just web-based interfaces (see Box 16.2). When interfaces are designed to be accessible, they not only work for people with disabilities, but they also provide flexibility to other users without disabilities who are faced with temporary or situational impairments, for example, a driver who is unable to look at a display screen or a train passenger watching a video without disturbing other passengers.

Interfaces that are not accessible can lead to various forms of discrimination. For instance, air fares are often lower if purchased through a website. If that website is inaccessible for blind consumers, for example, and those consumers have to use the call center, they may unknowingly be charged higher fares for the same flight (Lazar et al., 2010). Many companies use online job applications, but if the online site is inaccessible, job applicants may be forced to identify themselves as having a disability before applying for a job. Similarly, when job aggregator websites (with information about jobs at many different employers) are inaccessible and individuals who are blind are told to call on the phone as an accommodation, they frequently aren't told about many or even any of the jobs available (Lazar et al., 2015).

There are resources available to help design for inclusivity, accessibility, and flexibility, such as Microsoft's inclusive design toolkit.

Microsoft's inclusive design toolkit has some useful and interesting resources. You can find out more at <https://www.microsoft.com/design/inclusive>

Aspects of cross-cultural design include use of appropriate language(s), colors, icons and images, navigation, and information architecture (Rau et al., 2013). These are all important for concrete design; however, tensions between local cultures and HCI principles have been highlighted (Winschiers-Theophilus and Bidwell, 2013) together with a desire to reframe human-computer interaction (HCI) through local and indigenous perspectives (Abdelnour-Nocera et al., 2013). These concerns not only impact concrete design but also wider issues such as what to design and how to design in order to be accepted by the target user group. For example, Gary Marsden et al. (2008) warn of the problems in seeing a user's need and attempting to meet that need without first asking the community if they too recognize that need. For one approach on how to address this concern, see Case Study 12.2.

12.5 Generating Prototypes

This section illustrates how prototypes may be used in design, and it demonstrates how prototypes may be generated from the output of the requirements activity—producing a storyboard from a scenario and an index card-based prototype from a use case. Both of these are low-fidelity prototypes.

12.5.1 Generating Storyboards

A *storyboard* represents a sequence of actions or events that the user and the product go through to achieve a goal. A *scenario* is one story about how a product may be used to achieve that goal. A storyboard can be generated from a scenario by breaking the scenario into a series of steps that focus on interaction and creating one scene in the storyboard for each step. The purpose for doing this is twofold: first to produce a storyboard that can be used to get feedback from users and colleagues and second to prompt the design team to consider the scenario and the product's use in more detail. For example, consider the scenario for the travel organizer developed in Chapter 11. This can be broken down into six main steps.

1. Will, Sky, and Eamonn gather around the organizer, but Claire is at her mother's house.
2. Will tells the organizer their initial idea of a sailing trip in the Mediterranean.
3. The system's initial suggestion is that they consider a flotilla trip, but Sky and Eamonn aren't happy.
4. The travel organizer shows them some descriptions written by young people about flotilla trips.
5. Will confirms this recommendation and asks for details.
6. The travel organizer emails details of the different options.

Notice that the first step sets the context, and later steps focus more on the goal. Breaking the scenario into steps can be achieved in different ways. The purpose of working from the scenario is for the design team to think through the product and its use, and so the steps are not as important as the thinking that happens through the process. Also, notice that some of these events are focused solely on the travel organizer's interface, and some are concerned with the environment. For example, the first one talks about the family gathering around the organizer, while the fourth and sixth are focused on the travel organizer. Storyboards can focus on the screens or on the environment, or a mixture of both. Either way, sketching out the storyboard will prompt the design team to think about design issues.

For example, the scenario says nothing about the kinds of input and output devices that the system might use, but drawing the organizer forces the designer to think about these things. There is some information in the scenario about the environment within which the system will operate, but drawing the scene requires specifics about where the organizer will be located and how interaction will continue. When focusing on the screens, the designer is prompted to consider issues including what information needs to be available and what information needs to be output. This all helps to explore design decisions and alternatives, but it is also made more explicit because of the drawing act.

The storyboard in Figure 12.12 includes elements of the environment and some of the screens. While drawing this, various questions came to mind such as how can the interaction be designed for all of the family? Will they sit or stand? How to handle remote participants?

What kind of help needs to be available? What physical components does the travel organizer need? How to enable all of the family to interact with the system (notice that the first scene uses voice input while other scenes have a keyboard option as well)? and so on. In this exercise, the questions it prompts are just as important as the end product.




Figure 12.12 The storyboard for the travel organizer

ACTIVITY 12.4

Activity 11.4 in Chapter 11 developed a futuristic scenario for the one-stop car shop. Using this scenario, develop a storyboard that focuses on the environment of the user. As you draw this storyboard, write down the design issues that it prompts.

Comment

The following is based on the scenario in the comment for Activity 11.4. This scenario breaks down into five main steps.

1. The user arrives at the one-stop car shop.
2. The user is directed into an empty booth.
3. The user sits down in the racing car seat, and the display comes alive.
4. The user can view reports.
5. The user can take a virtual reality drive in their chosen car.

The storyboard is shown in Figure 12.13. Issues that arose while drawing this storyboard included how to display the reports, what kind of virtual reality equipment is needed, what input devices are needed—a keyboard or touchscreen, a steering wheel, clutch, accelerator, and brake pedals? How much like actual car controls do the input devices need to be? You may have thought of other issues.




Figure 12.13 The storyboard generated from the one-stop car shop scenario in Activity 11.4 ■

12.5.2 Generating Card-Based Prototypes

Card-based prototypes are commonly used to capture and explore elements of an interaction, such as dialog exchanges between the user and the product. The value of this kind of prototype lies in the fact that the interaction elements can be manipulated and moved around in order to simulate interaction with a user or to explore the user's end-to-end experience. This may be done as part of the evaluation or in conversations within the design team. If the storyboard focuses on the screens, this can be translated almost directly into a card-based prototype and used in this way. Another way to produce a card-based prototype is to generate one from a use case output from the requirements activity.

For example, consider the use cases for the visa requirements aspect of the group travel organizer presented in Section 11.6. The first, less-detailed use case provides an overview of the interaction, while the second one is more detailed.

This second use case can be translated into cards as follows. For each step in the use case, the travel organizer will need to have an interaction component to deal with it, for example, input via a button, menu option, or voice, and output via a display or sound. By stepping through the use case, a card-based prototype can be developed that covers the required behavior, and different designs can be considered. For example, Figure 12.14 shows six dialog elements on six separate cards. The set on the left has been written in friendlier language, while the set on the right is more official. These cover steps 1, 2, 3, 4, and 5.

The alternative courses, for example those dealing with error messages, would also each have a card, and the tone and information contained in the error message could be evaluated with users. For example, step 7.1 might translate into a simple “No visa information is available,” or a more helpful, “I am not able to find visa information for you to visit your chosen destination. Please contact the <destination country>'s embassy.”

These cards can be shown to potential users of the system or fellow designers to get informal feedback. In this case, we showed these cards to a colleague, and through discussion of the application and the cards, concluded that although the cards represent one interpretation




Figure 12.14 Cards 1–3 of a card-based prototype for the travel organizer

of the use case, they focus too much on an interaction model that assumes a WIMP/GUI interface. Our discussion was informed by several things including the storyboard and the scenario. One alternative would be to have a map of the world, and users can indicate their destination and nationality by choosing one of the countries on the map; another might be based around national flags. These alternatives could be prototyped using cards and further feedback obtained. Cards can also be used to elaborate other aspects of the concrete design, such as icons and other interface elements.

ACTIVITY 12.5

Look at the storyboard in Figure 12.4. This storyboard shows Christina exploring the Acropolis in search of information about the pottery trade. In the second scene in the top row, Christina “adjusts the preferences to find information about the pottery trade in Ancient Greece.” Many interaction icons have become standardized, but there isn’t a standard one for “pottery trade.” Suggest two alternative icons to represent this and draw them on separate

cards. Using the storyboard in Figure 12.4 and the two cards, try out the different icons with a friend or colleague to see what they understand by your two icons.

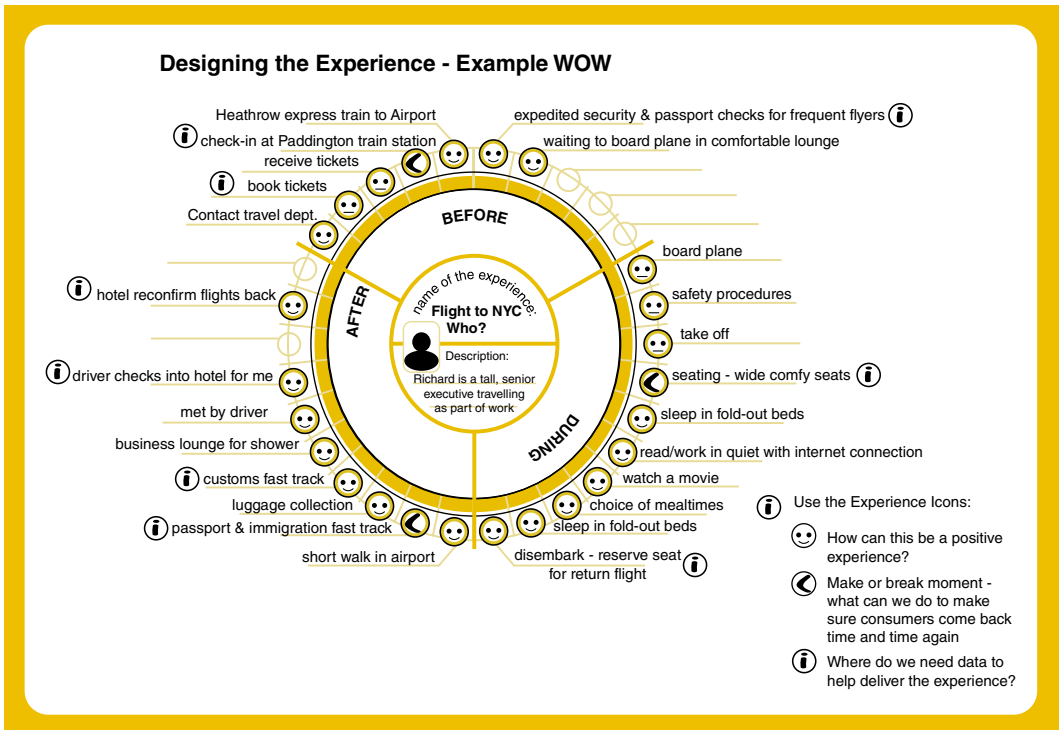
Comment

The two cards we drew are shown in Figure 12.15. The first is simply an Ancient Greek pot, while the second attempts to capture the idea of a pottery seller in the market. When we stepped through the storyboard with a colleague and showed them these alternatives, both were found to require improvement. The pot on its own did not capture the pottery trade, and it wasn't clear what the market seller represented, but there was a preference for the latter, and the users' feedback was useful.

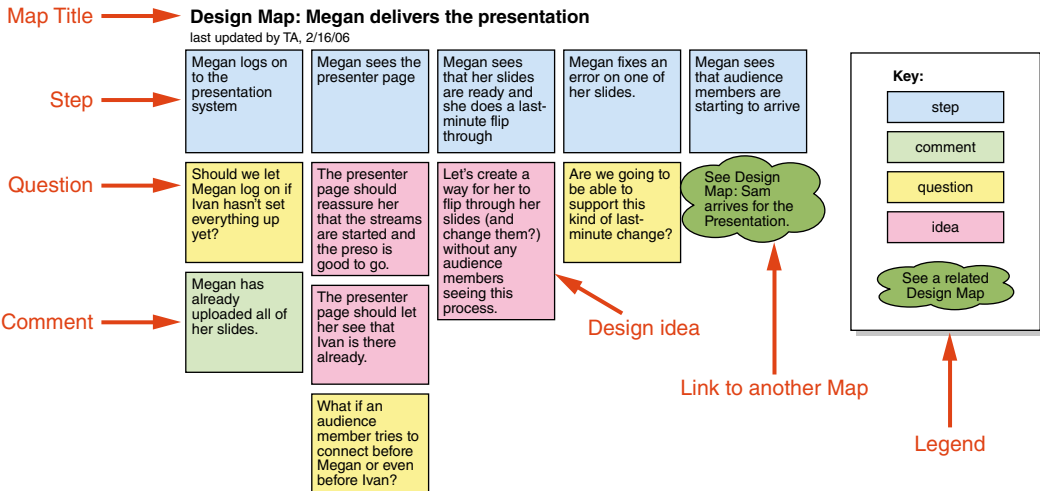


Figure 12.15 Two icons to represent “pottery trade” for the new mobile device for exploring historic sites depicted in the storyboard of Figure 12.4 ■

A set of card-based prototypes that cover a scenario from beginning to end may be the basis of a more detailed prototype, such as an interface or screen sketch, or it may be used in conjunction with personas to explore the user's end-to-end experience. This latter purpose is also achieved by creating a visual representation of the user's experience. These representations are variably called a *design map* (Adlin and Pruitt, 2010), a *customer journey map*



(a)



(b)

Figure 12.16 (a) An experience map using a wheel representation and (b) an example timeline design map illustrating how to capture different issues

Source: (a) LEGO (b) Adlin and Pruitt (2010), p. 134. Used courtesy of Morgan Kaufmann

(Ratcliffe and McNeill, 2012), or an *experience map*. They illustrate a user's path or journey through the product or service and are usually created for a particular persona and based on a particular scenario, hence giving the journey sufficient context and detail to bring the discussions to life. They support designers in considering the user's overall experience when achieving a particular goal and are used to explore and question the designed experience and to identify issues that have not been considered so far. They may be used to analyze existing products and to collate design issues, or as part of the design process.

There are many different types of representation of varying complexities. Two main ones are: the wheel and the timeline. The wheel representation is used when an interaction phase is more important than an interaction point, such as for a flight (see Figure 12.16(a) for an example). The timeline is used where a service is being provided that has a recognizable beginning and end point, such as purchasing an item through a website (an example of a timeline representation is shown in Figure 11.7(b)—look for the smiley faces). Figure 12.16(b) illustrates the structure of a timeline and how different kinds of issues may be captured, such as, questions, comments, and ideas.

To generate one of these representations, take one persona and two or three scenarios. Draw a timeline for the scenario and identify the interaction points for the user. Then use this as a discussion tool with colleagues to identify any issues or questions that may arise. Some people consider the user's mood and identify pain points, sometimes the focus will be on technical issues, and at other times this can be used to identify missing functionality or areas of under-designed interaction.

This video illustrates the benefits of experience mapping using a timeline:
http://youtu.be/eLT_Q8sRpyl.

To read about the main elements of a customer journey map, when you need them, and how to construct them, see this article:
<https://www.nngroup.com/articles/customer-journey-mapping/>.

BOX 12.3

Involving Users in Design: Participatory Design

Participatory design (PD) emerged in Scandinavia in the late 1960s and early 1970s. There were two influences on this early work: the desire to be able to communicate information about complex systems and the labor union movement pushing for workers to have democratic control over changes in their work. In the 1970s, new laws gave workers the right to have a say in how their working environment was changed, and such laws are still in force today.

(Continued)

The idea that those who use information technology will play a critical role in its design, and in particular that they will engage in active and genuine participation with the design itself, is still central to participatory design (Simonsen and Robertson, 2012). But the approach has evolved considerably in response to political, social, and technological changes (Bannon et al., 2018). In addition, many approaches to technology design include participation with users, so what makes participatory design different?

In a review of research in the PD conference 2002–2012, Kim Halskov et al. (2015) wanted to understand the different definitions of ‘participation’ as evidenced by those papers. They identified three approaches.

- Implicit, meaning that the paper wasn’t clear about the details of participation
- Users as full participants in the design process, which goes beyond simple involvement of users, but extends to understanding the user’s point of view, and regarding what users know as being important
- Mutual learning between users and designer

One of the key questions for participatory design today is how to handle scale: How to ensure participation by a community when that community includes several hundreds or thousands of people? How to ensure participation when the data collected can be from many different sources and without explicit agreement? How to ensure participation when users extend over several countries?

Daniel Gooch et al. (2018) designed an approach to facilitate citizen engagement in a smart city project. They used an integrated approach of online and offline activities that was tailored to local contexts and showed how it is possible to engage citizens in a way that addresses citizens’ current concerns. They also identified four key challenges to utilizing participatory design on an urban scale.

- Balancing scale with the personal. In particular, the need to engage face-to-face with potential participants.
- Who has control of the process? If participants are to have a meaningful say in what is designed, they need to have some of the power, but sometimes regulations mitigate against this.
- Who is participating? In a city, there are many diverse stakeholders, yet it is important to include all sections of the society to avoid bias.
- Integrating citizen-led work with local authorities. Regulations set by local authorities can be an obstacle to innovation within a city context.

Case Study 12.2 describes an extension to participatory design, called *community-based design*, developed for the local situation in South Africa. ■

CASE STUDY 12.2

Deaf Telephony

This case study by Edwin Blake, William Tucker, Meryl Glaser, and Adinda Freudenthal discusses their experiences of community-based design in South Africa. The process of community-based co-design is one that explores various solution configurations in a multidimensional design space whose axes are the different dimensions of requirements and the various dimensions of designer skills and technological capabilities. The bits of this space that one can “see” are determined by one’s knowledge of the user needs and one’s own skills. Co-design is a way of exploring that space in a way that alleviates the myopia of one’s own viewpoint and bias. As this space is traversed, a trajectory is traced according to one’s skills and learning and according to the users’ expressed requirements and their learning.

The project team set out to assist South African deaf people to communicate with each other, with hearing people, and with public services. The team has been working for many years with a deaf community that has been disadvantaged due both to poverty and hearing impairment. The story of this wide-ranging design has been one of continual fertile (and on occasion frustrating) co-design with this community. The team’s long-term involvement has meant that they have transformed aspects of the community and that they have themselves been changed in what they view as important and in how they approach design. Figure 12.17 illustrates one participant’s view of communication, captured during a community engagement event, and Figure 12.18 shows two participants discussing design using sign language.




Figure 12.17 One participant’s view of communication

Source: Edwin Blake

(Continued)



Figure 12.18 Participants discussing design in sign language

Source: Helen Sharp

Deaf users in this community started out knowing essentially nothing about computers. Their first language is South African Sign Language (SASL), and this use of SASL is a proud sign of their identity as a people. Many are also illiterate or semi-literate. There are a large number of deaf people using SASL; in fact, there are more using it than with some of the smaller official languages. Since the advent of democracy in 1994, there has been an increasing empowerment of deaf people, and it is accepted as a distinct language in its own right.

In this case study on id-book.com, a brief historical overview of the project and the various prototypes that formed nodes in a design trajectory are presented. The methodology of Action Research and its cyclical approach to homing in on an effective implementation is reviewed. An important aspect of the method is how it facilitates learning by both the researchers and the user community so that together they can form an effective design team. Lastly, such a long-term intimate involvement with a community raises important ethical issues, which are fundamentally concerns of reciprocity. ■

ACTIVITY 12.6

Design thinking has been described as an approach to problem-solving and innovative design that focuses on understanding what people want and what technology can deliver. It is derived from professional design practice, and it is often viewed as having five stages that together evolve a solution: empathize, define, ideate, prototype, and test. A slightly different view of design thinking, according to IDEO (<https://www.ideo.com/pages/design-thinking>), emphasizes human needs, empathy, and collaboration by looking at the situation through three lenses: desirability, feasibility, and viability.

Design thinking has become very popular, but some have questioned its benefits and implications. This activity invites you to decide for yourself.

Click the following links, and do some investigation yourself around the idea of design thinking. Based on what you find, do you think the turn toward design thinking overall is beneficial or damaging to interaction design?

Jon Kolko's (2018) article:

<http://interactions.acm.org/archive/view/may-june-2018/the-divisiveness-of-design-thinking>

Natasha Jen's (2017) presentation:

<https://vimeo.com/228126880>

Dan Nessler's (2016) article:

<https://medium.com/digital-experience-design/how-to-apply-a-design-thinking-hcd-ux-or-any-creative-process-from-scratch-b8786efbf812>

Comment

Design thinking is similar to the approaches espoused by user-centered design and the notion of design thinking has been embraced by many designers and organizations. Nevertheless, the way in which it has been popularized has resulted in some heavy criticism too. In her presentation, Natasha Jen criticizes the simple five-stage process and invites proponents to share the evidence of its success and its outcomes so that it can be improved.

Jon Kolko (2018) believes that this surge of interest in design thinking “will leave behind two benefits: validation of the design profession as real, intellectual, and valuable—and a very large need for designers who can make things.” However, he also points out that it has been popularized at a simplistic level of detail.

At the end of the day, what this suggests is that design is a creative activity supported by techniques, tools, and processes, but it cannot be boiled down into a particular process or set of techniques—design involves “the habit of continually doing things in new ways in order to make a difference,” as stated by Dan Nessler. ■

12.6 Construction

As prototyping and building alternatives progresses, development will focus more on putting together components and developing the final product. This may take the form of a physical product, such as a set of alarms, sensors, and lights, a piece of software, or both. Whatever the final form, it is unlikely that anything will need to be developed from scratch, as there are many useful (in some cases essential) resources to support development. Here we introduce two kinds of resources: physical computing kits and software development kits (SDKs).

12.6.1 Physical Computing

Physical computing is concerned with how to build and code prototypes and devices using electronics. Specifically, it is the activity of “creating physical artifacts and giving them behaviors through a combination of building with physical materials, computer programming, and circuit building” (Gubbels and Froehlich, 2014). Typically, it involves designing things, using a printed circuit board (PCB), sensors (for instance push buttons, accelerometers, infrared, or temperature sensors) to detect states, and output devices (such as displays, motors, or buzzers) that cause some effect. An example is a “friend or foe” cat detector that senses, via an accelerometer, any cat (or anything else for that matter) that tries to push through a family’s cat door. The movement triggers an actuator to take a photo of what came through the cat door using a webcam positioned on the back door. The photo is uploaded to a website that alerts the owner if the image does not match that of their own cat.

A number of physical computing toolkits have been developed for educational and prototyping purposes. One of the earliest is Arduino (see Banzi, 2009). The goal was to enable artists and designers to learn how to make and code physical prototypes using electronics in a couple of days, having attended a workshop. The toolkit is composed of two parts: the Arduino board (see Figure 12.19), which is the piece of hardware that is used to build objects, and the Arduino integrated development environment (IDE), which is a piece of software that makes it easy to program and upload a sketch (Arduino’s name for a unit of code) to the board. A sketch, for example, might turn on an LED when a sensor detects a change in the light level. The Arduino board is a small circuit that contains a tiny chip (the microcontroller). It has two rows of small electrical “sockets” that let the user connect sensors and actuators to its input and output pins. Sketches are written in the IDE using a simple processing language and then translated into the C programming language and uploaded to the board.




Figure 12.19 The Arduino board

Source: Used courtesy of Dr Nicolai Marquardt

There are other toolkits that have been developed, based on the basic Arduino kit. The most well-known is the LilyPad, which was co-developed by Leah Beuchley (see Figure 12.20 and her interview at the end of Chapter 7). It is a set of sewable electronic components for building fashionable clothing and other textiles. The Engduino is a teaching tool based on the Arduino LilyPad; it has 16 multicolor LEDs and a button, which can be used to provide

visual feedback and simple user input. It also has a thermistor (that senses temperature), a 3D accelerometer (that measures accelerations), and an infrared transmitter/receiver that can be used to transmit messages from one Engduino to another.




Figure 12.20 The LilyPad Arduino kit

Source: Used courtesy of Leah Beuchley

Watch this video that introduces Magic Cubes—a novel toolkit that is assembled from six sides that are slotted together to become an interactive cube that lights up in different colors, depending on how vigorously it is shaken. Intended to encourage children to learn, share, and fire their imagination to come up with new games and other uses, see it in action at <https://uclmagiccube.weebly.com/video.html>.

Other kinds of easy-to-use and quick-to-get-started physical toolkits, intended to provide new opportunities for people to be inventive and creative, are Senseboard (Richards and Woodthorpe, 2009), Raspberry Pi (<https://www.raspberrypi.org/>), .NET Gadgeteer (Villar et al., 2012), and MaKey MaKey (Silver and Rosenbaum, 2012). The MaKey MaKey toolkit is composed of a printed circuit board with an Arduino microcontroller, alligator clips, and a USB cable (see Figure 12.21). It communicates with a computer to send key presses, mouse clicks, and mouse movements. There are six inputs (the four arrow keys, the space bar, and a mouse click) positioned on the front of the board onto which alligator clips are clipped in order to connect with a computer via the USB cable. The other ends of the clips can be attached to any noninsulating object, such as a vegetable or piece of fruit. Thus, instead of using the computer keyboard buttons to interact with the computer, external objects such as bananas are used. The computer thinks MaKey MaKey is just like a keyboard or mouse. An example is to play a digital piano app using bananas as keys rather than keys on the computer keyboard. When they are touched, they make a connection to the board and MaKey MaKey sends the computer a keyboard message.




Figure 12.21 The MaKey MaKey toolkit

Source: Helen Sharp

One of the most recent physical computing systems is the BBC micro:bit (<https://microbit.org>, see Figure 12.22). Like Arduino, the micro:bit system consists of a physical computing device that is used in conjunction with an IDE. However, unlike Arduino, the micro:bit device contains a number of built-in sensors and a small display so that it is possible to create simple physical computing systems without attaching any components or wires. If desired, external components can still be added, but rather than the small electrical sockets of the Arduino, the micro:bit has an “edge connector” for this purpose. This is formed from a row of connection points that run along one edge of the device and allow it to be “plugged into” a range of accessories including larger displays, Xbox-style game controllers, and small robots. The micro:bit IDE, which runs in a web browser with no installation or setup process, supports a graphical programming experience based on visual “blocks” of code alongside text-based editing using a variant of JavaScript. This means that the micro:bit provides a great experience for young students and other beginner programmers, while also supporting more sophisticated programming. As a result, micro:bit has been widely adopted in schools around the world.




Figure 12.22 The BBC micro:bit

Source: <https://microbit.org/guide/features>. Used courtesy of Micro:bit Foundation

So far, physical toolkits have been aimed at children or designers to enable them to start programming through rapid creation of small electronic gadgets and digital tools (for example, Hodges et al., 2013, Sentance et al., 2017). However, Yvonne Rogers et al. (2014) demonstrated how retired people were equally able to be creative using the kit, turning “everyday objects into touchpads.” They ran a series of workshops where small groups of retired friends, aged between their early 60s and late 80s, assembled and played with the MaKey MaKey toolkit (see Figure 12.23). After playing music using fruit and vegetables as input, they saw many new possibilities for innovative design. Making and playing together, however childlike it might seem at first, can be a catalyst for imagining, free thinking, and exploring. People are sometimes cautious to volunteer their ideas, fearing that they are easily squashed, but in a positive environment they can flourish. The right kind of shared experience can create a positive and relaxed atmosphere in which people from all walks of life can freely bounce ideas off each other.



Figure 12.23 A group of retired friends playing with a MaKey MaKey toolkit

Source: Helen Sharp

BOX 12.4

The Rise of the Maker Movement

The maker movement emerged in the mid-2000s. Following in the footsteps of the personal computer revolution and the Internet, some viewed it as the next big transformation that would modernize manufacturing and production (Hatch, 2014). Whereas the explosion of the Web was all about what it could do for us virtually, with a proliferation of apps, social media, and services, the maker movement is transforming how we make, buy, consume, and recycle physical things, from houses to clothes and food to bicycles. At its core is DIY—crafting physical things using a diversity of machines, tools, and methods collaboratively in workshops and makerspaces. In a nutshell, it is about inventing the future through connecting technologies, the Internet, and physical things.

(Continued)

While there have always been hobbyists tinkering away making radios, clocks, and other devices, the world of DIY making has been opened up to many more people. Affordable, powerful, and easy-to-use tools, coupled with a renewed focus on locally-sourced products and community-based activities, and a desire for sustainable, authentic, and ethically-produced products, has led to a groundswell of interest in “making.” Fablabs (fabrication laboratories) first started appearing in cities throughout the world, offering a large physical space containing electronics and manufacturing equipment, including 3D printers, CNC milling machines, and laser cutters. Individuals bring their digital files to print and make things such as large 3D models, furniture, and installations—something that would have been impossible for them to do previously. Then smaller makerspaces started appearing in the thousands across the world, from Shanghai to rural India, again sharing production facilities for all to use and make. While some are small, for example sharing the use of a 3D printer, others are much larger and well resourced, offering an array of manufacturing machines, tools, and workspaces to make in.

Another development has been to build and program e-textiles using sewing machines and electronic thread. E-textiles comprise fabrics that are embedded with electronics, such as sensors, LEDs, and motors that are stitched together using conductive thread and conductive fabrics (Buechley and Qiu, 2014). An early example is the turn-signal biking jacket (developed by Leah Buechley and illustrated in Figure 1.4). Other e-textiles include interactive soft toys, wallpaper that sings when touched, and fashion clothing that reacts to the environment or events.

A central part of the maker movement involves tinkering (as discussed in section 12.2.4) and the sharing of knowledge, skills, know-how, and what you have made. The Instructables .com website is for anyone to explore, document, and share their DIY creations. Go to the Instructables site and take a look at a few of the projects that have been uploaded by makers. How many of them are a combination of electronics, physical materials, and pure invention? Are they fun, useful, or gadgety? How are they presented? Do they inspire you to make?

Another site, Etsy.com, is an online marketplace for people who make things to sell their crafts and other handmade items, which has grown in popularity over the past few years. It is designed to be easy for makers to use and to set up their store to sell to family, friends, and strangers across the world. Unlike corporate online sites, such as Amazon or eBay, Etsy is a place for craft makers to reach out to others and to show off their wares in ways that they feel best fit what they have made. This transition from “making” to “manufacturing,” albeit on the limited scale of craft production, is an interesting phenomenon. Some authors believe that the trend will continue and that increasingly new products and new businesses will emerge from activities rooted in maker culture (Hodges et al., 2014).

In essence, the maker movement is about taking the DIY movement online to make it public, and in doing so, massively increase who can take part and how it is shared (Anderson, 2013). In his interview at the end of this chapter, Jon Froehlich explains more about the maker movement. ■

12.6.2 SDKs: Software Development Kits

A *software development kit* (SDK) is a package of programming tools and components that supports the development of applications for a specific platform, for example, for iOS on iPhone and iPad and for Android on mobile phone and tablet apps. Typically, an SDK includes an integrated development environment, documentation, drivers, and sample programming code to illustrate how to use the SDK components. Some also include icons and buttons that can easily be incorporated into the design. While it is possible to develop applications without using an SDK, it is much easier using such a powerful resource and so much more can be achieved.

For example, the availability of Microsoft's Kinect SDK has made the device's powerful gesture recognition and body motion tracking capabilities accessible. This has led to the exploration of many applications including elderly care and stroke rehabilitation (Webster and Celik, 2014), motion tracking in immersive games (Manuel et al., 2012), user identification using body lengths (Hayashi et al., 2014), robot control (Wang et al., 2013), and virtual reality (Liu et al, 2018).

An SDK will include a set of application programming interfaces (APIs) that allows control of the components without the developer needing to know the intricacies of how they work. In some cases, access to the API alone is sufficient to allow significant work to be undertaken, for instance, Eiji Hayashi et al. (2014) only needed access to the APIs. The difference between APIs and SDKs is explained in Box 12.5.

See the following websites to learn about two different types of SDKs and their use:

- Building voice-based services with Amazon's Alexa Skills Kit:
<https://developer.amazon.com/alexa-skills-kit>.
- Constructing augmented reality experiences with Apple's ARKit:
<https://developer.apple.com/arkit/>.

BOX 12.5

APIs and SDKs

SDKs consist of a set of programming tools and components, while an API is the set of inputs and outputs, that is, the technical interface to those components. To explain this further, an API allows different-shaped building blocks of a child's puzzle to be joined together, while an SDK provides a workshop where all of the development tools are available to create whatever size and shape blocks you desire, rather than using preshaped building blocks. An API therefore allows the use of pre-existing building blocks, while an SDK removes this restriction and allows new blocks to be created or even to build something without blocks at all. An SDK for any platform will include all of the relevant APIs, but it adds programming tools, documentation, and other development support as well. ■

In-Depth Activity

This in-depth activity builds upon the requirements activities related to the booking facility introduced at the end of Chapter 11.

1. Based on the information gleaned from the activity in Chapter 11, suggest three different conceptual models for this system. Consider each of the aspects of a conceptual model discussed in this chapter: interface metaphor, interaction type, interface type, activities it will support, functions, relationships between functions, and information requirements. Of these conceptual models, decide which one seems most appropriate and articulate the reasons why.
2. Using the scenarios generated for the online booking facility, produce a storyboard for the task of booking a ticket for one of the conceptual models in step 1. Show it to two or three potential users and record some informal feedback.
3. Considering the product's concrete design, sketch out the application's initial interface. Consider the design issues introduced in Chapter 7 for the chosen interface type. Write one or two sentences explaining your choices and consider whether the choice is a usability consideration or a user experience consideration.
4. Sketch out an experience map for the product. Use the scenarios and personas you generated previously to explore the user's experience. In particular, identify any new interaction issues that had not been considered previously, and suggest what could be done to address them.
5. How does the product differ from applications that typically might emerge from the maker movement? Do software development kits have a role? If so, what is that role? If not, why not?

Summary

This chapter explored the activities of design, prototyping, and construction. Prototyping and scenarios are used throughout the design process to test ideas for feasibility and user acceptance. We have looked at different forms of prototyping, and the activities have encouraged you to think about and apply prototyping techniques in the design process.

Key points

- Prototyping may be low fidelity (such as paper-based) or high fidelity (such as software-based).
- High-fidelity prototypes may be vertical or horizontal.
- Low-fidelity prototypes are quick and easy to produce and modify, and they are used in the early stages of design.

- Ready-made software and hardware components support the creation of prototypes.
- There are two aspects to the design activity: conceptual design and concrete design.
- Conceptual design develops an outline of what people can do with a product and what concepts are needed to understand how to interact with it, while concrete design specifies the details of the design such as layout and navigation.
- We have explored three approaches to help you develop an initial conceptual model: interface metaphors, interaction styles, and interface styles.
- An initial conceptual model may be expanded by considering which functions the product will perform (and which the user will perform), how those functions are related, and what information is required to support them.
- Scenarios and prototypes can be used effectively in design to explore ideas.
- Physical computing kits and software development kits facilitate the transition from design to construction.

Further Reading

BANZI, M. and SHILOH, M. (2014) *Getting Started with Arduino* (3rd ed.). Maker Media Inc. This hands-on book provides an illustrated step-by-step guide to learning about Arduino with lots of ideas for projects to work on. It outlines what physical computing is in relation to interaction design and the basics of electricity, electronics, and prototyping using the Arduino hardware and software environment.

GREENBERG, S., CARPENDALE, S., MARQUARDT, N. and BUXTON, B. (2012) *Sketching User Experiences*. Morgan Kaufmann. This is a practical introduction to sketching. It explains why sketching is important, and it provides useful tips to get the reader into the habit of sketching. It is a companion book to Buxton, B. (2007) *Sketching User Experiences*. Morgan Kaufmann, San Francisco.

INTERACTIONS MAGAZINE (2018) *Designing AI*. ACM. This issue of the *Interactions* magazine is all about design and different aspects of it including sketching, human-centered design for children, collaborative art, design capabilities, and the special topic of designing for AI.

LAZAR, J., GOLDSTEIN, D., and TAYLOR, A. (2015). *Ensuring Digital Accessibility Through Process and Policy*. Waltham, MA: Elsevier/Morgan Kaufmann Publishers. This book is about accessibility, bringing together knowledge in technology, law, and research. It includes a range of standards, regulations, methods, and case studies.



INTERVIEW with Jon Froehlich

Jon Froehlich is an Associate Professor in the Paul G. Allen School of Computer Science and Engineering at the University of Washington (UW) where he directs the Makeability Lab (<http://makeabilitylab.io/>), a cross-disciplinary research group focused on applying computer science and HCI to high-value social domains such as environmental sustainability and STE(A)M education. He has published more than 50 peer-reviewed publications; 11 have been honored with awards, including Best Papers at ACM CHI and ASSETS and a 10-Year Impact Award at UbiComp. Jon is a father of two, and he is increasingly passionate about CS4All—both as an educator and a researcher.

Can you tell us a bit about your research, what you do, and why you do it?

The goal of my research is to develop interactive tools and techniques to address pressing global challenges in areas such as accessibility, STE(A)M education, and environmental sustainability. To succeed at this work, I collaborate across disciplines

with a focus on identifying long-term, ambitious research problems such as mapping the accessibility of the physical world via crowdsourcing plus computer vision that can also provide immediate, practical utility. Typically, my research involves inventing or reappropriating methods to sense physical or behavioral phenomena, leveraging techniques in computer vision (CV) and machine learning (ML) to interpret and characterize this data, and then building and evaluating interactive software or hardware tools uniquely enabled by these approaches. My research process is iterative, consisting of formative studies, which then inform the design and implementation of prototypes, followed by a series of evaluations, first in the lab and then eventually deployment studies of refined prototypes in the field.

What is the maker movement, and why are you so enthusiastic about it?

The maker movement emerged in the mid-2000s as an informal collection of hobbyists, engineers, artists, coders, and

craftspeople dedicated to playful creation, self-learning, and material design. While the movement builds on longstanding hobbyist and do-it-yourself (DIY) culture—for example, in woodworking and electronics—the movement was galvanized and accelerated by a series of socio-technical developments, including new, low-cost computational fabrication tools like CNC mills and 3D printers, the emergence of inexpensive and easy-to-use microcontroller platforms like Arduino and Raspberry Pi, online marketplaces like Adafruit and Sparkfun that made it easy to find and purchase parts, and social networks like Instructables, YouTube, and Thingiverse, which provided a forum for novices and experts alike to share and critique ideas, tutorials, and creations.

My enthusiasm for the maker movement stems both from my intrinsic excitement as a technologist in observing the creativity and creations of “makers” as well as from my perspectives as an educator and mentor in wondering how we can borrow from and adapt elements of the movement into formal education. While the maker movement is a relatively new phenomenon, its historical roots in education and learning science stretch back to pioneering educational thinkers like Maria Montessori, Jean Piaget, Seymour Papert, Lev Vygotsky, and others, all who emphasize the importance of learning through creation and experimentation, the role of peer mentorship, and how sharing work and soliciting feedback shapes thinking. For example, Papert’s Constructionism learning theory places a critical focus not just on learning through making but on the social nature of design—that is, that ideas are shaped by the knowledge of an audience and the feedback provided by others.

I have tried to inject this philosophy into my undergrad and graduate teaching. As one example, students in my Tangible Interactive Computing course explore the materiality of interactive computing via design prompts such as making a new input device for a computer using lo-fi materials like conductive clay and fabric, breaking and remaking an existing electronic technology to reformulate its physical interaction, and combining computer vision and video cameras to create whole-body, gestural input. Students share and critique each other’s work but also design outwardly beyond the confines of the classroom by sharing their results and design processes publicly (under pseudonyms, if preferred) via videos on YouTube, step-by-step tutorials on Instructables.com, and on the course website. Student-written Instructables in Tangible Interactive Computing, for example, have won awards and acquired more than 300,000 views and 1,900 favorites.

What are the advantages and challenges of working with communities to design products?

Much of my research involves designing and evaluating technologies for users who have different abilities, perspectives, and/or experiences from me and my research group—for example, early elementary school learners, people who use wheelchairs, or people with visual impairments. Thus, a key facet of our research and design process is employing methods from participatory design (or “co-design”), an approach to design that attempts to actively involve and empower target users throughout the design process from ideation to lo-fi prototyping to

(Continued)

summative evaluation. For example, in the MakerWear project (Kazemitabaar et al., 2017)—a wearable construction kit for children—we worked with children to gather design ideas and solicit critical feedback, to test initial designs, and to help co-design toolkit behavior and the overall look and feel. Similarly, we also involved professional STEM educators to help us improve our designs and think about corresponding learning activities. Finally, we ran a series of pilot studies followed by workshops in afterschool programs and a children’s museum to examine what and how children make with MakerWear, what challenges arise, and how their designs differ from creations made with other toolkits (for example, in robotics).

This human-centered, participatory design approach offers many advantages, including ensuring that we are addressing real user problems, helping ground our design decisions through use and feedback from target stakeholders, and empowering our users to have a real voice in shaping outcomes (from which our participants of all ages seem to gain satisfaction). There are trade-offs, however. Soliciting ideas from target users in an unstructured and unprincipled manner may lead to poorly defined outcomes and suboptimal designs. When working with children, we often follow Druin’s Cooperative Inquiry methodology (Guha et al., 2013), which provides a set of techniques and guidelines for co-design with children that helps to channel and focus their creativity and ideas. A second challenge is in recruiting and supporting co-design sessions: this is a resource-intensive process that requires time and effort from both stakeholders and the research team. To mitigate this challenge, we often work

on establishing and maintaining longitudinal relationships with community groups like local schools and museums. Finally, not all projects are amenable to these methods (such as when timelines are particularly aggressive).

Have you encountered any big surprises in your work?

The life of a researcher is full of surprises—one must get comfortable with ambiguity and ending a research journey at an unpredictable location. My most significant surprises, however, have come from people: from my students, from my mentors, and from my collaborators. My research methods and ideas have been profoundly influenced in unexpected ways by colleagues like Professor Tamara Clegg who made me rethink how we can personalize STEM learning through opportunities in everyday life (what she calls “scientizing” life) and Professor Allison Druin who introduced me to and immersed me in children-oriented participatory design methods. (I could hear the excited shouts and joyful exclamations of Kidsteam from my office, and I couldn’t resist finding out more, which fundamentally changed how I did research in STEM education.) My students never cease to surprise me, from 3D-printing gears to fix an aerial drone to developing an interactive sandbox that traces human movement using electromechanically controlled marbles to designing an e-textile shirt that senses and visualizes the wearer’s changing physiology via integrated anatomical models.

What are your hopes for the future?

As a graduate student, I recall being asked, “What are the biggest open questions in HCI, and how does your research work

toward addressing them?” I found this question both profoundly interesting and profoundly startling because it forced me to think about the most significant open areas in my field and to (somewhat uncomfortably) confront the relationship between this answer and my research. At the risk of sounding overly ambitious, I would like to adapt this question, which serves as a guiding principle for my research but that I also hope will inspire others: “What are the most significant societal challenges across the world? What role can computer

science, HCI, and design play in addressing those challenges? And where does your research/work fit?” As computation pervades nearly every aspect of our lives, I believe it is our role as technologists, designers, and practitioners to ask these questions of ourselves and to think about the political, economic, environmental, and social implications of our work. As a professor and educator, I am hopeful. This larger world-view framing of CS seems to resonate with younger generations and, I hope, will soon become the norm. ■

Chapter 13

INTERACTION DESIGN IN PRACTICE

13.1 Introduction

13.2 AgileUX

13.3 Design Patterns

13.4 Open Source Resources

13.5 Tools for Interaction Design

Objectives

The main goals of the chapter are to accomplish the following:

- Describe some of the key trends in practice related to interaction design.
- Enable you to discuss the place of UX design in agile development projects.
- Enable you to identify and critique interaction design patterns.
- Explain how open source and ready-made components can support interaction design.
- Explain how tools can support interaction design activities.

13.1 Introduction

As our interviewee at the end of Chapter 1, Harry Brignull, remarked, the field of interaction design changes rapidly. He says, “A good interaction designer has skills that work like expanding foam.” In other words, the practice of interaction design is quite messy, and keeping up with new techniques and developments is a constant goal. When placed within the wider world of commerce and business, interaction designers face a range of pressures, including restricted time and limited resources, and they need to work with people in a wide range of other roles, as well as stakeholders. In addition, the principles, techniques, and approaches introduced in other chapters of this book need to be translated into practice, that is, into real situations with sets of real users, and this creates its own pressures.

Many different names may be given to a practitioner conducting interaction design activities, including interface designer, information architect, experience designer, usability engineer, and user experience designer. In this chapter, we refer to *user experience designer* and *user experience design* because these are most commonly found in industry to describe someone who performs the range of interaction design tasks such as interface design, user evaluations, information architecture design, visual design, persona development, and prototyping.

Other chapters of this book may have given the impression that designers create their designs from scratch, with little or no help from anyone except users and immediate colleagues, but in practice, user experience (UX) designers draw on a range of support. Four main areas of support that impact the job of UX designers are described in this chapter.

- Working with software and product development teams operating an agile model of development (introduced in Chapter 2, “The Process of Interaction Design”) has led to technique and process adaptation, resulting in agileUX approaches.
- Reusing existing designs and concepts is valuable and time-saving. Interaction design and UX design patterns provide the blueprint for successful designs, utilizing previous work and saving time by avoiding “reinventing the wheel.”
- Reusable components—from screen widgets and source code libraries to full systems, and from motors and sensors to complete robots—can be modified and integrated to generate prototypes or full products. Design patterns embody an interaction idea, but reusable components provide implemented chunks of code or widgets.
- There is a wide range of tools and development environments available to support designers in developing visual designs, wireframes, interface sketches, interactive prototypes, and more.

This chapter introduces each of these four areas.

In this video, Kara Pernice suggests three challenges for UX in practice. It is available at <https://www.youtube.com/watch?v=qV5ILjmlL278>.

Here is a concrete view of what a UX designer does in practice:
<https://www.interaction-design.org/literature/article/7-ux-deliverables-what-will-i-be-making-as-a-ux-designer>

BOX 13.1

Technical Debt in UX

Technical debt is a term commonly used in software development, coined originally by Ward Cunningham in 1992, which refers to making technical compromises that are expedient in the short term but that create a technical context that increases complexity and cost in the long term. As with financial debt, technical debt is acceptable as a short-term approach to overcoming an immediate shortfall, provided that the debt will be repaid quickly. Leaving a debt for longer results in significant extra costs. Technical debt can be incurred unintentionally, but pressures associated with time and complexity also lead to design trade-offs that may prove to be expensive in the longer term.

UX debt is created much like technical debt in the sense that trade-offs are made for the needs of the project.

To address technical debt, a discipline of *refactoring* is needed, that is, correcting any pragmatic trade-offs quickly after the immediate pressure has receded. Significant difficulties arise if these trade-offs are not identified, understood, and corrected in a timely manner. Two interrelated situations can lead to significant user experience debt that is then extremely costly to correct.

- *If an organization did not, in the past, understand the value of good user experience design and products or software systems with poor user experiences persist.* This can be particularly prevalent for internal systems and products, where the drive for a good user experience is less acute than for externally marketed products that face more competition from other providers.
- *If an organization has a large portfolio of products, each of which was developed independently.* This can be the result of acquisitions and mergers of companies, each with their own UX brand, leading to a proliferation of designs.

In severe cases, UX debt can lead to the revamping of infrastructure and complete renewal of products. ■

For an interesting take on UX debt, see this article: <https://www.nngroup.com/articles/ux-debt>.

13.2 AgileUX

Since the rise of agile software development during the 2000s, UX designers have been concerned about the impact that it will have on their own work (Sharp et al., 2006), and the debate is ongoing (McInerney, 2017). *AgileUX* is the collective label given to efforts that aim to resolve these concerns by integrating techniques and processes from interaction design and those from agile methods. While agile software development and UX design have some characteristics in common such as iteration, a focus on measurable completion criteria, and user involvement, agileUX requires a reorganization and some rethinking of UX design activities and products. A recent reflection on the evolution of agileUX concluded that integrating agile and UX requires mutual team understanding across three dimensions, and those dimensions are variably understood (Da Silva et al., 2018): the “process and practice” dimension is understood; the “people and social” dimension is nearly understood; but the “technology and artifact” dimension—that is, use of technology to coordinate teams’ activities and artifacts to mediate teams’ communication—has yet to be properly understood. A key aspect is for agile development teams to understand that user experience design is not a role but is a discipline and mind-set. This account makes it clear that using agileUX in practice is far from straightforward. The key is to find a suitable balance that preserves both the research and reflection

needed for good UX design, as well as rapid iterations that incorporate user feedback and allow technical alternatives to be tested.

In a plan-driven (waterfall) software development process, requirements are specified as completely as possible before any implementation begins. In an agile software development process, requirements are specified only in enough detail for implementation to begin. Requirements are then elaborated as implementation proceeds, according to a set of priorities that change on a regular basis in response to changing business needs.

To integrate UX design into an agile workflow, it also needs to progress in a similar fashion. Reprioritization may happen as frequently as every two weeks, at the beginning of each iterative cycle. The shift from developing complete requirements up front to “just-in-time” or just enough requirements aims to reduce wasted effort, but it means that UX designers (along with their software engineer colleagues) have had to rethink their approach. All of the techniques and principles that UX designers use are just as relevant, but how much of each activity needs to be completed at what point in the iterative cycle and how the results of those activities feed into implementation need to be adjusted in an agile development context. This can be unsettling for designers, as the design artifacts are their main deliverable and hence may be viewed as finished, whereas for agile software engineers, they are consumables and will need to change as implementation progresses and requirements become elaborated.

Consider the group travel organizer example introduced in Chapter 11, and assume that it is being developed using agileUX. Four epics (large user stories) for the product are identified in Chapter 11, as follows:

1. As a <group traveler>, I want <to choose from a range of potential vacations that suit the group’s preferences> so that <the whole group can have a good time>.
2. As a <group traveler>, I want <to know the visa restrictions for everyone in the group> so that <visas can be arranged for everyone in the group in plenty of time>.
3. As a <group traveler>, I want <to know the vaccinations required to visit the chosen destination> so that <vaccinations can be arranged for everyone in the group in plenty of time>.
4. As a <travel agent>, I want <up-to-date information displayed> so that <my clients receive accurate information>.

At the beginning of the project, these epics will be prioritized, and the central goal of the product (to identify potential vacations) will be the top priority. This will then initially be the focus of development activities. To allow users to choose a vacation, epic 4, supporting the travel agent to update travel details, will also need to be implemented (otherwise travel details will be out of date), so this is also likely to be prioritized. Establishing the detailed requirements and the design of the other two areas will be postponed until after a product that allows users to choose a vacation has been delivered. Indeed, once this product is delivered, the customer may decide that offering help for vaccinations and visas does not result in sufficient business value for it to be included at all. In this case, referring users to other, more authoritative sources of information may be preferable.

Conducting UX activities within an agile framework requires a flexible point of view that focuses more on the end product as the deliverable than on the design artifacts as deliverables. It also requires cross-functional teams where specialists from a range of disciplines, including UX design and engineering, work closely together to evolve an understanding of both the users and their context, as well as the technical capabilities and practicalities of the technology. In particular, agileUX requires attention to three practices, each of which is elaborated in the following sections.

- What user research to conduct, how much, and when
- How to align UX design and agile working practices
- What documentation to produce, how much, and when



Source: Leo Cullum / Cartoon Stock

13.2.1 User Research

The term *user research* refers to the data collection and analysis activities necessary to characterize the users, their tasks, and the context of use before product development begins. Field studies and ethnography are often used in these investigations, but agile development works on short “timeboxes” of activity (up to four weeks in length, but often only two weeks in length) and hence does not support long periods of user research. (Different names are given by different agile methods to the iteration, or timeframe, the most common being *sprint*, *timebox*, and *cycle*.) Even a month to develop a set of personas or to conduct a detailed investigation into online purchasing habits (for example) is too long for some agile development cycles. User-focused activities evaluating elements of the design, or interviews to clarify requirements or task context, can be done alongside technical development (see the parallel tracks approach discussed in a moment), but planning to conduct extensive user research once iterative development starts will result in shallow user research, which is impossible to react to, as there just isn’t enough time.

One way to address this is for user research to be conducted before the project begins, or indeed before it is announced, as suggested by Don Norman (2006), who argues that it is better to be on the team that decides which project will be done at all, hence avoiding the constraints caused by limited timeboxes. This period is often called *iteration zero* (or Cycle 0, as you’ll see later in Figure 13.2), and it is used to achieve a range of up-front activities including software architecture design as well as user research.

Another approach to conducting user research for each project is to have an ongoing program of user research that revises and refines a company’s knowledge of their users over a longer time span. For example, Microsoft actively recruits users of their software to sign up and take part in user research that is used to inform future developments. In this case, the specific data gathering and analysis needed for one project would be conducted during iteration zero, but done in the context of a wider understanding of users and their goals.

ACTIVITY 13.1

Consider the “one-stop car shop” introduced in Activity 11.4. What kind of user research would be helpful to conduct before iterative development begins? Of these areas, which would be useful to conduct in an ongoing program?

Comment

Characterizing car drivers and the hybrid driving experience would be appropriate user research before iterative development begins. Although many people drive, the driving experience is different depending on the car itself and according to the individual’s capabilities and experiences. Collecting and analyzing suitable data to inform the product’s development is likely to take longer than the timebox constraints would allow. Such user research could develop a set of personas (maybe one set for each class of car) and a deeper understanding of the hybrid driving experience.

Car performance and handling is constantly evolving, however, and so an understanding of the driving experience would benefit from ongoing user research. ■

Lean UX (see Box 13.2) takes a different approach to user research by focusing on getting products into the market and capturing user feedback on products that are in the marketplace. It specifically focuses on designing and developing innovative products.

BOX 13.2

Lean UX (Adapted from Gothelf and Seiden (2016))

Lean UX is designed to create and deploy innovative products quickly. It is linked to agileUX because agile software development is one of its underlying philosophies and it champions the importance of providing a good user experience. Lean UX builds upon UX design, design thinking, agile software development, and the Lean Startup ideas (Ries, 2011). All four perspectives emphasize iterative development, collaboration between all stakeholders, and cross-functional teams.

Lean UX is based on tight iterations of build-measure-learn, a concept central to the lean startup idea, which in turn was inspired by the lean manufacturing process from Japan. The lean UX process is illustrated in Figure 13.1. It emphasizes waste reduction, the importance of experimentation to learn, and the need to articulate outcomes, assumptions, and hypotheses about a planned product. Moving the focus from outputs (for example, a new smartphone app) to outcomes (for example, more commercial activity through mobile channels) clarifies the aims of the project and provides metrics for defining success. The importance of identifying assumptions was discussed in Chapter 3, “Conceptualizing Interaction.” An example assumption might be that young people would rather use a smartphone app to access local event information than any other media. Assumptions can be expressed as hypotheses that can be put to the test more easily through research or by building a *minimum viable product* (MVP) that can be released to the user group.

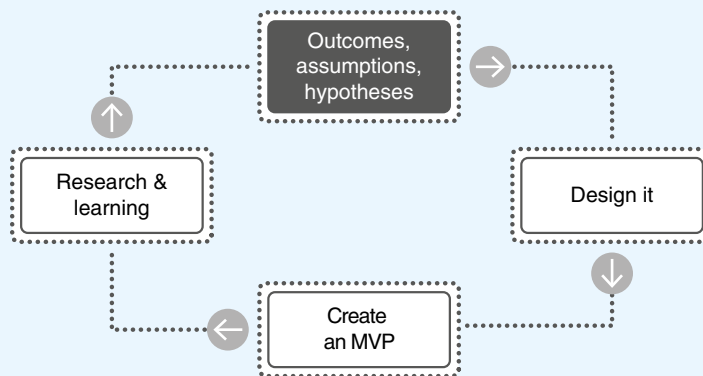


Figure 13.1 The Lean UX process

Source: Gothelf and Seiden (2016). Used courtesy of O'Reilly Media

Testing hypotheses, and hence assumptions, is done through experimentation, but before undertaking an experiment, the evidence required to confirm or refute each assumption needs to be characterized. An MVP is the smallest product that can be built that allows assumptions to be tested by giving it to a user group and seeing what happens. Experimentation and the evidence collected are therefore based on actual use of the product, and this allows the team to learn something.

As an example, Gothelf and Seiden (2016, pp. 76-77) describes an example of a company that wanted to launch a monthly newsletter. Their assumption was that a monthly newsletter would be attractive to their customers. To test this assumption, they spent half a day designing and coding a sign-up form on their website and collected evidence in the form of the number of sign-ups received. This form was an MVP that allowed them to collect evidence to support or refute their assumption, that is, that a monthly newsletter would be attractive to their customers. Having collected enough data, they planned to continue their experiments with further MVPs that experimented with formats and content for the newsletter. ■

In this video, Laura Klein explains Lean UX, at <http://youtu.be/7NkMm5WefBA>.

13.2.2 Aligning Work Practices

If requirements are specified before implementation begins, there is a tendency for designers to develop complete UX designs at the beginning of a project to ensure a coherent design throughout. In agile terms, this is referred to as *big design up front* (BDUF), and this is an anathema to agile working. Agile development emphasizes regular delivery of working software through evolutionary development and the elaboration of requirements as implementation proceeds. In this context, BDUF leads to practical problems since the reprioritization

of requirements means that interaction elements (features, workflows, and options) may no longer be needed or may require redesigning. To avoid unnecessary work on detailed design, UX design activities need to be conducted alongside and around agile iterations. The challenge is how to organize this so that a good user experience is achieved and the product vision is maintained (Kollman et al., 2009).

In response to this challenge, Miller (2006) and Sy (2007) proposed that UX design work is done one iteration ahead of development work in parallel tracks (see Figure 13.2). The parallel tracks approach to integrating UX design and agile processes originated at Alias—now part of Autodesk. Note that in this diagram, iteration is referred to as *Cycle*. The principle of parallel tracks development is quite simple: that design activity and user data collection for Cycle $n+1$ is performed during Cycle n . This enables the design work to be completed just ahead of development work, yet to be tightly coupled to it as the product evolves. Completing it much sooner than this can result in wasted effort, as the product and understanding about its use evolves.

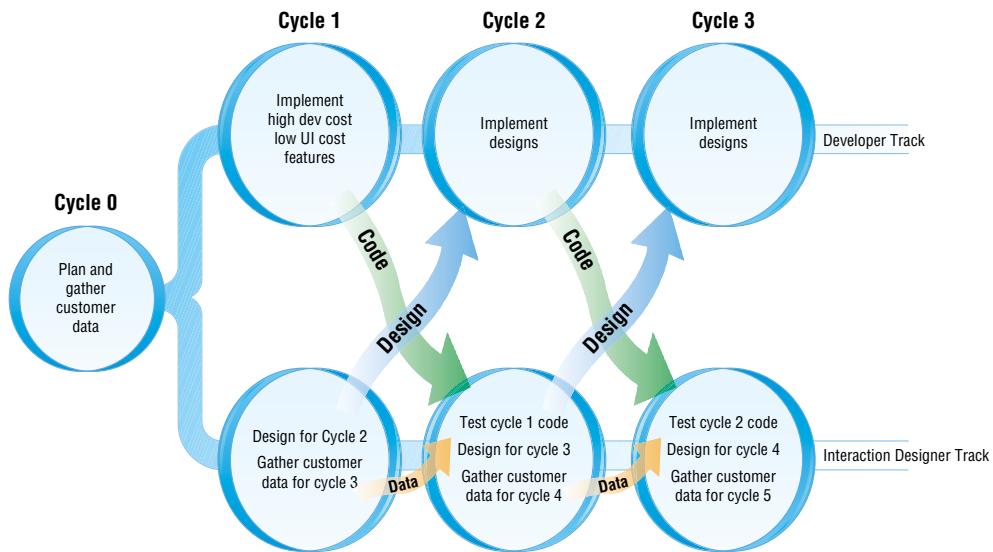


Figure 13.2 Cycle 0 and its relationship to later cycles

Source: Sy (2017)

Cycle 0 and Cycle 1 are different from subsequent cycles because, before evolutionary development can begin, the product vision needs to be created. This is handled in different ways in different agile methods, but all agree that there needs to be some kind of work up front to understand the product, its scope, and its overall design (both technical and UX). Some general data about customers and their behavior may have been collected before Cycle 0, but the vision and overall design is completed for the current project by the end of Cycle 0. The work required will depend on the nature of the product: whether it is a new version of an existing product, a new product, or a completely new experience. Cycle 0 can also be longer than other cycles to accommodate differing needs, but producing pixel-perfect designs of the product before evolutionary development starts is not the aim for Cycle 0.

One of the originators of the parallel tracks development idea, Desiree Sy (2007), explained this in the context of two different products. The first product is SketchBook Pro v2.0, a sophisticated sketching, annotating, and presentation tool to support digital artists. The second is Autodesk's Showcase which, though no longer available, was a real-time automotive 3D visualization product. For SketchBook Pro v2.0, the team conducted a survey of users who had downloaded v1.0 (a free trial version) but had not purchased v2.0. The results of the survey helped the team to refine 100 features into five major work streams, and this information informed development and prioritization throughout the development process. For Showcase, during Cycle 0, the team interviewed potential purchasers who performed work that the tool was going to be designed to support. This data formed the foundation for the design principles of the product as well as prioritization and design decisions as development progressed.

Cycle 1 usually involves technical setup activities in the developer track, which allows the UX designers to get started on the design and user activities for Cycle 2. For subsequent cycles, the team gets into a rhythm of design and user activities in Cycle $n-1$ and corresponding technical activity in Cycle n .

For example, imagine that development of a smartphone app to support attendance at a music festival is in Cycle n , and that Cycle n is scheduled to work on capturing reviews of the acts performing at the festival. During Cycle $n-1$, UX designers will have produced initial designs for capturing reviews of the acts by designing detailed icons, buttons, or other graphics, and prototyping different interaction types. During Cycle n , they will answer specific queries about these concrete designs, and they will revise them if necessary based on implementation feedback. Cycle n design work will be to develop concrete designs for the next cycle, which might be focusing on identifying and displaying reviews on demand. Also during Cycle n , UX designers will evaluate the implementation coming out of Cycle $n-1$. So, in any one cycle, UX designers are handling three different types of activity: evaluating implementations from the previous cycle, producing concrete designs for the next cycle, and answering queries on the designs being implemented in the current cycle.

The team at Alias found that the UX designers worked closely with the developers during design and implementation to make sure that they designed something that could be implemented and also what was indeed implemented was what had actually been designed. The interaction designers felt that there were three big advantages to this process. First, no design time was wasted on features that would not be implemented. Second, usability testing (for one set of features) and contextual inquiry (for the next set) could be done on the same customer visit, thus saving time. Third, the interaction designers received timely feedback from all sides—both users and developers. More importantly, they had time to react to that feedback because of the agile way of working. For example, the schedule could be changed if something was going to take longer to develop than first thought, or a feature could be dropped if it became apparent from the users that something else had higher priority. In summary, “Agile user-centered design resulted in better-designed software than waterfall user-centered design” (Sy, 2007, p. 130).

These advantages have been realized by others too, and this parallel tracks way of working has become a popular way to implement agileUX. Sometimes, the UX designers work two iterations ahead, depending on the work to be done, the length of the iteration, and external factors such as time required to obtain appropriate user input. Working in this way does not diminish the need for UX designers and other team members to collaborate closely together, and although the tracks are parallel, they should not be seen as separate processes. This does, however, raise a dilemma, as discussed in the Dilemma box.

DILEMMA

To Co-locate or Not to Co-locate, That Is the Question

UX designers in most large organizations are not numerous enough to have one UX designer for every team, so where should the UX designer be located? Agile development emphasizes regular communication and the importance of being informed about the project as it evolves. Hence, it would be good for the UX designer to be located with the rest of the team. But which team? Maybe a different agile team every day? Or each team for one iteration? Some organizations, however, believe that it is better for UX designers to sit together in order to provide discipline coherence: “UX designers work best when they are separated from the issues of software construction because these issues hamper creativity” (Ferreira et al., 2011). Indeed, this view is shared by some UX designers. If you, as a UX designer, were part of several agile teams, needing to engage with each of them, where would you prefer to be located? What might be the advantages and disadvantages of each, or maybe using a social awareness tool such those introduced in Chapter 5, “Social Interaction,” would be more appropriate? ■

This video describes some case studies on the UX techniques used by Android within agile iterations: <http://youtu.be/6MOeVNbh9cY>.

ACTIVITY 13.2

Compare Lean UX, agileUX, and evolutionary prototyping (introduced in Chapter 12, “Design, Prototyping, and Construction”). In what ways are they similar and how do they differ?

Comment

Lean UX produces a MVP to test assumptions by releasing it to users as a finished product and collecting evidence of users’ reactions. This evidence is then used to evolve subsequent (larger) products based on the results of this experimentation. In this sense, Lean UX is a form of evolutionary development, and it has similarities with evolutionary prototyping. However, not all the MVPs developed to test assumptions may be incorporated into the final product, just the results of the experiment.

AgileUX is an umbrella term for all efforts that focus on integrating UX design with agile development. Agile software development is an evolutionary approach to development, and hence agileUX is also evolutionary. Additionally, agileUX projects can employ prototyping to answer questions and test ideas as with any other approach, as described in Chapter 12. ■

13.2.3 Documentation

The most common way for UX designers to capture and communicate their design has been through documentation, for instance, user research results and resulting personas, detailed interface sketches, and wireframes. Because UX designers view the design as their main deliverable, a key indicator that their work is ready for sign-off is the availability of comprehensive documentation to show that their goals have been achieved. This may include other forms of design capture, such as prototypes and simulations, but documentation is still common. Agile development encourages only minimal documentation so that more time can be spent on design, thus producing value to the user via a working product.

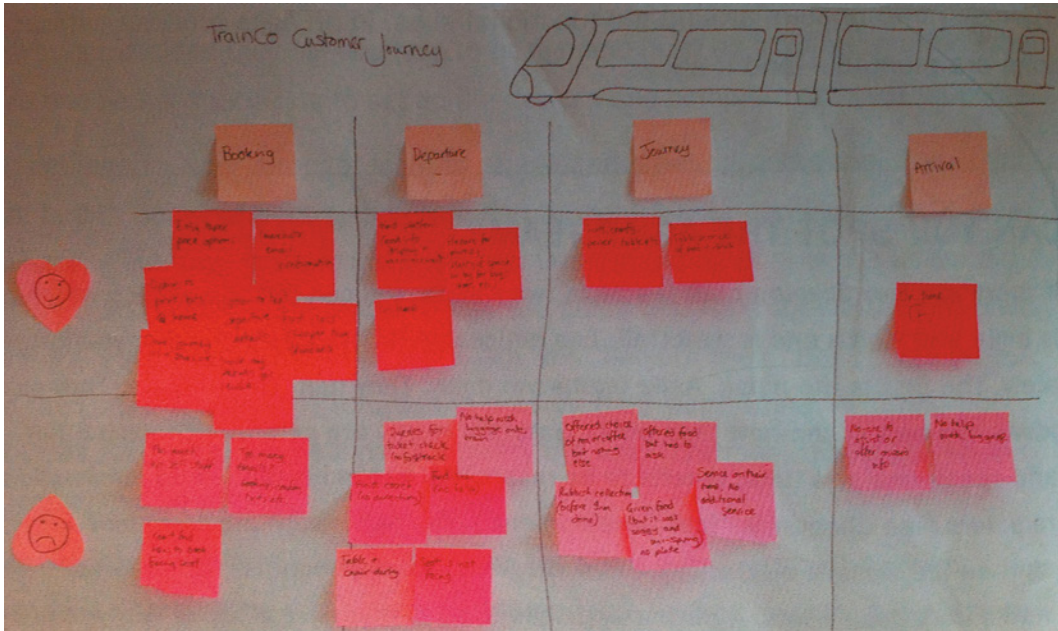
Minimal documentation does not mean “no documentation,” and some documentation is desirable in most projects. However, a key principle in agileUX is that documentation should not replace communication and collaboration. To help identify the right level of documentation, Lindsay Ratcliffe and Marc McNeill (2012, p. 29) suggest asking a set of questions of any documentation process.

1. How much time do you spend on documentation? Try to decrease the amount of time spent on documentation and increase design time.
2. Who uses the documentation?
3. What is the minimum that customers need from the documentation?
4. How efficient is your sign-off process? How much time is spent waiting for documentation to be approved? What impact does this have on the project?
5. What evidence is there of document duplication? Are different parts of the business documenting the same things?
6. If documentation is only for the purpose of communication or development, how polished does it need to be?

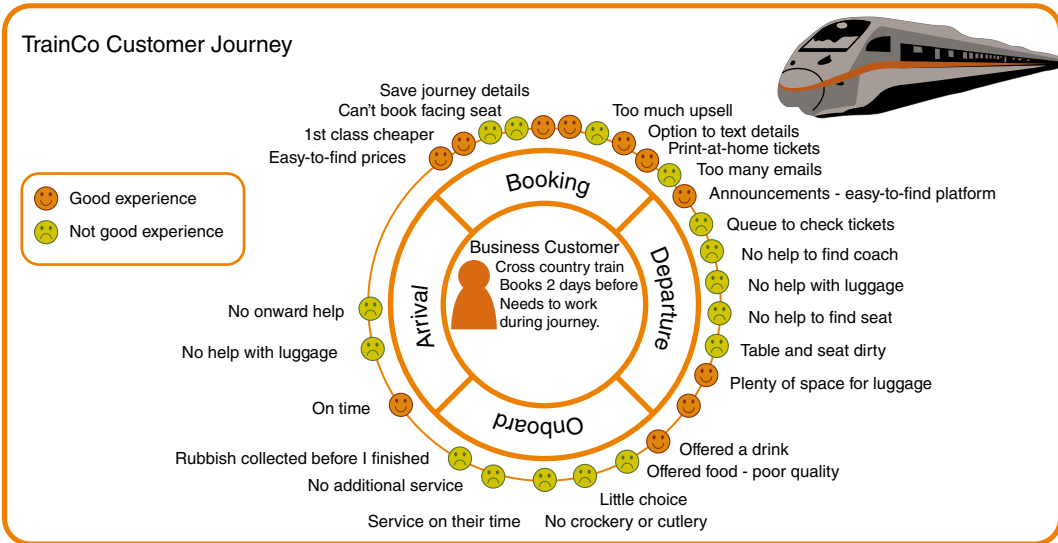
They also use the example in Figure 13.3 to illustrate these points. Both images capture a user journey, that is, one path a user might take through the product. The sketch in Figure 13.3(a) is constructed with sticky notes and string, and it was generated by all of the team members during a discussion. The sketch in Figure 13.3(b) took hours of designer time to draw and polish. It looks good, but that time could have been used to design the product rather than the user journey sketch.

The question of how much documentation is needed in an agile project is not limited to agileUX. Scott Ambler (Ambler, 2002) provides a detailed description of best practices for agile documentation. These support the production of “good enough” documentation in an efficient way, and they are intended to determine what documentation is needed. He proposes questions such as these:

- What is the purpose of the documentation?
- Who is the customer of the documentation?
- When should documents be updated?”



(a)



(b)

Figure 13.3 (a) A low-fidelity user journey, (b) a high-fidelity user journey

Source: Ratcliffe and McNeill (2012)

DILEMMA

Quick, Quick, Slow?

One of the challenges for UX practice is how best to integrate with software and product development conducted using an agile approach. Taking an agile approach is seen as beneficial for a range of reasons, including an emphasis on producing something of use, customer (and user) collaboration, rapid feedback, and minimal documentation—only areas of the product that are definitely going to be implemented are designed in detail. However, focusing on short timeboxes can lead to an impression that everything is being rushed. Creating an appropriate balance between short timeboxes and a reflective design process requires careful planning so that important aspects of UX design are not rushed.

Slow design is part of the *slow movement*, which advocates a cultural shift toward slowing down the pace of life (Grosse-Hering et al., 2013). The main intent of *slow design* is to focus on promoting well-being for individuals, society, and the natural environment by promoting the design of products that are long-lived and sustainable. Working more slowly does not, *per se*, address the impression of rushing, but slow design also emphasizes the importance of providing time to reflect and think, for the user to engage and create their own products, and for products and their use to evolve over time.

The agile movement is here to stay, but the importance of taking time to reflect and think, when necessary, and not rushing to make decisions remains. The dilemma here is finding the right balance between rapid feedback to identify solutions that work and providing the time to stop and reflect. ■

CASE STUDY 13.1

Integrating UX Design into a Dynamic Systems Development Method Project

Challenges, Working Practices, and Lessons Learned

This case study presents a portion of one organization's journey to integrate UX design into one agile software development approach: the Dynamic Systems Development Method (DSDM) framework (see <https://www.agilebusiness.org/what-is-dsdm> for more details). It describes the difficulties they faced, the working practices adopted, and the lessons learned from their experiences of integrating UX designers into their DSDM agile process.

LShift is a high-tech software development company that works across a broad range of industries, languages, and platforms. The company faced four main challenges while integrating UX design into the DSDM framework.

- *Communication between developers and UX designers:* What is the relevant information that needs to be communicated, how best to communicate it, how to keep communication channels open, and how to keep the emerging design implementation visible for feedback. Difficulties in these areas can cause frustration, problems with the technical feasibility of design solutions, and mistaken expectations by the client.

- *Level of precision in up-front design:* Developers suggested five main reasons why “less is more” when it comes to design documentation ready for the start of developer involvement.
 - Prioritization and de-scoping can lead to a waste of pixel-perfect designs.
 - Some design issues will be found only once you start implementing.
 - Pixel-perfect designs may increase resistance to making design changes.
 - It is better to focus on functionality first and design as you go along.
 - Quality of designs can benefit from early input by developers.
- *Design documentation:* The amount and detail of documentation needs to be discussed early on so that it meets both developers’ and designers’ requirements.
- *User testing:* User testing can be a challenge in a product development setting if the business does not yet have any customers. This can be addressed at least partially using personas and user representatives.

This case study describes the background to these challenges, provides more detail about these challenges, and introduces some practices that the company used to address them. The case study is available in full at <http://tinyurl.com/nee nbk>. ■

13.3 Design Patterns

Design patterns capture design experience, but they have a different structure and a different philosophy from other forms of guidance or specific methods. One of the intentions of the patterns community is to create a vocabulary based on the names of the patterns, which designers can use to communicate with one another and with users. Another is to produce literature in the field that documents experience in a compelling form.

The idea of patterns was first proposed by the architect Christopher Alexander, who described patterns in architecture (Alexander, 1979). His hope was to capture the “quality without a name” that is recognizable in something when you know it is good.

But what is a design pattern? One simple definition is that it is a solution to a problem in a context; that is, a pattern describes a problem, a solution, and where this solution has been found to work. Users of the pattern can therefore not only see the problem and solution but can also understand the circumstances under which the idea has worked before and access a rationale for why it worked. A key characteristic of design patterns is that they are generative; that is, they can be instantiated or implemented in many different ways. The application of patterns to interaction design has grown steadily since the late 1990s (for instance, Borchers, 2001; Tidwell, 2006; Crumlish and Malone, 2009) and have continued to be actively developed (for example, Josh et al., 2017).

Patterns on their own are interesting, but they are not as powerful as a pattern language. A *pattern language* is a network of patterns that reference one another and work together to create a complete structure. Pattern languages are not common in interaction design, but there are several pattern collections, that is, sets of patterns that are independent of each other. Patterns are attractive to designers because they are tried and tested solutions to common problems. It is common (although not obligatory) for pattern collections to be associated with software components that can be used with little modification, and as they are common solutions, many users are already familiar with them, which is a great advantage for a new app or product on the market. See Box 13.3 for an example pattern: Swiss Army Knife Navigation.

BOX 13.3

Swiss Army Knife Navigation: An Example Design Pattern for Mobile Devices (Nudelman (2013). Used courtesy of John Wiley & Sons, Inc.)

The principle behind the Swiss Army Knife Navigation design pattern is to maximize productive use of the screen space and keep users engaged in the content of what they are doing. For example, in a game design, the user does not want to be side-tracked by navigation bars and menu pop-ups. Having a mechanism that allows the controls to fade in and out of view is a much more engaging design.

This design pattern is commonly instantiated as “off-canvas” or “side drawer” navigation (Neil, 2014), where the control bar slides in, overlaying the main screen contents. It is useful because it is a “transient” navigation bar and takes up screen space only temporarily; that is, it can be swiped in over the top of the main app screen and then swiped back once the user has finished the action. It is good from an interaction design point of view because it supports the use of both text and icons to represent actions. It is also good from the point of view of screen layout because it takes up space only when the menu is needed. It can also be used for interactions other than navigation (Peatt, 2014). This is an example of a common design pattern that is also evolving in a range of different directions. Exactly how this navigation bar is implemented varies with different platforms.

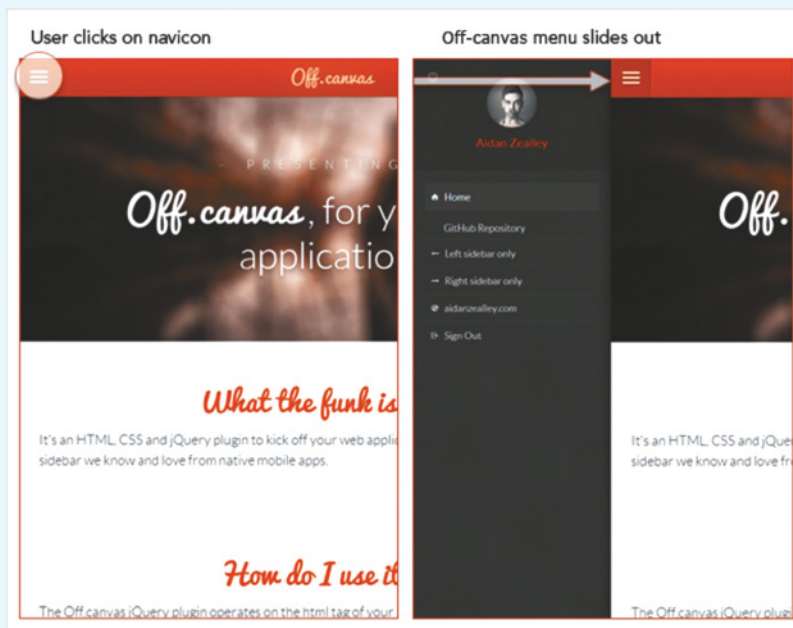


Figure 13.4 Example of Swiss Army Knife Navigation pattern, instantiated as off-canvas navigation
Source: Used courtesy of Aidan Zealley

In figure 13.4, the left image, the menu is represented as a list of lines at the top left of the screen, while in the right image, the menu items have pushed the user view to the right. ■

Examples of interaction design guidelines and pattern libraries plus downloadable collections of screen elements are available at:

Windows: <https://developer.microsoft.com/en-us/windows/apps/design>

Mac: <https://developer.apple.com/design/human-interface-guidelines/>

General UI Design Patterns: <https://www.interaction-design.org/literature/article/10-great-sites-for-ui-design-patterns>

Google Material Design: <https://design.google/>

Pattern collections, libraries, and galleries relevant to interaction design are commonly used in practice (for instance, Nudelman, 2013) and are often accompanied by code snippets available through open source repositories such as GitHub (<https://github.com/>) or through platform websites such as <https://developer.apple.com/library/iOS/documentation/userexperience/conceptual/mobilehig/> for iOS on an iPhone.

Patterns are a “work in progress,” because they continue to evolve as more people use them, experience increases, and users’ preferences change. Patterns can continue to evolve for some time, but they can also be deprecated, that is, become outdated and no longer considered good interaction design. Reusing ideas that have proved to be successful in the past is a good strategy in general, particularly as a starting point, but it should not be used blindly. In custom applications, the design team may also create their own libraries. As with many areas of design, there is disagreement about which patterns are current and which are outdated.

For a humorous discussion of the hamburger icon, see:

<https://icons8.com/articles/most-hated-ui-ux-design-pattern/>

For a discussion of the tab bar, see:

<https://uxplanet.org/tab-bars-are-the-new-hamburger-menus-9138891e98f4>

ACTIVITY 13.3

One design pattern for mobile devices that is deprecated by some and not others is the Carousel navigation pattern, in which the user is presented with several images (of products, for example) horizontally across the screen, or one at a time in the same screen location. Swiping (or clicking) left or right displays other images, just like a carousel.

This design pattern has provoked different reactions by different designers. Search for information on this design pattern using your favorite browser and read at least two articles or blog posts about it: one arguing that it should be deprecated and one that explains how it can be used successfully. Decide for yourself whether the Carousel pattern should be labeled outdated or kept alive.



(a)



(b)

Figure 13.5 Two example carousel navigation styles (a) shows pictures of a house for sale. Note the arrows to the left and right of the row of photos at the bottom. (b) shows a weather application for a mobile phone that can be swiped left and right for other locations. Note the line of dashes in the bottom middle of the screen that indicate there are other screens.

Source: Helen Sharp

Comment

The Nielsen Norman Group has two articles about the carousel on its website (see the URLs at the end of this paragraph).

One presents evidence from a usability trial with one user that shows carousels can fail, and the other presents a balanced view of how to design a good carousel. This second article focuses on the version of the carousel where several images are displayed at the same location on the screen, one at a time. They identify the greatest advantages as being the good use of screen space (because several elements occupy the same space) and that having information at the top of the screen means that visitors are more likely to see it. Disadvantages include that users often navigate past the carousel and that even if users do see the image, it is usually only the first one. The article does suggest using an alternative design, and it goes on to provide some useful examples and guidelines for good carousels.

There is a thread of posts and articles arguing that the carousel should not be used. These also point to evidence that users rarely use the carousel, but if they do they focus only on the first image. Nevertheless, there seems to be no solid set of data to support or refute the usability of the carousel in all of its various forms.

On balance, it seems that some forms of carousel meet the product's goals more readily than others, for example, because only the first image in a series is viewed by most users. Assuming appropriate design and, maybe more importantly, an evaluation with potential users and your content, it seems plausible that the carousel navigation pattern is not yet ready to be deprecated. ■

www.nngroup.com/articles/designing-effective-carousels/
www.nngroup.com/articles/auto-forwarding/

Design patterns are a distillation of previous common practice, but one of the problems with common practice is that it is not necessarily good practice. Design approaches that represent poor practice are referred to as *anti-patterns*. The quality of interaction design and user experience in general has improved immensely since the first edition of this book in 2002, so why are anti-patterns still a problem? Basically, the technology is changing and design solutions that work on one platform don't necessarily work on another. A common source of anti-patterns for mobile devices is where websites or other software have been migrated from a large screen, such as a laptop, to a smartphone. One example of this is the untappable phone number that displays on a smartphone pop-up (see Figure 13.6).

Another kind of pattern that was introduced in Chapter 1 (see Figure 1.10) is the *dark pattern*. Dark patterns are not necessarily poor design, but they have been designed carefully to trick people, championing stakeholder value over user value, for instance. Some apparent dark patterns are just mistakes, in which case they will be corrected relatively quickly once identified. However, when a UX designer's knowledge of human behavior is deliberately used to implement deceptive functionality that is not in the user's best interests, that is a dark pattern. Colin Gray et al. (2018) collated and analyzed a set of 118 dark pattern examples identified by practitioners and identified five strategies: nagging, obstruction, sneaking, interface interference, and forced action.




Figure 13.6 An untappable phone number for help when smartphone installation goes wrong

13.4 Open Source Resources

Open source software refers to source code for components, frameworks, or whole systems that is available for reuse or modification free of charge. Open source development is a community-driven endeavor in which individuals produce, maintain, and enhance code, which is then given back to the community through an open source repository for further development and use. The community of open source committers (that is, those who write and maintain this software) are mostly software developers who give their time for free. The components are available for (re)use under software licenses that allow anyone to use and modify the software for their own requirements without the standard copyright restrictions.

Many large pieces of software underlying our global digital infrastructure are powered by open source projects. For example, the operating system Linux, the development environment Eclipse, and the NetBeans development tools are all examples of open source software.

Perhaps more interesting for interaction designers is that there is a growing proportion of open source software available for designing good user experiences. The design pattern implementation libraries introduced in section 13.3 are but one example of how open source software is affecting user experience design. Another example is the Bootstrap framework for front-end web development, released as open source in August 2011 and actively updated on a regular basis; see Figure 13.7 for an example of its use. This framework contains reusable code snippets, a screen layout grid that supports multiple screen sizes, and pattern libraries that include predefined sets of navigational patterns, typefaces, buttons, tabs, and so on. The framework and documentation are available through the GitHub open source repository (<https://github.com/twbs/bootstrap#community>).

Open source resources require a suitable repository, that is, somewhere for the source code to be stored and made accessible to others. More than this, the repository needs to serve a huge number of users (GitHub was reported to have 31 million users in 2018) who will want to build, review, modify, and extend software products. Managing this level of activity also requires version control, such as a mechanism that retains and can reinstate previous versions of the software. For example, GitHub is based on the version control system called Git. Communities form around these repositories, and submitting code to a repository requires an account. For example, each developer on GitHub can set up a profile that will keep track of their activity for others to see and comment upon.



Figure 13.7 An example website built using the Bootstrap framework

Source: Didier Garcia/Larson Associates.

Most repositories support both public and private spaces. Submitting code to a public space means that anyone in the community can see and download the code, but in a private space the source will be “closed.” One of the advantages of putting code on an open source repository is that many eyes can see, use, and modify your work—spotting security vulnerabilities or inefficient coding practices as well as contributing to, extending, or improving its functionality. Other popular open source repositories are BitBucket, Team Foundation Server, and GitLab.

Some advantages of using GitHub over other source repositories are discussed here:

<https://www.thebalancecareers.com/what-is-github-and-why-should-i-use-it-2071946>

The GitHub repository itself may look a little daunting for those who first come across it, but there is a community of developers behind it who are happy to help and support newcomers.

An introduction to using GitHub is available here:

<https://product.hubspot.com/blog/git-and-github-tutorial-for-beginners>

13.5 Tools for Interaction Design

Many types of digital tools are used in practice by UX designers, and the tooling landscape changes all the time (Putnam et al., 2018). These tools support creative thinking, design sketching, simulation, video capture, automatic consistency checking, brainstorming, library search, and mind mapping. In fact, any aspect of the design process will have at least one associated support tool. For example, Microsoft Visio and OmniGraffle support the creation of a wide range of drawings and screen layouts, while FreeMind is an open source, mind-mapping tool. In and of themselves, these tools provide significant support for UX design, but they can also work together to speed up the process of creating prototypes of various levels of fidelity.

Elsewhere in this book, we have emphasized the value of low-fidelity prototyping and its use in getting user feedback. As with any prototype, however, paper-based prototypes have their limitations, and they do not support user-driven interaction (see, for example, Lim et al., 2006). In recognition of this, developing interactive, low-fidelity prototypes has been investigated through research for many years (see Lin et al., 2000, or Segura et al., 2012). In recent years, tools to support the creation of interactive prototypes from static graphical elements have become available commercially. For example, overflow.io supports the production of playable user flow diagrams.

Commercial packages that support the quick and easy development of interactive wireframes, or mock-ups, are widely used in practice for demonstration and evaluation. Some commonly used tools are Balsamiq® (<https://balsamiq.com/>), Axure RP (<https://www.axure.com/>), and Sketch (<https://sketchapp.com/>). Activity 13.4 invites you to try one or more of the tools available to create a simple prototype.

Tools available for UX designers, many of which have free trial versions and tutorials, are available at the following links:

<https://support.balsamiq.com/tutorials/>

www.axure.com/learn/core/getting-started

<https://www.digitalartsonline.co.uk/features/interactive-design/16-best-ux-tools-for-designers/>

<https://blog.prototypr.io/meet-overflow-9b2d926b6093>

Having created an interactive wireframe using one of these tools, it is then possible to generate a higher-fidelity prototype by implementing the next prototype using a ready-made pattern library or framework, like those introduced in section 13.3 and section 13.4, to provide a coherent look and feel. This means going from a low-fidelity mockup to a working, styled prototype in one step. Other open source resources can also be used to provide a wider choice of interface elements or design components with which to create the product.

Paper-based prototypes are also not very good if technical performance issues such as component interfaces need to be prototyped—software-based prototypes are better. For example, Gosper et al. (2011) describes how, at SAP, employees often use a drawing or graphics package to mock up key use cases and their interfaces, interactions, and task flows and then output that to PowerPoint. This creates a set of slides that can be viewed to give an

overall sense of a user session. However, when they developed a business intelligence tool with key performance and “backend” implications, this form of prototyping was not sufficient for them to assess their product goals. Instead, the UX designer worked with a developer who prototyped some of the elements in Java.

ACTIVITY 13.4

Choose one of the commercially available tools that supports the generation of interactive wireframes or low-fidelity prototypes and generate a wireframe for a simple app, for instance, one that allows visitors to a local music festival to review the acts. Explore the different features offered by the tool, and note those that were particularly useful. Unless your employer or university has a license for these tools, you may not have access to all of their features.

Comment

We used the moqups web demo at <https://app.moqups.com/edit/page/ad64222d5> to create the design in Figure 13.8. Two features of note for this tool are (1) that it appears to be similar to other generic graphics packages, and hence it was easy to get started and produce an initial wireframe, and (2) the ruler settings automatically allowed precise positioning of interface elements in relation to each other. ■

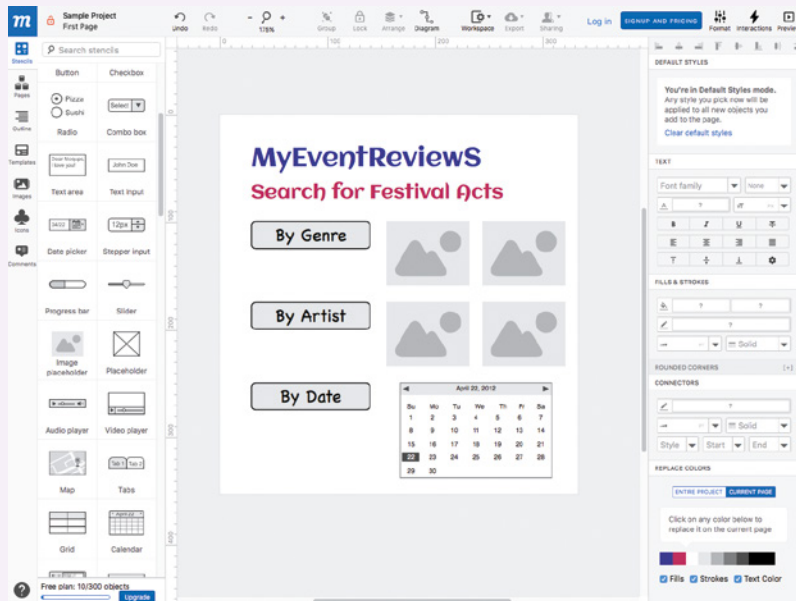


Figure 13.8 A screenshot of our interactive wireframe for the event review app, generated using moqups

Source: <https://moqups.com>

In-Depth Activity

This in-depth activity continues the work begun on the booking facility introduced at the end of Chapter 11.

1. Assume that you will produce the online booking facility using an agile approach.
 - a. Suggest the type of user research to conduct before iteration cycles begin.
 - b. Prioritize requirements for the product according to business value, in particular, which requirements are likely to provide the greatest business benefit, and sketch out the UX design work you would expect to undertake during the first four iteration cycles, that is, Cycle 0 and Cycles 1 to 3.
2. Using one of the mock-up tools introduced, generate a mock-up of the product's initial interface, as developed in the assignment for Chapter 12.
3. Using one of the patterns websites listed previously, identify suitable interaction patterns for elements of the product and develop a software-based prototype that incorporates all of the feedback and the results of the user experience mapping achieved at the end of Chapter 12. If you do not have experience in using any of these, create a few HTML web pages to represent the basic structure of the product.

Summary

This chapter explored some of the issues faced when interaction design is carried out in practice. The move toward agile development has led to a rethinking of how UX design techniques and methods may be integrated into and around agile's tight iterations. The existence of pattern and code libraries, together with open source components and automated tools, means that interactive prototypes with a coherent and consistent design can be generated quickly and easily, ready for demonstration and evaluation.

Key Points

- AgileUX refers to approaches that integrate UX design activities with an agile approach to product development.
- A move to agileUX requires a change in mind-set because of repeated reprioritization of requirements and short timeboxed implementation, which seeks to avoid wasted effort.
- AgileUX requires a rethinking of UX design activities: when to perform them, how much detail to undertake and when, and how to feed back results into implementation cycles.
- Design patterns present a solution to a problem in a context, and there are many UX design pattern libraries available.
- Dark patterns are designed to trick users into making choices that have unexpected consequences, for instance, by automatically signing them up for marketing newsletters.
- Open source resources, such as those on GitHub, make the development of standard applications and libraries with consistent interfaces easier and quicker.
- A variety of digital tools to support interaction design in practice are available.

Further Reading

GOTHELF, J., and SEIDEN, J. (2016) *Lean UX: Designing Great Products with Agile Teams* (2nd ed.), O'Reilly. This book focuses on the lean UX approach to development (see Box 13.2), but it also includes a wide range of practical examples and experiences from readers of the first edition of the book as to how agile development and UX design can work well together.

KRUCHTEN, P., NORD, R.L. and OZKAYA, I. (2012) “Technical Debt: From Metaphor to Theory and Practice,” *IEEE Software*, November/December (2nd ed.). This paper is the editors' introduction to a special issue on technical debt. This topic has been largely discussed and written about in the context of software development, with very little mention of interaction design or UX debt. However, these issues are relevant to interaction design practice today, and this paper provides an accessible starting point for anyone wanting to investigate this area further.

PUTNAM, C., BUNGUM, M., SPINNER, D., PARELKAR, A.N., VIPPARTI, S. and CASS, P. (2018), “How User Experience Is Practiced: Two Case Studies from the Field,” *Proceedings of CHI 2018*. This short paper provides some useful insights into UX practice based on two case studies from consumer-facing companies in Illinois.

RAYMOND, E.S. (2001) *The Cathedral and the Bazaar*. O'Reilly. This seminal book is a set of essays introducing the open source movement.

SANDERS, L. and STAPPERS, P. J. (2014) “From Designing to Co-Designing to Collective Dreaming: Three Slices in Time,” *interactions*, Nov–Dec, p. 25–33. This provides a fascinating account of the changes in design practice over the last 30 years, a reflection on what design practice is like in 2014, and then a projection into the future to see what design practice may be like 30 years from now. It considers the role of the customer and the designer and how the object being designed emerges from the design process.

SY, D. (2007) “Adapting Usability Investigations for Development,” *Journal of Usability Studies* 2(3), May, 112–130. This short paper is a good introduction to some of the key issues faced when trying to perform UX design alongside an agile project. It describes the well-established dual-track process model for agileUX, shown in Figure 13.2.

Chapter 14

INTRODUCING EVALUATION

14.1 Introduction

14.2 The Why, What, Where, and When of Evaluation

14.3 Types of Evaluation

14.4 Evaluation Case Studies

14.5 What Did We Learn from the Case Studies?

14.6 Other Issues to Consider When Doing Evaluation

Objectives

The main goals of this chapter are to accomplish the following:

- Explain the key concepts and terms used in evaluation.
- Introduce a range of different types of evaluation methods.
- Show how different evaluation methods are used for different purposes at different stages of the design process and in different contexts of use.
- Show how evaluation methods are mixed and modified to meet the demands of evaluating novel systems.
- Discuss some of the practical challenges of doing evaluation.
- Illustrate through short case studies how methods discussed in more depth in Chapters 8, 9, and 10 are used in evaluation and describe some methods that are specific to evaluation.
- Provide an overview of methods that are discussed in detail in the next two chapters.

14.1 Introduction

Imagine that you designed an app for teenagers to share music, gossip, and photos. You prototyped your first design and implemented the core functionality. How would you find out whether it would appeal to them and whether they will use it? You would need to evaluate it—but how? This chapter presents an introduction to the main types of evaluation and the methods that you can use to evaluate design prototypes and design concepts.

Evaluation is integral to the design process. It involves collecting and analyzing data about users' or potential users' experiences when interacting with a design artifact such as a screen sketch, prototype, app, computer system, or component of a computer system. A central goal of evaluation is to improve the artifact's design. Evaluation focuses on both the usability of the system (that is, how easy it is to learn and to use) and on the users' experiences when interacting with it (for example, how satisfying, enjoyable, or motivating the interaction is).

Devices such as smartphones, iPads, and e-readers, together with the pervasiveness of mobile apps and the emergence of IoT devices, have heightened awareness about usability and interaction design. However, many designers still assume that if they and their colleagues can use a product and find it attractive, others will too. The problem with this assumption is that designers may then design only for themselves. Evaluation enables them to check that their design is appropriate and acceptable for the target user population.

There are many different evaluation methods. Which to use depends on the goals of the evaluation. Evaluations can occur in a range of places such as in labs, people's homes, outdoors, and work settings. Evaluations usually involve observing participants and measuring their performance during usability testing, experiments, or field studies in order to evaluate the design or design concept. There are other methods, however, that do not involve participants directly, such as modeling users' behavior and analytics. Modeling users' behavior provides an approximation of what users might do when interacting with an interface; these models are often done as a quick way of assessing the potential of different interface configurations. Analytics provide a way of examining the performance of an already existing product, such as a website, so that it can be improved. The level of control on what is evaluated varies; sometimes there is none, such as for studies in the wild, and in others there is considerable control over which tasks are performed and the context, such as in experiments.

In this chapter, we discuss why evaluation is important, what needs to be evaluated, where evaluation should take place, and when in the product lifecycle evaluation is needed. Some examples of different types of evaluation studies are then illustrated by short case studies.

14.2 The Why, What, Where, and When of Evaluation

Conducting evaluations involves understanding not only why evaluation is important but also what aspects to evaluate, where evaluation should take place, and when to evaluate.

14.2.1 Why Evaluate?

User experience involves all aspects of the user's interaction with the product. Nowadays users expect much more than just a usable system—they also look for a pleasing and engaging experience from more products. Simplicity and elegance are valued so that the product is a joy to own and use.

From a business and marketing perspective, well-designed products sell. Hence, there are good reasons for companies to invest in evaluating the design of products. Designers can focus on real problems and the needs of different user groups and make informed decisions about the design, rather than on debating what each other likes or dislikes. It also enables problems to be fixed before the product goes on sale.

© 1999 Randy Glasbergen.



“it’s the latest innovation in office safety.
When your computer crashes, an air bag is activated
so you won’t bang your head in frustration.”

Source: © Glasbergen. Reproduced with permission of Glasbergen
Cartoon Service

ACTIVITY 14.1

Identify two adults and two teenagers prepared to talk with you about their Facebook usage (these may be family members or friends). Ask them questions such as these: How often do you look at Facebook each day? How many photos do you post? What kind of photos do you have in your albums? What kind of photo do you have as your profile picture? How often do you change it? How many Facebook friends do you have? What books and music do you list? Are you a member of any groups?

Comment

As you may know, some teenagers are leaving Facebook, while some adults, often parents and grandparents, continue to be avid users, so you may have had to approach several teenagers before finding two who would be worth talking to. Having found people who use Facebook, you probably heard about different patterns of use between the adults and the teenagers. Teenagers are more likely to upload selfies and photos of places they have just visited on sites such as Instagram or that they sent to friends on WhatsApp. Adults tend to spend time discussing family issues, the latest trends in fashion, news, and politics. They frequently post pictures and descriptions for family members and friends about where they went on vacation or of their children and grandchildren. After doing this activity, you should be aware that different kinds of users may use the same software in different ways. It is therefore important to include a range of different types of users in your evaluations. ■

14.2.2 What to Evaluate

What to evaluate ranges from low-tech prototypes to complete systems, from a particular screen function to the whole workflow, and from aesthetic design to safety features. Developers of a new web browser may want to know whether users find items faster with their product. Developers of an ambient display may be interested in whether it changes people’s

behavior. Game app developers will want to know how engaging and fun their games are compared with those of their competitors and how long users will play them. Government authorities may ask if a computerized system for controlling traffic lights results in fewer accidents or if a website complies with the standards required for users with disabilities. Makers of a toy may ask whether 6-year-olds can manipulate the controls, whether they are engaged by its furry cover, and whether the toy is safe for children. A company that develops personal, digital music players may want to know whether people from different age groups and living in different countries like the size, color, and shape of the casing. A software company may want to assess market reaction to its new home page design. A developer of smartphone apps for promoting environmental sustainability in the home may want to know if their designs are enticing and whether users continue to use their app. Different types of evaluations will be needed depending on the type of product, the prototype or design concept, and the value of the evaluation to the designers, developers, and users. In the end, the main criteria are whether the design does what the users need and want it to do; that is, will they use it?

ACTIVITY 14.2

What aspects would you want to evaluate for the following systems:

1. A personal music service?
2. A website for selling clothes?

Comment

1. You would need to discover how well users can select tracks from potentially thousands of tunes and whether they can easily add and store new music.
2. Navigation would be a core concern for both examples. Users of a personal music service will want to find tracks to select quickly. Users wanting to buy clothes will want to move quickly among pages displaying clothes, comparing them, and purchasing them. In addition, do the clothes look attractive enough to buy? Other core aspects include how trustworthy and how secure the procedure is for taking customer credit card details. ■

14.2.3 Where to Evaluate

Where evaluation takes place depends on what is being evaluated. Some characteristics, such as web accessibility, are generally evaluated in a lab because it provides the control necessary to investigate systematically whether all of the requirements are met. This is also true for design choices, such as choosing the size and layout of keys for a small handheld device for playing games. User experience aspects, such as whether children enjoy playing with a new toy and for how long before they get bored, can be evaluated more effectively in natural settings, which are often referred to as *in-the-wild studies*. Unlike a lab study, seeing children play in a natural setting will reveal when the children get bored and stop playing with the toy. In a lab study, the children are told what to do, so the UX researchers cannot easily see how the children naturally engage with the toy and when they get

bored. Of course, the UX researchers can ask the children whether they like it or not, but sometimes children will not say what they really think because they are afraid of causing offense.

Remote studies of online behavior, such as social networking, can be conducted to evaluate natural interactions of participants in the context of their interaction, for example, in their own homes or place of work. Living labs (see Box 14.1) have also been built that are a compromise between the artificial, controlled context of a lab and the natural, uncontrolled nature of in-the-wild studies. They provide the setting of a particular type of environment, such as the home, a workplace, or a gym, while also giving the ability to control, measure, and record activities through embedding technology in them.

ACTIVITY 14.3

A company is developing a new car seat to monitor whether a person is starting to fall asleep while driving and to provide a wake-up call using olfactory and haptic feedback. Where would you evaluate it?

Comment

It would be initially important to conduct lab-based experiments using a car simulator to see the effectiveness of the new type of feedback—in a safe setting, of course! You would also need to find a way to try to get the participants to fall asleep at the wheel. Once established as an effective mechanism, you would then need to test it in a more natural setting, such as a race track, airfield, or safe training circuit for new drivers, which can be controlled by the experimenter using a dual-control car. ■

14.2.4 When to Evaluate

The stage in the product lifecycle when evaluation takes place depends on the type of product and the development process being followed. For example, the product being developed could be a new concept, or it could be an upgrade to an existing product. It could also be a product in a rapidly changing market that needs to be evaluated to see how well the design meets current and predicted market needs. If the product is new, then considerable time is usually invested in market research and discovering user requirements. Once these requirements have been established, they are used to create initial sketches, a storyboard, a series of screens, or a prototype of the design ideas. These are then evaluated to see whether the designers have interpreted the users' requirements correctly and embodied them in their designs appropriately. The designs will be modified according to the evaluation feedback and new prototypes developed and subsequently evaluated.

When evaluations are conducted during design to check that a product continues to meet users' needs, they are known as *formative evaluations*. Formative evaluations cover a broad range of design processes, from the development of early sketches and prototypes through to tweaking and perfecting a nearly finished design.

Evaluations that are carried out to assess the success of a finished product are known as *summative evaluations*. If the product is being upgraded, then the evaluation may not focus on discovering new requirements but may instead evaluate the existing product to ascertain what needs improving. Features are then often added, which can result in new usability problems. At other times, attention is focused on improving specific aspects, such as enhanced navigation.

As discussed in earlier chapters, rapid iterations of product development that embed evaluations into short cycles of design, build, and test (evaluate) are common. In these cases, the evaluation effort may be almost continuous across the product's development and deployment lifetime. For example, this approach is sometimes adopted for government websites that provide information about Social Security, pensions, and citizens' voting rights.

Many agencies, such as the National Institute of Standards and Technology (NIST) in the United States, the International Standards Organization (ISO), and the British Standards Institute (BSI) set standards by which particular types of products, such as aircraft navigation systems and consumer products that have safety implications for users, have to be evaluated. The Web Content Accessibility Guidelines (WCAG) 2.1 describe how to design websites so that they are accessible. WCAG 2.1 is discussed in more detail in Box 16.2.

14.3 Types of Evaluation

We classify evaluations into three broad categories, depending on the setting, user involvement, and level of control. These are as follows:

- *Controlled settings directly involving users* (examples are usability labs and research labs): Users' activities are controlled to test hypotheses and measure or observe certain behaviors. The main methods are usability testing and experiments.
- *Natural settings involving users* (examples are online communities and products that are used in public places): There is little or no control of users' activities to determine how the product would be used in the real world. The main method used is field studies (for example in-the-wild studies).
- *Any settings not directly involving users*: Consultants and researchers critique, predict, and model aspects of the interface to identify the most obvious usability problems. The range of methods includes inspections, heuristics, walk-throughs, models, and analytics.

There are pros and cons of each evaluation category. For example, lab-based studies are good at revealing usability problems, but they are poor at capturing context of use; field studies are good at demonstrating how people use technologies in their intended setting, but they are often time-consuming and more difficult to conduct (Rogers et al., 2013); and modeling and predicting approaches are relatively quick to perform, but they can miss unpredictable usability problems and subtle aspects of the user experience. Similarly, analytics are good for tracking the use of a website but are not good for finding out how users feel about a new color scheme or why they behave as they do.

Deciding on which evaluation approach to use is determined by the goals of the project and on how much control is needed to find out whether an interface or device meets those goals and can be used effectively. For example, in the case of the music service mentioned earlier, this includes finding out how users use it, whether they like it, and what problems they experience with the functions. In turn, this requires determining how they carry out various tasks using

the interface operations. A degree of control is needed when designing the evaluation study to ensure participants try all of the tasks and operations for which the service is designed.

14.3.1 Controlled Settings Involving Users

Experiments and user tests are designed to control what users do, when they do it, and for how long. They are designed to reduce outside influences and distractions that might affect the results, such as people talking in the background. The approach has been extensively and successfully used to evaluate software applications running on laptops and other devices where participants can be seated in front of them to perform a set of tasks.

Usability Testing

This approach to evaluating user interfaces involves collecting data using a combination of methods in a controlled setting, for example, experiments that follow basic experimental design, observation, interviews, and questionnaires. Often, *usability testing* is conducted in labs, although increasingly interviews and other forms of data collection are being done remotely via phone and digital communication (for instance, through Skype or Zoom) or in natural settings. The primary goal is to determine whether an interface is usable by the intended user population to carry out the tasks for which it was designed. This involves investigating how typical users perform on typical tasks. By typical, we mean the users for whom the system is designed (for example, teenagers, adults, and so on) and the activities that it is designed for them to be able to do (such as, purchasing the latest fashions). It often involves comparing the number and kinds of errors that users make between versions and recording the time that it takes them to complete the task. As users perform the tasks, they may be recorded on video. Their interactions with the software may also be recorded, usually by logging software. User satisfaction questionnaires and interviews can also be used to elicit users' opinions about how they liked the experience of using the system. This data can be supplemented by observation at product sites to collect evidence about how the product is being used in the workplace or in other environments. Observing users' reactions to an interactive product has helped developers reach an understanding of usability issues, which would be difficult for them to glean simply by reading reports or listening to presentations. The qualitative and quantitative data that is collected using these different techniques are used in conjunction with each other to form conclusions about how well a product meets the needs of its users.

Usability testing is a fundamental, essential HCI process. For many years, usability testing has been a staple of companies, which is used in the development of standard products that go through many generations, such as word processing systems, databases, and spreadsheets (Johnson, 2014; Krug, 2014; Redish, 2012). The findings from usability testing are often summarized in a *usability specification* that enables developers to test future prototypes or versions of the product against it. Optimal performance levels and minimal levels of acceptance are generally specified, and current levels are noted. Changes in the design can then be implemented, such as navigation structure, use of terms, and how the system responds to users. These changes can then be tracked.

While usability testing is well established in UX design, it has also started to gain more prominence in other fields such as healthcare, particularly as mobile devices take an increasingly central role (Schnall et al., 2018) in hospitals and for monitoring one's own health (for instance, Fitbit and the Polar series, Apple Watch, and so forth). Another trend reported by Kathryn Whinton and Sarah Gibbons (2018) from the Nielsen/Norman (NN/g) Usability

Consulting Group is that while usability guidelines have tended to be stable over time, audience expectations change. For example, Whitenon and Gibbons report that since the last major redesign of the NN/g home page a few years ago, both the content and the audience's expectations about the attractiveness of the visual design have evolved. However, they stress that even though visual design has assumed a bigger and more important role in design, it should never replace or compromise basic usability. Users still need to be able to carry out their tasks effectively and efficiently.

ACTIVITY 14.4

Look at Figure 14.1, which shows two similarly priced devices for recording activity and measuring heart rate: (a) Fitbit Charge and (b) Polar A370. Assume that you are considering buying one of these devices. What usability issues would you want to know about, and what aesthetic design issues would be important to you when deciding which one to purchase?



Figure 14.1 Devices for monitoring activity and heart rate (a) Fitbit Charge and (b) Polar A370

Source: Figure 14.1 (a) Fitbit (b) Polar

Comment

There are several usability issues to consider. Some that you might be particularly interested in finding out about include how comfortable the device is to wear, how clearly the information is presented, what other information is presented (for example, time), how long the battery lasts before it needs to be recharged, and so forth. Most important of all might be how accurate the device is, particularly for recording heart rate if that is a concern for you.

Since these devices are worn on your wrist, they can be considered to be fashion items. Therefore, you might want to know whether they are available in different colors, whether they are bulky and likely to rub on clothes and cause damage, and whether they are discrete or clearly noticeable. ■

Experiments are typically conducted in research labs at universities or commercial labs to test such hypotheses. These are the most controlled evaluation settings, where researchers try to remove any extraneous variables that may interfere with the participant's performance. The reason for this is so that they can reliably say that the findings arising from the experiment are due to the particular interface feature being measured. For example, in an experiment comparing which is the best way for users to enter text when using an iPad or other type of tablet interface, researchers would control all other aspects of the setting to ensure that they do not affect the user's performance. These include providing the same instructions to all of the participants, using the same tablet interface, and asking the participants to do the same tasks. Depending on the available functionality, conditions that might be compared could be typing using a virtual keyboard, typing using a physical keyboard, and swiping using a virtual keyboard. The goal of the experiment would be to test whether one type of text input is better than the others in terms of speed of typing and number of errors. A number of participants would be brought into the lab separately to carry out the predefined set of text entry tasks, and their performance would be measured in terms of time consumed and how many errors are made, for example, selecting the wrong letter. The data collected would then be analyzed to determine whether the scores for each condition were significantly different. If the performance measures obtained for the virtual keyboard were significantly faster than those for the other two and had the least number of errors, one could say that this method of text entry is the best. Testing in a laboratory may also be done when it is too disruptive to evaluate a design in a natural setting, such as in a military conflict.

BOX 14.1

Living Labs

Living labs have been developed to evaluate people's everyday lives, which would be simply too difficult to assess in usability labs, for example, to investigate people's habits and routines over a period of several months. An early example of a living lab was the Aware Home (Abowd et al., 2000) in which the house was embedded with a complex network of sensors and audio/video recording devices that recorded the occupants' movements throughout the house and their use of technology. This enabled their behavior to be monitored and analyzed, for example, their routines and deviations. A primary motivation was to evaluate how real families would respond and adapt to such a setup over a period of several months. However, it proved difficult to get families to agree to leave their own homes and live in a living lab home for that long.

Ambient-assisted homes have also been developed where a network of sensors is embedded throughout someone's home rather than in a special, customized building. One rationale is to enable disabled and elderly people to lead safe and independent lives by providing a non-intrusive system that can remotely monitor and provide alerts to caregivers in the event of an accident, illness, or unusual activities (see Fernández-Luque et al., 2009). The term *living lab* is also used to describe innovation networks in which people gather in person and virtually to explore and form commercial research and development collaborations (Ley et al., 2017).

Nowadays, many living labs have become more like commercial enterprises, which offer facilities, infrastructure, and access to participating communities, bringing together users, developers, researchers, and other stakeholders. Living labs are being developed that form an

integral part of a smart building that can be adapted for different conditions in order to investigate the effects of different configurations of lighting, heating, and other building features on the inhabitant's comfort, work productivity, stress levels, and well-being. The Smart Living Lab in Switzerland, for example, is developing an urban block, including office buildings, apartments, and a school, to provide an infrastructure that offers opportunities for researchers to investigate different kinds of human experiences within built environments (Verma et al., 2017). Some of these spaces are large and may house hundreds and even thousands of people. People themselves are also being fitted with mobile and wearable devices that measure heart rate, activity levels, and so on, which can then be aggregated to assess the health of a population (for example, students at a university) over long periods of time. Hofte et al. (2009) call this a *mobile living lab* approach, noting how it enables more people to be studied for longer periods and at the times and locations where observation by researchers is difficult.

Citizen science, in which volunteers work with scientists to collect data on a scientific research issue, such as biodiversity (for instance, iNaturalist.org), monitoring the flowering times of plants over tens or hundreds of years (Primak, 2014), or identifying galaxies online (for example, <https://www.zooniverse.org/projects/zookeeper/galaxy-zoo/>), can also be thought of as a type of living lab, especially when the behavior of the participants, their use of technology, and the design of that technology are also being studied. Lab in the Wild (<http://www.LabintheWild.org>), for example, is an online site that hosts volunteers who participate in a range of projects. Researchers analyzed more than 8,000 comments from volunteers involved in four experiments. They concluded that such online sites have potential as online research labs (Oliveira et al., 2017) that can be studied over time and hence form a type of living lab. ■

DILEMMA

Is a Living Lab Really a Lab?

The concept of a living lab differs from a traditional view of a lab insofar as it is trying to be both natural and experimental and the goal is to bring the lab into the home (or other natural setting) or online. The dilemma is how artificial to make the more natural setting; where does the balance lie in setting it up to enable the right level of control to conduct research and evaluation without losing the sense of it being natural? ■

14.3.2 Natural Settings Involving Users

The goal of field studies is to evaluate products with users in their natural settings. Field studies are used primarily to

- Help identify opportunities for new technology
- Establish the requirements for a new design
- Facilitate the introduction of technology or inform deployment of existing technology in new contexts

Methods that are typically used are observation, interviews, and interaction logging (see Chapters 8 and 9). The data takes the form of events and conversations that are recorded by the researchers as notes, or through audio or video recording, or by the participants as diaries and notes. The goal is to be unobtrusive and not to affect what people do during the evaluation. However, it is inevitable that some methods will influence how people behave. For example, diary studies require people to document their activities or feelings at certain times, and this can make them reflect on and possibly change their behavior.

During the last 15 years, there has been a trend toward conducting in-the-wild studies. These are essentially user studies that look at how new technologies or prototypes have been deployed and used by people in various settings, such as outdoors, in public places, and in homes. Sometimes, a prototype that is deployed is called a *disruptive technology*, where the aim is to determine how it displaces an existing technology or practice. In moving into the wild, researchers inevitably have to give up control of what is being evaluated in order to observe how people approach and use—or don't use—technologies in their everyday lives. For example, a researcher might be interested in observing how a new mobile navigation device will be used in urban environments. To conduct the study, they would need to recruit people who are willing to use the device for a few weeks or months in their natural surroundings. They might then tell the participants what they can do with the device. Other than that, it is up to the participants to decide how to use it and when, as they move among work or school, home, and other places.

The downside of handing over control is that it makes it difficult to anticipate what is going to happen and to be present when something interesting does happen. This is in contrast to usability testing where there is always an investigator or camera at hand to record events. Instead, the researcher has to rely on the participants recording and reflecting on how they use the product, by writing up their experiences in diaries, filling in online forms, and/or taking part in intermittent interviews.

Field studies can also be virtual, where observations take place in multiuser games such as *World of Warcraft*, online communities, chat rooms, and so on. A goal of this kind of field study is to examine the kinds of social processes that occur in them, such as collaboration, confrontation, and cooperation. The researcher typically becomes a participant and does not control the interactions (see Chapters 8 and 9). Virtual field studies have also become popular in the geological and biological sciences because they can supplement studies in the field. Increasingly, online is partnered with a real-world experience so that researchers and students get the best of both situations (Cliffe, 2017).

14.3.3 Any Settings Not Involving Users

Evaluations that take place without involving users are conducted in settings where the researcher has to imagine or model how an interface is likely to be used. Inspection methods are commonly employed to predict user behavior and to identify usability problems based on knowledge of usability, users' behavior, the contexts in which the system will be used, and the kinds of activities that users undertake. Examples include heuristic evaluation that applies knowledge of typical users guided by rules of thumb and walkthroughs that involve stepping through a scenario or answering a set of questions for a detailed prototype. Other techniques include analytics and models.

The original heuristics used in heuristic evaluation were for screen-based applications (Nielsen and Mack, 1994; Nielsen and Tahir, 2002). These have been adapted to develop new sets of heuristics for evaluating web-based products, mobile systems, collaborative

technologies, computerized toys, information visualizations (Forsell and Johansson, 2010), and other new types of systems. One of the problems with using heuristics is that designers can sometimes be led astray by findings that are not as accurate as they appeared to be at first (Tomlin, 2010). This problem can arise from different sources, such as a lack of experience and the biases of UX researchers who conduct the heuristic evaluations.

Cognitive walk-throughs involve simulating a user's problem-solving process at each step in the human-computer dialogue and checking to see how users progress from step to step in these interactions (see Wharton et al., 1994 in Nielsen and Mack, 1994). During the last 15 years, cognitive walk-throughs have been used to evaluate smartphones (Jadhav et al., 2013), large displays, and other applications, such as public displays (Parker et al., 2017). A key feature of cognitive walk-throughs is that they focus on evaluating designs for ease of learning.

Analytics is a technique for logging and analyzing data either at a customer's site or remotely. *Web analytics* is the measurement, collection, analysis, and reporting of Internet data to understand and optimize web usage. Examples of web analytics include the number of visitors to a website home page over a particular time period, the average time users spend on the home page, which other pages they visit, or whether they leave after visiting the homepage. For example, Google provides a commonly used approach for collecting analytics data that is a particularly useful method for evaluating design features of a website (<https://marketingplatform.google.com/about/analytics/>). As part of the massive open online courses (MOOCs) and open educational resources (OERs) movement, *learning analytics* has evolved and gained prominence for assessing the learning that takes place in these environments. The Open University in the United Kingdom, along with others, has published widely on this topic, describing how learning analytics are useful for guiding course and program design and for evaluating the impact of pedagogical decision-making (Toetenel and Bart, 2016).

This web page provides information about learning analytics and learning design:
<https://iet.open.ac.uk/themes/learning-analytics-and-learning-design>

Models have been used primarily for comparing the efficacy of different interfaces for the same application, for example, the optimal arrangement and location of features. A well-known approach uses *Fitts' law* to predict the time it takes to reach a target using a pointing device (MacKenzie, 1995) or using the keys on a mobile device or game controller (Ramcharitar and Teather, 2017).

14.3.4 Selecting and Combining Methods

The three broad categories identified previously provide a general framework to guide the selection of evaluation methods. Often, combinations of methods are used across the categories to obtain a richer understanding. For example, sometimes usability testing conducted in labs is combined with observations in natural settings to identify the range of usability problems and find out how users typically use a product.

There are both pros and cons for controlled and uncontrolled settings. The benefits of controlled settings include being able to test hypotheses about specific features of the interface where the results can be generalized to the wider population. A benefit of uncontrolled settings is that unexpected data can be obtained that provides quite different insights into people's perceptions and their experiences of using, interacting, or communicating through the new technologies in the context of their everyday and working lives.

14.3.5 Opportunistic Evaluations

Evaluations may be detailed, planned studies, or *opportunistic explorations*. The latter explorations are generally done early in the design process to provide designers with feedback quickly about a design idea. Getting this kind of feedback early in the design process is important because it confirms whether it is worth proceeding to develop an idea into a prototype. Typically, these early evaluations are informal and do not require many resources. For example, the designers may recruit a few local users and ask their opinions. Getting feedback this early in design provides feedback early on when it is easier to make changes to an evolving design. Opportunistic evaluations with users can also be conducted to hone the target audience so that subsequent evaluation studies can be more focused. Opportunistic evaluations can also be conducted in addition to more formal evaluations.

14.4 Evaluation Case Studies

Two contrasting case studies are described in this section to illustrate how evaluations can take place in different settings with different amounts of control over users' activities. The first case study (section 14.4.1) describes a classic experiment that tested whether it was more exciting playing against a computer versus playing against a friend in a collaborative computer game (Mandryk and Inkpen, 2004). Though published more than 15 years ago, we are keeping this case study in this edition of the book because it provides a concise and clear description about a variety of measures that were used in the experiment. The second case study (section 14.4.2) describes an ethnographic field study in which a bot, known as Ethnobot, was developed to prompt participants to answer questions about their experiences while walking around a large outdoor show (Tallyn et al., 2018).

14.4.1 Case Study 1: An Experiment Investigating a Computer Game

For games to be successful, they must engage and challenge users. Criteria for evaluating these aspects of the user experience are therefore needed. In this case study, physiological responses were used to evaluate users' experiences when playing against a friend and when playing alone against the computer (Mandryk and Inkpen, 2004). Regan Mandryk and Kori Inkpen conjectured that physiological indicators could be an effective way of measuring a player's experience. Specifically, they designed an experiment to evaluate the participants' engagement while playing an online ice-hockey game.

Ten participants, who were experienced game players, took part in the experiment. During the experiment, sensors were placed on the participants to collect physiological data. The data collected included measurements of the moisture produced by sweat glands of their hands and feet and changes in heart and breathing rates. In addition, they videoed the

participants and asked them to complete user satisfaction questionnaires at the end of the experiment. To reduce the effects of learning, half of the participants played first against a friend and then against the computer, and the other half played against the computer first. Figure 14.2 shows the setup for recording data while the participants were playing the game.

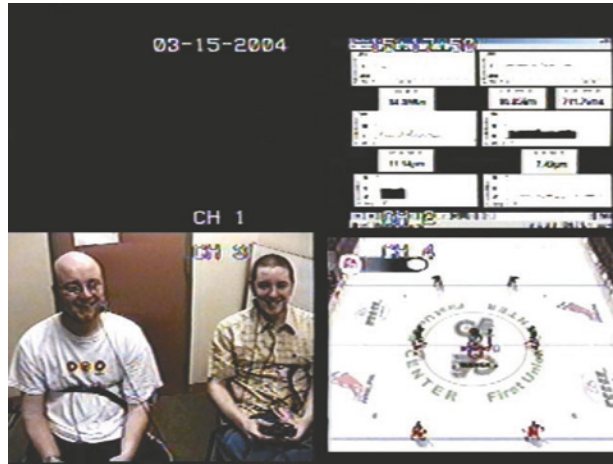


Figure 14.2 The display shows the physiological data (top right), two participants, and a screen of the game they played.

Source: Mandryk and Inkpen (2004). *Physiological Indicators for the Evaluation of Co-located Collaborative Play*, CSCW'2004, pp. 102–111. Reproduced with permission of ACM Publications

Results from the user satisfaction questionnaire revealed that the mean ratings on a 1–5 scale for each item indicated that playing against a friend was the favored experience (Table 14.1). Data recorded from the physiological responses was compared for the two conditions and in general revealed higher levels of excitement when participants played against a friend than when they played against the computer. The physiological recordings were also compared across participants and, in general, indicated the same trend. Figure 14.3 shows a comparison for two participants.

| | Playing Against Computer | | Playing Against Friend | |
|-------------|--------------------------|----------|------------------------|----------|
| | Mean | St. Dev. | Mean | St. Dev. |
| Boring | 2.3 | 0.949 | 1.7 | 0.949 |
| Challenging | 3.6 | 1.08 | 3.9 | 0.994 |
| Easy | 2.7 | 0.823 | 2.5 | 0.850 |
| Engaging | 3.8 | 0.422 | 4.3 | 0.675 |
| Exciting | 3.5 | 0.527 | 4.1 | 0.568 |
| Frustrating | 2.8 | 1.14 | 2.5 | 0.850 |
| Fun | 3.9 | 0.738 | 4.6 | 0.699 |

Table 14.1 Mean subjective ratings given on a user satisfaction questionnaire using a five-point scale, in which 1 is lowest and 5 is highest for the 10 players

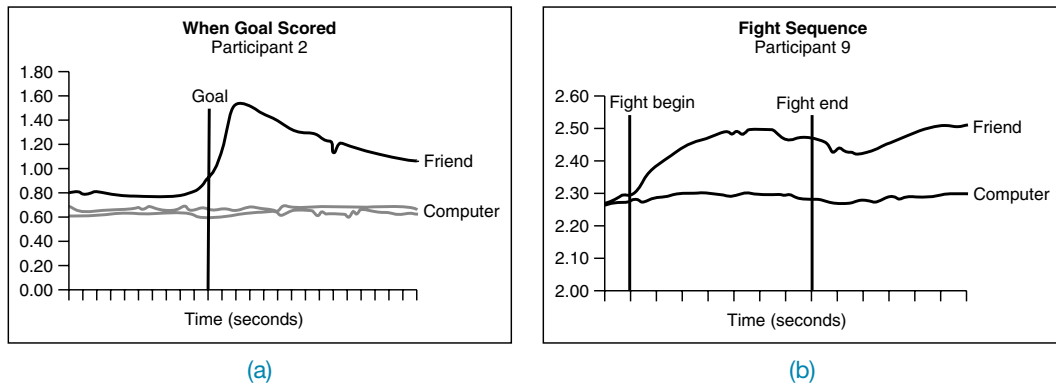


Figure 14.3 (a) A participant's skin response when scoring a goal against a friend versus against the computer, and (b) another participant's response when engaging in a hockey fight against a friend versus against the computer

Source: Mandryk and Inkpen (2004). *Physiological Indicators for the Evaluation of Co-located Collaborative Play, CSCW'2004*, pp. 102–111. Reproduced with permission of ACM Publications

Identifying strongly with an experience state is indicated by a higher mean. The standard deviation indicates the spread of the results around the mean. Low values indicate little variation in participants' responses; high values indicate more variation.

Because of individual differences in physiological data, it was not possible to compare directly the means of the two sets of data collected: subjective questionnaires and physiological measures. However, by normalizing the results, it was possible to correlate the results across individuals. This indicated that the physiological data gathering and analysis methods were effective for evaluating levels of challenge and engagement. Although not perfect, these two kinds of measures offer a way of going beyond traditional usability testing in an experimental setting to get a deeper understanding of user experience goals.

ACTIVITY 14.5

1. What kind of setting was used in this experiment?
2. How much control did the researchers exert?
3. Which types of data were collected?

Comment

1. The experiment took place in a research lab, which is a controlled setting.
2. The evaluation was strongly controlled by the evaluators. They specified which of the two gaming conditions was assigned to each participant. The participants also had sensors placed on them to collect physiological data as they played the game, for example to monitor changes in heart rate and breathing.
3. Physiological measures of the participants while playing the game were collected together with data collected afterward using a user satisfaction questionnaire that asked questions about how satisfied they were with the game and how much they enjoyed it. ■

14.4.2 Case Study 2: Gathering Ethnographic Data at the Royal Highland Show

Field observations, including in-the-wild and ethnographic studies, provide data about how users interact with technology in their natural environments. Such studies often provide insights not available in lab settings. However, it can be difficult to collect participants' thoughts, feelings, and opinions as they move about in their everyday lives. Usually, it involves observations and asking them to reflect after an event, for example through interviews and diaries. In this case study, a novel evaluation approach—a live chatbot—was used to address this gap by collecting data about people's experiences, impressions, and feelings as they visited and moved around the Royal Highland Show (RHS) (Tallyn et al., 2018). The RHS is a large agricultural show that runs every June in Scotland. The chatbot, known as Ethnobot, was designed as an app that runs on a smartphone. In particular, Ethnobot was programmed to ask participants pre-established questions as they wandered around the show and to prompt them to expand on their answers and take photos. It also directed them to particular parts of the show that the researchers thought would interest the participants. This strategy also allowed the researchers to collect data from all of the participants in the same place. Interviews were also conducted by human researchers to supplement the data collected online by the Ethnobot.

The overall purpose of the study was to find out about participants' experiences of, and feelings about, the show and of using Ethnobot. The researchers also wanted to compare the data collected by the Ethnobot with the interview data collected by the human researchers.

The study consisted of four data collection sessions using the Ethnobot over two days and involved 13 participants, who ranged in age and came from diverse backgrounds. One session occurred in the early afternoon and the other in the late afternoon on each day of the study. Each session lasted several hours. To participate in the study, each participant was given a smartphone and shown how to use the Ethnobot app (Figure 14.4), which they could experience on their own or in groups as they wished.

Two main types of data were collected.

- The participants' online responses to a short list of pre-established questions that they answered by selecting from a list of prewritten comments (for example, "I enjoyed something" or "I learned something") presented by the Ethnobot in the form of buttons called *experience buttons*, and the participants' additional open-ended, online comments and photos that they offered in response to prompts for more information from Ethnobot. The participants could contribute this data at any time during the session.
- The participants' responses to researchers' in-person interview questions. These questions focused on the participants' experiences that were not recorded by the Ethnobot, and their reactions to using the Ethnobot.

A lot of data was collected that had to be analyzed. The pre-established comments collected in the Ethnobot chatlogs were analyzed quantitatively by counting the responses. The in-person interviews were audio-recorded and transcribed for analysis, and that involved coding them, which was done by two researchers who cross-checked each other's analysis for consistency. The open-ended online comments were analyzed in a similar way to the in-person interview data.

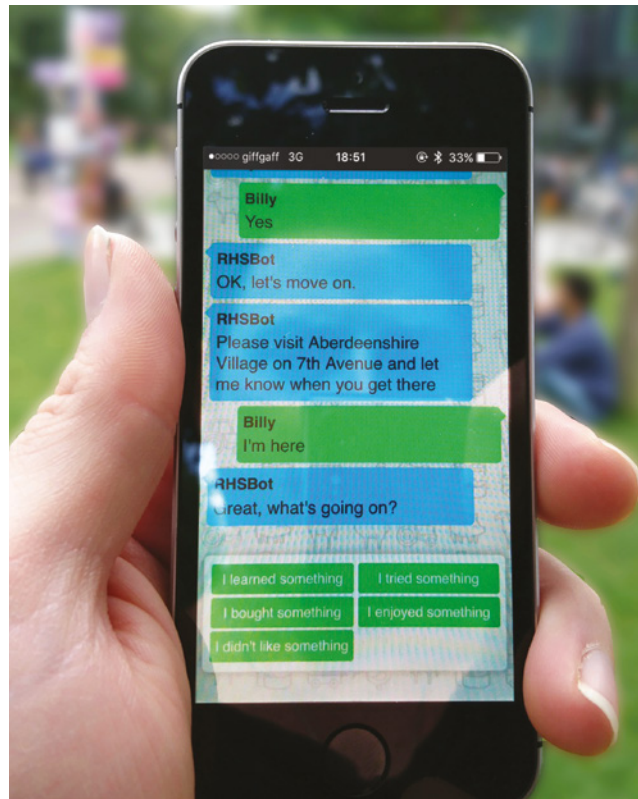


Figure 14.4 The Ethnobot used at the Royal Highland Show in Scotland. Notice that the Ethnobot directed participant Billy to a particular place (that is, Aberdeenshire Village). Next, Ethnobot asks “. . . What’s going on?” and the screen shows five of the experience buttons from which Billy needs to select a response

Source: Tallyn et al. (2018). Reproduced with permission of ACM Publications

Overall, the analyses revealed that participants spent an average of 120 minutes with the Ethnobot on each session and recorded an average of 71 responses, while submitting an average of 12 photos. In general, participants responded well to prompting by the Ethnobot and were eager to add more information. For example, P9 said, “I really enjoyed going around and taking pictures and [to the question] ‘have you got something to add’ [said] yeah! I have, I always say ‘yes’. . . .” A total of 435 pre-established responses were collected, including 70 that were about what the participants did or experienced (see Figure 14.5). The most frequent response was “I learned something” followed by “I tried something” and “I enjoyed something.” Some participants also supplied photos to illustrate their experiences.

When the researchers asked the participants about their reactions to selecting prewritten comments, eight participants remarked that they were rather restrictive and that they would like more flexibility to answer the questions. For example, P12 said, “maybe there should have been more options, in terms of your reactions to the different parts of the show.” However, in general participants enjoyed their experience of the RHS and of using Ethnobot.

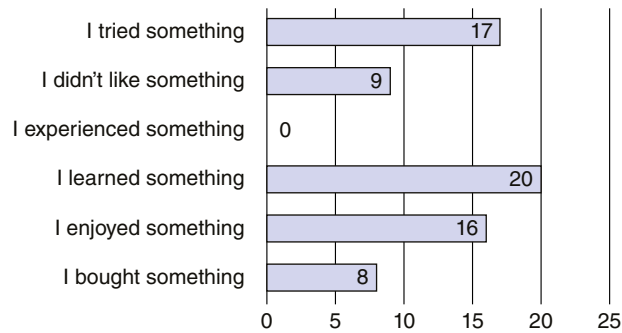


Figure 14.5 The number of prewritten experience responses submitted by participants to the pre-established questions that Ethnobot asked them about their experiences

Source: Tallyn et al. (2018). Reproduced with permission of ACM Publications

When the researchers compared the data collected by Ethnobot with that from the interviews collected by the human researchers, they found that the participants provided more detail about their experiences and feelings in response to the in-person interview questions than to those presented by Ethnobot. Based on the findings of this study, the researchers concluded that while there are some challenges to using a bot to collect in-the-wild evaluation data, there are also advantages, particularly when researchers cannot be present or when the study involves collecting data from participants on the move or in places that are hard for researchers to access. Collecting data with a bot and supplementing it with data collected by human researchers appears to offer a good solution in circumstances such as these.

ACTIVITY 14.6

1. What kind of setting was used in this evaluation?
2. How much control did the researchers exert?
3. Which types of data were collected?

Comment

1. The evaluation took place in a natural outdoor setting at the RHS.
2. The researchers imposed less control on the participants than in the previous case study, but the Ethnobot was programmed to ask specific questions, and a range of responses was provided from which participants selected. The Ethnobot was also programmed to request additional information and photos. In addition, the Ethnobot was programmed to guide the participants to particular areas of the show, although some participants ignored this guidance and went where they pleased.
3. The Ethnobot collected answers to a specific set of predetermined questions (closed questions) and prompted participants for additional information and photographs. In addition,

participants were interviewed by the researchers using semi-structured, open-ended interviews. The data collected was qualitative, but counts of the response categories produced quantitative data (see Figure 14.5). Some demographic data was also quantitative (for instance, participants' ages, gender, and so forth), which is provided in the full paper (Talyn et al., 2018). ■

BOX 14.2

Crowdsourcing

The Internet makes it possible to gain access to hundreds of thousands of participants who will perform tasks or provide feedback on a design or experimental task quickly and almost immediately. Mechanical Turk is a service hosted by Amazon that has thousands of people registered (known as Turkers), who have volunteered to take part by performing various activities online, known as *human intelligence tasks* (HITs), for a very small reward. HITs are submitted by researchers or companies that pay a few cents for simple tasks (such as tagging pictures) to a few dollars (for taking part in an experiment). Advantages of using *crowdsourcing* in HCI is that it is more flexible, relatively inexpensive, and often much quicker to enroll participants than with traditional lab studies. Another benefit is that many more participants can be recruited.

Early in the history of online crowdsourcing, Jeff Heer and Michael Bostock (2010) used it to determine how reliable it was to ask random people over the Internet to take part in an experiment. Using Mechanical Turk, they asked the Turkers to perform a series of perception tasks using different visual display techniques. A large number agreed, enabling them to analyze their results statistically and generalize from their findings. They also developed a short set of test questions that generated 2,880 responses. They then compared their findings from using crowdsourcing with those reported in published lab-based experiments. They found that while the results from their study using Turkers showed wider variance than in the lab study, the overall results across the studies were the same. They also found that the total cost of their experiment with Turkers was one-sixth the cost of a typical lab study involving the same number of people. While these results are important, online crowdsourcing studies have also raised ethical questions about whether Turkers are being appropriately rewarded and acknowledged—an important question that continues to be discussed (see, for example, the Brookings Institute article by Vanessa Williamson, 2016).

Since Jeff Heer's and Michael Bostock's 2010 study, crowdsourcing has become increasingly popular and has been used in a wide range of applications including collecting design ideas, such as ideation, for developing a citizen science app (Maher et al., 2014); managing volunteers for disaster relief (Ludwig et al., 2016); and delivering packages (Kim, 2015). Both the number and diversity of useful contributions and ideas generated make crowdsourcing particularly attractive for getting timely feedback from the public. For example, in a study to collect and improve the design of a street intersection, a system called *CommunityCrit* was used to collect opinions from members of the community and to draw on their skills and availability (Mahyar et al., 2018). Those who contributed were empowered by getting to see the planning process.

Citizen science is also a form of crowdsourcing. Originally dating back to the days of Aristotle and Darwin, the data was collected by humans, who were sometimes referred to as *sensors*. During the last 10 years, the volume of data collected has increased substantially, leveraged by technology, particularly smartphones, and a range of other digital devices (Preece, 2017). For example, iSpotNature.org, iNaturalist.com, and eBird.com are apps that are used across the world for collecting biodiversity data and data about bird behavior.

These examples illustrate how crowdsourcing can be a powerful tool for improving, enhancing, and scaling up a wide range of tasks. Crowdsourcing makes it possible to recruit participants to generate a large pool of potential ideas, collect data, and make other useful inputs that would be difficult to achieve in other ways. Several companies, including Google, Facebook, and IDEO, use crowdsourcing to try ideas and to gather evaluation feedback about designs. ■

14.5 What Did We Learn from the Case Studies?

The case studies along with Box 14.1 and Box 14.2 provide examples of how different evaluation methods are used in different physical settings that involve users in different ways to answer various kinds of questions. They demonstrate how researchers exercise different levels of control in different settings. The case studies also show how it is necessary to be creative when working with innovative systems and when dealing with constraints created by the evaluation setting (for example, online, distributed, or outdoors where people are on the move) and the technology being evaluated. In addition, the case studies and boxes discussed illustrate how to do the following:

- Observe users in the lab and in natural settings
- Develop different data collection and analysis techniques to evaluate user experience goals, such as challenge and engagement and people on the move
- Run experiments on the Internet using crowdsourcing, thereby reaching many more participants while being straightforward to run
- Recruit a large number of participants who contribute to a wide range of projects with different goals using crowdsourcing

BOX 14.3

The Language of Evaluation

Sometimes terms describing evaluation are used interchangeably and have different meanings. To avoid this confusion, we define some of these terms here in alphabetical order. (You may find that other books use different terms.)

Analytics Data analytics refers to examining large volumes of raw data with the purpose of drawing inferences about a situation or a design. Web analytics is commonly used to measure website traffic through analyzing users' click data.

Analytical evaluation This type of evaluation models and predicts user behavior. This term has been used to refer to heuristic evaluation, walk-throughs, modeling, and analytics.

Bias The results of an evaluation are distorted. This can happen for several reasons. For example, selecting a population of users who have already had experience with the new system and describing their performance as if they were new users.

Controlled experiment This is a study that is conducted to test hypotheses about some aspect of an interface or other dimension. Aspects that are controlled typically include the task that participants are asked to perform, the amount of time available to complete the tasks, and the environment in which the evaluation study occurs.

Crowdsourcing This can be done in person (as was typical in citizen science for decades) or online via the web and mobile apps. Crowdsourcing provides the opportunity for hundreds, thousands, or even millions of people to evaluate a product or take part in an experiment. The crowd may be asked to perform a particular evaluation task using a new product or to rate or comment on the product.

Ecological validity This is a particular kind of validity that concerns how the environment in which an evaluation is conducted influences or even distorts the results.

Expert review or crit This is an evaluation method in which someone (or several people) with usability expertise and knowledge of the user population reviews a product looking for potential problems.

Field study This type of evaluation study is done in a natural environment such as in a person's home or in a work or leisure place.

Formative evaluation This type of evaluation is done during design to check that the product fulfills requirements and continues to meet users' needs.

Heuristic evaluation This is an evaluation method in which knowledge of typical users is applied, often guided by heuristics, to identify usability problems.

Informed consent form This form describes what a participant in an evaluation study will be asked to do, what will happen to the data collected about them, and their rights while involved in the study.

In-the-wild study This is a type of field study in which users are observed using products or prototypes within their everyday context.

Living lab This place is configured to measure and record people's everyday activities in a natural setting, such as in the home.

Predictive evaluation This type of evaluation is where theoretically based models are used to predict user performance.

Reliability The reliability or consistency of a method is how well it produces the same results on separate occasions under the same circumstances.

Scope This refers to how much the findings from an evaluation can be generalized.

Summative evaluation This evaluation is done when the design is complete.

Usability lab This lab is specially designed for usability testing.

Usability testing This involves measuring how well a design supports users' performance on various tasks.

User studies This generic term covers a range of evaluations involving users, including field studies and experiments.

Users or participants In this context, these terms are used interchangeably to refer to the people who take part in evaluation studies.

Validity Validity is concerned with whether the evaluation method measures what it is intended to measure. ■

14.6 Other Issues to Consider When Doing Evaluation

Reading the case studies may have raised other issues, such as the importance of asking good question to focus the evaluation. A good question is important because it helps to focus the evaluation and decide on the best approach and methods to use. Another issue is how to find suitable participants and, having found them, how to approach them. Can you just ask children in a café to participate, or do you need permission from their parents? What do you have to tell participants, and what if they decide part way through the study that they don't want to continue to the end? Can they stop, or do they have to continue? Two central issues are:

- Informing participants about their rights
- Making sure you take into account biases and other influences that impact how you describe your evaluation findings

14.6.1 Informing Participants About Their Rights and Getting Their Consent

Most professional societies, universities, government, and other research offices require researchers and those performing evaluation studies to provide information about activities in which human participants will be involved. They do this to protect participants by ensuring that they are not endangered physically or emotionally and that their right to privacy is protected, particularly the details about how participants' data is collected and will be treated. Drawing up such an agreement is mandatory in many universities and major organizations. Indeed, special review boards generally prescribe the format required, and many provide a detailed form that must be completed. Once the details are accepted, the review board checks periodically to oversee compliance. In American universities, these are known as *institutional review boards* (IRBs).

Institutions in other countries use different names, forms, and processes to protect users, and some countries have different laws that govern areas such as users' privacy, mentioned in Chapter 8, "Data Gathering." For example, the *General Data Protection Regulation* (GDPR) was introduced in 2018 to strengthen data protection and privacy for all individuals living within the European Union. Such laws influence not just the countries directly involved but also people in other countries who collaborate with EU countries on research projects or commercial software development.

Over the years, IRB forms have become increasingly detailed, particularly now that much research involves the Internet and people's interaction via social media and other communications technologies. IRB reviews are especially stringent when a research or evaluation study involves people who could be considered vulnerable (such as children, older adults, and people with disabilities).

Several lawsuits at prominent universities have heightened attention to IRB and similar compliance laws and standards to the extent that it sometimes takes several months and multiple amendments to get IRB acceptance. Not only are IRB reviewers interested in the more obvious issues of how participants will be treated and what they will be asked to do; they also want to know how the data will be analyzed and stored. For example, data about participants must be coded and stored to prevent linking participants' names with that data.

Participants must be told what they will be asked to do, the conditions under which data will be collected, and what will happen to their data when they finish the task. Participants must also be told their rights, for instance, that they may withdraw from the study at any time if they want. This information is usually presented to participants on a form, often referred to as a *consent form*, that each participant reads and signs before the study starts. When new laws come into existence, such as the EU's GDPR mentioned earlier, it is particularly important to be aware of how such laws will be enacted and their potential impact on research and evaluation studies.

Some companies have “boilerplate” templates that UX researchers and designers can use that describe how participants will be treated and how the data collected will be used so that new documents do not have to be created for each evaluation study. Many companies also ask the evaluation participants to sign a nondisclosure agreement, which requires that they do not talk about the product and their experience of evaluating it with anyone after completing the evaluation. Companies require this because they do not want their competitors and the public to know about the product before it is launched or modified.

DILEMMA

When is a person considered vulnerable, and how might this affect them?

Who is vulnerable? The answer is all of us at various times and stages in our lives. At any particular time, however, some people are more vulnerable than others (for example, children and people with emotional and certain physical disabilities). Furthermore, definitions of people who are vulnerable varies from country to country, state to state, and policy to policy, so the following two scenarios are broad categories to get you thinking about this important issue. At what age can children read and sign their own consent forms? Is it when they are considered to be old enough to understand what they are being asked to do? This could be 12 years of age, or at other times and places 16 or even 18 or 21. It also depends on the kind of study. In some parts of the world, a 17-year-old can get married but may need their parents to sign a form saying that they can take part in an evaluation study to rate the realism of a social robot's expressions. What is the balance here between seeking reasonable consent and respecting individuals' rights to privacy for themselves and their families? ■

14.6.2 Issues That Influence the Choice of Method and How the Data Is Interpreted

Decisions have to be made about what data is needed to answer the study questions, how the data will be analyzed, and how the findings will be presented (see Chapters 8 and 9). To a great extent, the method used determines the type of data collected, but there are still some choices. For example, should the data be treated statistically? Ideally, this question is

addressed before the data is collected, but if unexpected data arises, for instance from an in-the-wild study, this question may need to be considered afterward. For example, in-the-wild studies sometimes generate demographic data or counts (categorical data) that can be analyzed and presented using descriptive statistics (for example, the percentage of people in different age groups). Some general questions also need to be asked. Is the method reliable? Has the method produced the kind of data intended? Is the evaluation study ecologically valid, or is the fundamental nature of the process being changed by studying it? Are biases creeping in that will distort the results? Will the results be generalizable; that is, what is their scope?

Reliability

The *reliability* or consistency of a method is how well it produces the same results on separate occasions under the same circumstances. Another evaluator or researcher who follows the same procedure should get similar results. Different evaluation methods have different degrees of reliability. For example, a carefully controlled experiment will have high reliability, whereas observing users in their natural setting will be variable. An unstructured interview will have low reliability—it would be difficult if not impossible to repeat exactly the same discussion.

Validity

Validity is concerned with whether the evaluation method measures what it is intended to measure. This encompasses both the method itself and the way it is implemented. If, for example, the goal of an evaluation study is to find out how users use a new product in their homes, then it is not appropriate to plan a lab experiment. An ethnographic study in users' homes would be more appropriate. If the goal is to find average performance times for completing a task, then a method that only recorded the number of user errors would be invalid. These examples are deliberately extreme, but subtler mistakes can be made, and it's good to consider these questions for each study.

Ecological Validity

Ecological validity is a particular kind of validity that concerns how the environment in which an evaluation is conducted influences or even distorts the results. For example, lab experiments are controlled, so what the participants do and how they behave is quite different from what happens naturally in their workplace, at home, or in leisure environments. Lab experiments therefore have low ecological validity because the results are unlikely to represent what happens in the real world. In contrast, ethnographic studies do not impact the participants or the study location as much, so they have high ecological validity.

Ecological validity is also affected when participants are aware of being studied. This is sometimes called the *Hawthorne effect* after a series of experiments at the Western Electric Company's Hawthorne factory in the United States in the 1920s and 1930s. The studies investigated changes in length of working day, heating, lighting, and so on; however, eventually it was discovered that the workers were reacting positively to being given special treatment rather than just to the experimental conditions. Similar findings sometimes

occur in medical trials. Patients given the placebo dose (a false dose in which no drug is administered) show improvement that is due to receiving extra attention that makes them feel good.

Bias

Bias occurs when the results are distorted. For example, expert evaluators performing a heuristic evaluation may be more sensitive to certain kinds of design flaws than others, and this will be reflected in the results. When collecting observational data, researchers may consistently fail to notice certain types of behavior because they do not deem them important. Put another way, they may selectively gather data that they think is important. Interviewers may subconsciously influence responses from interviewees by their tone of voice, their facial expressions, or the way questions are phrased, so it is important to be sensitive to the possibility of biases.

Scope

The *scope* of an evaluation study refers to how much its findings can be generalized. For example, some modeling methods, like Fitts' Law (also discussed in Chapter 16, "Evaluation: Inspections, Analytics, and Models"), which is used to evaluate keypad design, have a narrow, precise scope. (The problems of overstating results are discussed in Chapter 9, "Data Analysis").

In-Depth Activity

In this activity, think about the case studies and reflect on the evaluation methods used.

1. For the two case studies discussed in this chapter, think about the role of evaluation in the design of the system and note the artifacts that were evaluated: *when* during the design were they evaluated, *which* methods were used, and *what* was learned from the evaluations? Note any issues of particular interest. You may find that constructing a table like the one shown here is a helpful approach.

| Name of the study or artifact evaluated | When during the design the evaluation occurred? | How the controlled the study and what role did users have? | Which methods were used? | What kind of data was collected, and how was it analyzed? | What was learned from the study? | Notable issues |
|---|---|--|--------------------------|---|----------------------------------|----------------|
| | | | | | | |

2. What were the main constraints that influenced the evaluations?
3. How did the use of different methods build on and complement each other to give a broader picture of the evaluations?
4. Which parts of the evaluations were directed at usability goals and which at user experience goals?

Summary

The goal of this chapter was to introduce the main approaches to evaluation and the methods typically used. These will be revisited in greater depth in the next two chapters. This chapter stressed how evaluation is done throughout design by collecting information about users' or potential users' experiences when interacting with a prototype, a computer system, a component of a computer system, or a design artifact (such as a screen sketch) to improve its design.

The pros and cons of running lab-based evaluations versus field studies were outlined in terms of participant reach, cost, effort, constraints, and the types of results that can be elicited. Choosing which approach to use will depend on the goals of the evaluation, the researcher's or evaluator's expectations, and the resources available to them.

Crowdsourcing was presented as a creative way of involving a wide range of people with different ideas and skills. Finally, we briefly mentioned the ethical issues relating to how evaluation participants are treated and their rights to privacy. We also raised questions about data interpretation including the need to be aware of biases, reliability, data and ecological validity, and the scope of the study.

Key Points

- Evaluation and design are closely integrated.
- Some of the same data gathering methods are used in evaluation as for discovering requirements and identifying users' needs, for instance, observation, interviews, and questionnaires.
- Evaluations can be done in controlled settings such as labs, less-controlled field settings, or where users are not present.
- Usability testing and experiments involve a high level of control over both what users do and what is tested, whereas field and in-the-wild evaluations typically impose little or no control on participants.
- Different methods are usually combined to provide different perspectives within a study.
- Participants need to be made aware of their rights. This is often done through informed consent forms.
- It is important not to over-generalize findings from an evaluation.

Further Reading

KRUGE, S. (2014) *Don't Make Me Think: A Common Sense Approach to Web Usability* (3rd ed.). New Riders. This book provides many useful practical examples of usability issues and how best to avoid them.

LAZAR, J., FENG, J. H. and HOCHHEISER, H. (2017) *Research Methods in Human-Computer Interaction* (2nd ed.). Cambridge, MA: Elsevier/Morgan Kaufmann Publishers. This book provides a useful overview of qualitative and quantitative methods. Chapter 15, "Working with Human Subjects," discusses ethical issues of working with human participants. PowerPoint slides are also available at:

<https://www.elsevier.com/books-and-journals/book-companion/9780128053904>

ROGERS, Y., YUILL, N. and MARSHALL, P. (2013) “Contrasting Lab-Based and In-the-Wild Studies for Evaluating Multi-User Technologies.” In B. Price (2013) *The SAGE Handbook on Digital Technology Research*. SAGE Publications: 359–173. This chapter explores the pros and cons of lab-based and in-the-wild evaluation studies with reference to different types of technology platforms including tabletops and large wall displays.

SHNEIDERMAN, B., PLAISANT, C., COHEN, M., JACOBS, S., ELMQUIST, N. and DIAKOPOULOS, N. (2016) *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (6th ed.). Addison-Wesley, Pearson. Chapter 5 provides an alternative way of categorizing evaluation methods and offers a useful overview.

TULLIS, T. and ALBERT, B. (2013) *Measuring the User Experience* (2nd ed.). Morgan Kaufmann. This book provides a more general treatment of usability testing. It focuses more strongly on evaluating the user experience and UX design.

Chapter 15

EVALUATION STUDIES: FROM CONTROLLED TO NATURAL SETTINGS

15.1 Introduction

15.2 Usability Testing

15.3 Conducting Experiments

15.4 Field Studies

Objectives

The main goals of the chapter are to accomplish the following:

- Explain how to do usability testing.
- Outline the basics of experimental design.
- Describe how to do field studies.

15.1 Introduction

Imagine that you have designed a new app to allow school children ages 9 or 10 and their parents to share caring for the class hamster over the school holidays. The app will schedule which children are responsible for the hamster and when, and it will also record when it is fed. The app will also provide detailed instructions about when the hamster is scheduled to go to another family and the arrangements about when and where it will be handed over. In addition, both teachers and parents will be able to access the schedule and send and leave messages for each other. How would you find out whether the children, their teacher, and their parents can use the app effectively and whether it is satisfying to use? What evaluation methods would you employ?

In this chapter, we describe evaluation studies that take place in a spectrum of settings, from controlled laboratories to natural settings. Within this range we focus on the following:

- *Usability testing*, which takes place in usability labs and other controlled lab-like settings
- *Experiments*, which take place in research labs
- *Field studies*, which take place in natural settings, such as people's homes, schools, work, and leisure environments



Source: <http://geek-and-poke.com>.
Licensed under CC-BY 3.0

15.2 Usability Testing

The usability of products has traditionally been tested in controlled laboratory settings. This approach emphasizes how usable a product is. Initially, it was most commonly used to evaluate desktop applications, such as websites, word processors, and search tools. It is also important now, however, to test the usability of apps and other digital products. Performing usability testing in a laboratory, or in a temporarily assigned controlled environment, enables designers to control what users do and allows them to control the environmental and social influences that might impact the user's performance. The goal is to test whether the product being developed is usable by the intended user in order to achieve the tasks for which it was designed and whether users are satisfied with their experience. For some products, such as games, designers will also want to know whether their product is enjoyable and fun to use. (Chapter 1, "What is Interaction Design," discusses usability and user experience goals.)

15.2.1 Methods, Tasks, and Users

Collecting data about users' performance on predefined tasks is a central component of usability testing. As mentioned in Chapter 14, "Introducing Evaluation," a combination of methods is often used to collect data. The data includes video recordings of the users, including their facial expressions, logged keystrokes, and mouse and other movements, such as swiping and dragging objects. Sometimes, participants are asked to describe what they are thinking and doing out loud (the "think aloud" technique) while carrying out tasks as a way of revealing what they are thinking and planning. In addition, a user satisfaction questionnaire is used to find out how users actually feel about using the product by asking them to rate it using a number of scales after they interact with it. Structured or semistructured

interviews may also be conducted with users to collect additional information about what they liked and did not like about the product. Sometimes, designers also collect data about how the product is used in the field.

Examples of the tasks that are typically given to users include searching for information, reading different typefaces (for example, Helvetica and Times), navigating through different menus, and uploading apps. Performance times and the number of the different types of actions carried out by users are the two main performance measures. Obtaining these two measures involves recording the time it takes typical users to complete a task, such as finding a website, and the number of errors that users make, such as selecting incorrect menu options when creating a visual display. The following quantitative performance measures, which were identified in the late 1990s, are still used as a baseline for collecting user performance data (Wixon and Wilson, 1997):

- Number of users completing a task successfully
- Time to complete a task
- Time to complete a task after a specified time away from the product
- Number and type of errors per task
- Number of errors per unit of time
- Number of navigations to online help or manuals
- Number of users making a particular error

A key concern when doing usability testing is the number of users that should be involved: early research suggests that 5 to 12 is an acceptable number (Dumas and Redish, 1999), though more is often regarded as being better because the results represent a larger and often broader selection of the target user population. However, sometimes it is reasonable to involve fewer users when there are budget and schedule constraints. For instance, quick feedback about a design idea, such as the initial placement of a logo on a website, can be obtained from only two or three users reporting on how quickly they spot the logo and whether they like its design. Sometimes, more users can be involved early on by distributing an initial questionnaire online to collect information about users' concerns. The main concerns can then be examined in more detail in a follow-up lab-based study with a small number of typical users.

This link provides a practical introduction to usability testing and describes how it relates to UX design: <https://icons8.com/articles/usability-practical-definition-ux-design/>.

15.2.2 Labs and Equipment

Many large companies, such as Microsoft, Google, and Apple, test their products in custom-built usability labs that consist of a main testing lab with recording equipment and an observation room where the designers can watch what is going on and how the data collected is being analyzed. There may also be a reception area where users can wait, a storage area, and a viewing room for observers. These lab spaces can be arranged to mimic superficially features of the real world. For example, when testing an office product or for use in a hotel

reception area, the lab can be set up to resemble those environments. Soundproofing and lack of windows, co-workers, and other workplace and social distractions are eliminated so that the users can concentrate on the tasks that have been set up for them to perform. While controlled environments like these enable researchers to capture data about users' uninterrupted performance, the impact that real-world interruptions can have on usability is not captured.

Typically, there are two to three wall-mounted video cameras that record the users' behavior, such as hand movements, facial expressions, and general body language. Microphones are placed near where the participants will be sitting to record their comments. Video and other data is fed through to monitors in an observation room, which is usually separated from the main lab or workroom by a one-way mirror so that designers can watch what participants are doing but not be seen by the participants. The observation room can be a small auditorium with rows of seats at different levels or, more simply, a small backroom consisting of a row of chairs facing the monitors.

Figure 15.1 shows a typical arrangement in which designers in an observation room are watching a usability test through a one-way mirror, as well as watching the data being recorded on a video monitor.



Figure 15.1 A usability laboratory in which designers watch participants on a monitor and through a one-way mirror

Source: Helen Sharp

Usability labs can be expensive and labor-intensive to run and maintain. Therefore, less expensive and more versatile alternatives started to become popular in the early and mid-1990s. The development of mobile and remote usability testing equipment also corresponded with the need to do more testing in small companies and in other venues. Mobile usability equipment typically includes video cameras, laptops, eye-tracking devices, and other measuring equipment that can be set up temporarily in an office or other space, converting it into a makeshift usability laboratory. An advantage of this approach is that the equipment can be taken into work settings, enabling testing to be done on-site, which makes it less artificial and more convenient for the participants.

An increasing number of products are specifically designed for performing mobile evaluations. Some are referred to as *lab-in-a-box* or *lab-in-a-suitcase* because they pack away neatly into a convenient carrying case. The portable lab equipment typically consists of off-the-shelf components that plug into a laptop that can record video directly to hard disk, eye-trackers

(some of which take the form of glasses for recording the user's gaze, as shown in Figure 15.2), and facial recognition systems for recording changes in the user's emotional responses.



Figure 15.2 The Tobii Glasses Mobile Eye-Tracking System

Source: Dalton et al. (2015), p. 3891. Reproduced with permission of ACM Publications

An example of a recent study in which eye-tracking glasses were used to record the eye-gaze of people in a shopping mall is reported by Nick Dalton and his colleagues (Dalton et al., 2015). The goal of this study was to find out whether shoppers pay attention to large-format plasma screen displays when wandering around a large shopping mall in London. The displays varied in size, and some contained information about directions to different parts of the mall, while others contained advertisements. Twenty-two participants (10 males and 12 females, aged 19 to 73 years old) took part in the study in which they were asked to carry out a typical shopping task while wearing Tobii Glasses Mobile Eye Tracking glasses (see Figure 15.2). These participants were told that the researchers were investigating what people look at while shopping; no mention was made of the displays. Each participant was paid £10 to participate in the study. They were also told that there would be a prize drawing after the study and that participants who won would receive a gift of up to £100 in value. Their task was to find one or more items that they would purchase if they won the prize drawing. The researchers did this so that the study was an ecologically valid in-the-wild shopping task, in which the participants focused on shopping for items that they wanted.

As the participants moved around the mall, their gaze was recorded and analyzed to determine the percentage of time that they were looking at different things. This was done by using software that converted eye-gaze movements so that they could be overlaid on a video of the scene. The researchers then coded the participants' gazes based on where they were looking (for instance, at the architecture of the mall, products, people, signage, large text, or displays). Several other quantitative and qualitative analyses were also performed. The findings from these analyses revealed that participants looked at displays, particularly large plasma screens, more than had been previously reported in earlier studies by other researchers.

Another trend in usability testing is to conduct remote, unmoderated usability testing in which users perform a set of tasks with a product in their own setting, and their interactions are logged remotely (Madathil and Greenstein, 2011). An advantage of this approach is that many

users can be tested at the same time in real-world settings, and the logged data can be automatically compiled for data analysis. For example, clicks can be tracked and counted per page when users search for specific information on websites. This approach is particularly popular in large companies such as Microsoft and Google and in companies specializing in user testing (for example, Userzoom.com) that test products used across the world. With remote testing, large numbers of participants can be recruited who are able participate at convenient times within their own time zones. As more and more products are designed for global markets, designers and researchers appreciate this flexibility. Remote testing also allows individuals with disabilities to be involved, as they can work from their own homes (Petrie et al., 2006).

15.2.3 Case Study: Testing the iPad Usability

When Apple's iPad first came onto the market, usability specialists Raluca Budiu and Jakob Nielsen from the Nielsen Norman Group conducted user tests to evaluate participants' interactions with websites and apps specifically designed for the iPad (Budiu and Nielsen, 2010). This classic study is presented here because it illustrates how usability tests are carried out and the types of modifications that are made to accommodate real-world constraints, such as having a limited amount of time to evaluate the iPad as it came onto the market. Completing the study quickly was important because Raluca Budiu and Jakob Nielsen wanted to get feedback to third-party developers, who were creating apps and websites for the iPad. These developers were designing products with little or no contact with the iPad developers at Apple, who needed to keep details about the design of the iPad secret until it was launched. There was also considerable "hype" among the general public and others before the launch, so many people were eager to know if the iPad would really live up to expectations. Because of the need for a quick first study, and to make the results public around the time of the iPad launch, a second study was carried out in 2011, a year later, to examine some additional usability issues. (Reports of both studies are available on the Nielsen Norman Group website, which suggests reading the second study first. However, in this case study, the reports are discussed in chronological order: <http://www.nngroup.com/reports/ipad-app-and-website-usability>.)

15.2.3.1 iPad Usability: First Findings from User Testing

In the first study of iPad usability, Raluca Budiu and Jakob Nielsen (Budiu and Nielsen, 2010) used two usability evaluation methods: usability testing with think-aloud in which users said what they were doing and thinking as they did it (discussed earlier in Chapter 8, "Data Gathering") and an expert review, which will be discussed in the next chapter. A key question they asked was about whether user expectations were different for the iPad as compared to the iPhone. They focused on this issue because a previous study of the iPhone showed that people preferred using apps to browsing the web because the latter was slow and cumbersome at that time. They wondered whether this would be the same for the iPad, where the screen was larger and web pages were more similar to how they appeared on the laptops or desktop computers that most people were accustomed to using at the time.

The usability testing was carried out in two cities in the United States: Fremont, California, and Chicago, Illinois. The test sessions were similar: the goal of both was to understand the typical usability issues that users encounter when using applications and accessing websites on the iPad. Seven participants were recruited. All were experienced iPhone users who had owned their phones for at least three months and who had used a variety of apps.

One reason for selecting participants who used iPhones was because they would have had previous experience in using apps and the web with a similar interaction style as the iPad.

The participants were considered to be typical users who represented the range of those who might purchase an iPad. Two participants were in their 20s, three were in their 30s, one was in their 50s, and one was in their 60s. Three were males, and four were females.

Before taking part, the participants were asked to read and sign an informed consent form agreeing to the terms and conditions of the study. This form described the following:

- What the participant would be asked to do
- The length of time needed for the study
- The compensation that would be offered for participating in the study
- The participants' right to withdraw from the study at any time
- A promise that the person's identity would not be disclosed
- An agreement that the data collected from each participant would be confidential and would not be made available to marketers or anyone other than the researchers

The Tests The session started with participants being invited to explore any application they found interesting on the iPad. They were asked to comment on what they were looking for or reading, what they liked and disliked about a site, and what made it easy or difficult for them to carry out a task. A moderator sat next to each participant, observed, and took notes. The sessions were video-recorded, and they lasted about 90 minutes each. Participants worked on their own.

After exploring the iPad, the participants were asked by the researchers to open specific apps or websites, explore them, and then carry out one or more tasks as they would have if they were on their own. Each participant was assigned the tasks in a random order. All of the apps that were tested were designed specifically for the iPad, but for some tasks the users were asked to do the same task on a website that was not specifically designed for the iPad. For these tasks, the researchers took care to balance the presentation order so that the app would be the first presented for some participants and the website would be first presented for others. More than 60 tasks were chosen from more than 32 different sites. Examples are shown in Table 15.1.

| App or website | Task |
|----------------------|--|
| iBook | Download a free copy of <i>Alice's Adventures in Wonderland</i> and read through the first few pages. |
| Craigslist | Find some free mulch for your garden. |
| <i>Time</i> Magazine | Browse through the magazine, and find the best pictures of the week. |
| Epicurious | You want to make an apple pie tonight. Find a recipe and see what you need to buy in order to prepare it. |
| Kayak | You are planning a trip to Death Valley in May this year. Find a hotel located in the park or close to the park. |

Table 15.1 Examples of some of the user tests used in the iPad evaluation (adapted from Budiu and Nielsen, 2010)

Source: <http://www.nngroup.com/reports/ipad-app-and-website-usability>. Used courtesy of the Nielsen Norman Group

ACTIVITY 15.1

1. What was the main purpose of this study?
2. What aspects are considered to be important for good usability and user experience in this study?
3. How representative do you consider the tasks outlined in Table 15.1 to be for a typical iPad user?

Comment

1. The main purpose of the study was to find out how participants interacted with the iPad by examining how they interacted with the apps and websites that they used on the iPad. The findings were intended to help designers and developers determine whether specific websites need to be developed for the iPad.
2. The definition of usability in Chapter 1 suggests that the iPad should be efficient, effective, safe, easy to learn, easy to remember, and have good utility (that is, good usability). The definition of user experience suggests that it should also support creativity and be motivating, helpful, and satisfying to use (that is, to offer a good user experience). The iPad is designed for the general public, so the range of users is broad in terms of age and experience with technology.
3. The tasks are a small sample of the total set prepared by the researchers. They cover shopping, reading, planning, and finding a recipe, which are common activities that people engage in during their everyday lives. ■

The Equipment The testing was done using a setup (see Figure 15.3) similar to the mobile usability kit described earlier. A camera recorded the participant's interactions and gestures when using the iPad and streamed the recording to a laptop computer. A webcam was also used to record the expressions on the participants' faces and their think-aloud commentary. The laptop ran software called *Morae*, which synchronized these two data streams. Up to three observers (including the moderator sitting next to the participant) watched the video streams (rather than observing the participants directly) on their laptops situated on the table so that they did not invade the participants' personal space.

Usability Problems The main findings from the study showed that the participants were able to interact with websites on the iPad but that it was not optimal. For example, links on the pages were often too small to tap on reliably, and the fonts were sometimes difficult to read. The various usability problems identified in the study were classified according to a number of well-known interaction design principles and concepts, including mental models, navigation, quality of images, problems of using a touchscreen with small target areas, lack of affordances, getting lost in the application, effects of changing orientations, working memory, and feedback received.

Getting lost in an application is an old but important problem for designers of digital products, and some participants got lost because they tapped the iPad too much and could not find a back button and could not get back to the home page. One participant said “. . . I like having everything there [on the home page]. That's just how my brain works” (Budiu and




Figure 15.3 The setup used in the Chicago usability testing sessions

Source: <http://www.nngroup.com/reports/ipad-app-and-website-usability>. Used courtesy of the Nielsen Norman Group

Nielsen, 2010, p. 58). Other problems arose because applications appeared differently in the two views possible on the iPad: portrait and landscape.

Interpreting and Presenting the Data Based on the findings of their study, Budiu and Nielsen made a number of recommendations, including supporting standard navigation. The results of the study were written up as a report that was made publicly available to app developers and the general public. It provided a summary of key findings for the general public as well as specific details of the problems the participants had with the iPad so that developers could decide whether to make specific websites and apps for the iPad.

While revealing how usable websites and apps are on the iPad, this user testing did not address how the iPad would be used in people's everyday lives. This required a field study where observations were made of how people use iPads in their own homes, at school, in the gym, and when traveling, but this did not happen because of lack of time.

ACTIVITY 15.2

1. Was the selection of participants for the iPad study appropriate? Justify your comments.
2. What might have been some of the problems with asking participants to think out loud as they completed the tasks?

Comments

1. The researchers tried to get a representative set of participants across an age and gender range with similar skill levels, that is, participants who had already used an iPhone. Ideally, it would have been good to have had additional participants to see whether the findings were more generalizable across the broad range of users for whom the iPad was designed. However, it was important to do the study as quickly as possible and get the results out to developers and to the general public.

2. If a person is concentrating hard on a task, it can be difficult to talk at the same time. This can be overcome by asking participants to work in pairs so that they talk to each other about the problems that they encounter. ■

15.2.3.2 iPad Usability: Year One

Having rushed to get their first report out when the iPad first came onto the market, Raluca Budiu and Jakob Nielsen did more tests a year later in 2011. Even though many of their recommendations (for example, designing apps with back buttons, broader use of search, and direct access to news articles by touching headlines on the front page) were implemented, there were still some problems. For example, users accidentally touched something and couldn't find their way back to their starting point. There were also magazine apps that required many steps to access a table of contents, and that led users to make mistakes when navigating through the magazine.

Normally, a second usability study would not be done just a year after the first. However, the first study was done with participants who did not have direct experience with an iPad. A year later, the researchers were able to recruit participants with at least two months' experience of using an iPad. Another reason for doing a second study so close to the first one was that many of the apps and websites with usability problems were developed without the direct involvement of the Apple iPad development team due to the need for secrecy until the iPad was officially launched onto the market.

This time, user testing was done with 16 iPad users, half men and half women. Fourteen were between 25–50 years of age, and two were older than 50. The new findings included splash screens that became boring after a while for regular users, too much information on the screen, fonts that were too small, and swiping the wrong items when several options were presented on-screen.

The first set of tests in 2010 illustrates how the researchers had to adapt their testing method to fit within a tight time period. Designers and researchers often have to modify how they go about user testing for a number of reasons. For example, in a study in Namibia, the researchers reported that questionnaires did not work well because the participants gave the responses that they thought the researchers wanted to hear (Paterson et al., 2011). However, “the interviews and observations, revealed that many participants were unable to solve all tasks and that many struggled . . . Without the interviews and observations many issues would not have been laid open during the usability evaluation” (Paterson et al., 2011, p. 245). This experience suggests that using multiple methods can reveal different usability problems. Even more important, it illustrates the importance of not taking for granted that a method used with one group of participants will work with another group, particularly when working with people from different cultures.

For another example of usability testing, see the report entitled “Case Study: Iterative Design and Prototype Testing of the *NN/g* Homepage” by Kathryn Whinton and Sarah Gibbons from the Nielsen Norman Group (August 26, 2018), which describes how user testing with prototypes is integrated into the design process. You can view this report at: <https://www.nngroup.com/articles/case-study-iterative-design-prototyping/>.

A video illustrates the usability problems that a woman had when navigating a website to find the best deal for renting a car in her neighborhood. It illustrates how usability testing can be done in person by a designer sitting with a participant. The video is called Rocket Surgery Made Easy by Steve Krug, and you can view it here: <https://www.youtube.com/watch?v=QckIzHC99Xc>.

15.3 Conducting Experiments

In research contexts, specific hypotheses are tested that make a prediction about the way users will perform with an interface. The benefits are more rigor and confidence that one interface feature is easier to understand or faster to use than another. An example of a hypothesis is that context menus (that is, menus that provide options related to the context determined by the users' previous choices) are easier to select as compared to cascading menus.

Hypotheses are often based on a theory, such as Fitts' Law (see Chapter 16, "Evaluation: Inspections, Analytics, and Models"), or previous research findings. Specific measurements provide a way of testing the hypothesis. In the previous example, the accuracy of selecting menu options could be compared by counting the number of errors made by participants when selecting from each menu type.

15.3.1 Hypotheses Testing

Typically, a hypothesis involves examining a relationship between two things, called *variables*. Variables can be independent or dependent. An *independent variable* is what the researcher manipulates (that is, selects), and in the previous example, it is the different menu types. The other variable is called the *dependent variable*, and in our example this is the time taken to select an option. It is a measure of user performance and, if our hypothesis is correct, will vary depending on the different types of menus.

When setting up a hypothesis to test the effect of the independent variable(s) on the dependent variable, it is normal to derive a *null hypothesis* and an alternative one. The null hypothesis in our example would state that there is no difference in the time it takes users to find items (that is, the selection time) between context and cascading menus. The *alternative hypothesis* would state that there is a difference between the two regarding selection time. When a difference is specified but not what it will be, it is called a *two-tailed hypothesis*. This is because it can be interpreted in two ways: either it is faster to select options from the context menu or the cascading menu. Alternatively, the hypothesis can be stated in terms of one effect. This is called a *one-tailed hypothesis*, and it would state that "it is faster to select options from context menus," or vice versa. A one-tailed hypothesis would be preferred if there was a strong reason to believe it to be the case. A two-tailed hypothesis would be chosen if there was no reason or theory that could be used to support the case that the predicted effect would go one way or the other.

You might ask why you need a null hypothesis, since it seems to be the opposite of what the experimenter wants to find out. It is put forward so that the data can reject a statement without necessarily supporting the opposite statement. If the experimental data shows a big difference between selection times for the two menu types, then the null hypothesis that the menu type has no effect on selection time can be rejected, which is different from saying that there is an effect. Conversely, if there is no difference between the two, then the null hypothesis cannot be rejected (that is, the claim that it is faster to select options from context menus is not supported).

To test a hypothesis, the researcher has to set up the conditions and find ways to keep other variables constant to prevent them from influencing the findings. This is called the *experimental design*. Examples of other variables that need to be kept constant for both types of menus might include size and screen resolution. For example, if the text is in 10-point font size in one condition and 14-point font size in the other, then it could be this difference that causes the effect (that is, differences in selection speed are due to font size). More than one condition can also be compared with the control, for example Condition 1 = Context menu; Condition 2 = Cascading menu; and Condition 3 = Scrolling.

Sometimes, a researcher might want to investigate the relationship between two independent variables, for example, age and educational background. A hypothesis might be that young people are faster at searching the web than older people and that those with a scientific background are more effective at searching the web. An experiment would be set up to measure the time it takes to complete the task and the number of searches carried out. The analysis of the data would focus on the effects of the main variables (age and background) and also look for any interactions among them.

Hypothesis testing can also be extended to include even more variables, but it makes the experimental design more complex. An example is testing the effects of age and educational background on user performance for two methods of web searching: one using a search engine and the other manually navigating through links on a website. Again, the goal is to test the effects of the main variables (age, educational background, and web searching method) and to look for any interactions among them. However, as the number of variables increases in an experimental design, it makes it more difficult to work out what is causing the results from the data.

15.3.2 Experimental Design

A concern in experimental design is to determine which participants to involve for which conditions in an experiment. The experience of participating in one condition will affect the performance of those participants if asked to participate in another condition. For example, having learned about the way the heart works using multimedia, if one group of participants was exposed to the same learning material via another medium, for instance, virtual reality, and another group of participants was not, the participants who had the additional exposure to the material would have an unfair advantage. Furthermore, it would create bias if the participants in one condition within the same experiment had seen the content and the others had not. The reason for this is that those who had the additional exposure to the content would have had more time to learn about the topic, and this would increase their chances of answering more questions correctly. In some experimental designs, however, it is possible to use the same participants for all conditions without letting such training effects bias the results.

The names given for the different designs are different-participant design, same-participant design, and matched-pairs design. In *different-participant design*, a single group of participants is allocated randomly to each of the experimental conditions so that different participants perform in different conditions. Another term used for this experimental design is *between-subjects*

design. An advantage is that there are no ordering or training effects caused by the influence of participants' experience on one set of tasks to their performance on the next set, as each participant only ever performs under one condition. A disadvantage is that large numbers of participants are needed so that the effect of any individual differences among participants, such as differences in experience and expertise, is minimized. Randomly allocating the participants and pretesting to identify any participants that differ strongly from the others can help.

In *same-participant design* (also called *within-subjects design*), all participants perform in all conditions so that only half the number of participants is needed; the main reason for this design is to lessen the impact of individual differences and to see how performance varies across conditions for each participant. It is important to ensure that the order in which participants perform tasks for this setup does not bias the results. For example, if there are two tasks, A and B, half the participants should do task A followed by task B, and the other half should do task B followed by task A. This is known as *counterbalancing*. Counterbalancing neutralizes possible unfair effects of learning from the first task, known as the *order effect*.

In *matched-participant design* (also known as *pair-wise design*), participants are matched in pairs based on certain user characteristics such as expertise and gender. Each pair is then randomly allocated to each experimental condition. A problem with this arrangement is that other important variables that have not been considered may influence the results. For example, experience in using the web could influence the results of tests to evaluate the navigability of a website. Therefore, web expertise would be a good criterion for matching participants. The advantages and disadvantages of using different experimental designs are summarized in Table 15.2.

| Design | Advantages | Disadvantages |
|--|--|--|
| Different participants (between-subjects design) | No order effects. | Many participants are needed. Individual differences among participants are a problem, which can be offset to some extent by randomly assigning to groups. |
| Same participants (within-subjects design) | Eliminates individual differences between experimental conditions. | Need to counterbalance to avoid ordering effects. |
| Matched participants (pair-wise design) | No order effects. The effects of individual differences are reduced. | Can never be sure that subjects are matched across variables that might affect performance. |

Table 15.2 The advantages and disadvantages of different allocations of participants to conditions

The data collected to measure user performance on the tasks set in an experiment usually includes response times for subtasks, total times to complete a task, and number of errors per task. Analyzing the data involves comparing the performance data obtained across the different conditions. The response times, errors, and so on, are averaged across conditions to see whether there are any marked differences. Statistical tests are then used, such as *t*-tests that statistically compare the differences between the conditions, to reveal if these are significant. For example, a *t*-test will reveal whether it is faster to select options from context or cascading menus.

15.3.3 Statistics: *t*-tests

There are many types of statistics that can be used to test the probability of a result occurring by chance, but *t*-tests are the most widely used statistical test in HCI and related fields, such as psychology. The scores, for example, time taken for each participant to select items from a menu in each condition (that is, context and cascading menus), are used to compute the means (\bar{x}) and standard deviations (SDs). The *standard deviation* is a statistical measure of the spread or variability around the mean. The *t*-test uses a simple equation to test the significance of the difference between the means for the two conditions. If they are significantly different from each other, we can reject the null hypothesis and in so doing infer that the alternative hypothesis holds. A typical *t*-test result that compared menu selection times for two groups with 9 and 12 participants each might be as follows:

$$t=4.53, p<0.05, df=19$$

The *t*-value of 4.53 is the score derived from applying the *t*-test; *df* stands for degrees of freedom, which represents the number of values in the conditions that are free to vary. This is a complex concept that we will not explain here other than to mention how it is derived and that it is always written as part of the result of a *t*-test. The *df* values are calculated by summing the number of participants in one condition minus 1 and the number of participants in the other condition minus 1. It is calculated as $df=(N_a-1)+(N_b-1)$, where N_a is the number of participants in one condition and N_b is the number of participants in the other condition. In our example, $df=(9-1)+(12-1)=19$, *p* is the probability that the effect found did not occur by chance. So, when $p<0.05$, it means that the effect found is probably not due to chance and that there is only a 5 percent possibility that it could be by chance. In other words, there most likely is a difference between the two conditions. Typically, a value of $p<0.05$ is considered good enough to reject the null hypothesis, although lower levels of *p* are more convincing, for instance, $p<0.01$ where the effect found is even less likely to be due to chance, there being only a 1 percent chance of that being the case.

15.4 Field Studies

Increasingly, more evaluation studies are being done in natural settings with either little or no control imposed on participants' activities. This change is largely a response to technologies being developed for use outside office settings. For example, mobile, ambient, IoT, and other technologies are now available for use in the home, outdoors, and in public places. Typically, field studies are conducted to evaluate these user experiences.

As mentioned in Chapter 14, evaluations conducted in natural settings are very different from those conducted in controlled environments, where tasks are set and completed in an orderly way. In contrast, studies in natural settings tend to be messy in the sense that activities often overlap and are constantly interrupted by events that are not predicted or controlled such as phone calls, texts, rain if the study is outside, and people coming and going. This follows the way that people interact with products in their everyday messy worlds, which is generally different from how they perform on fixed tasks in a laboratory setting. Evaluating how people think about, interact with, and integrate products within the settings in which they will ultimately be used, gives a better sense of how successful the products will be in the real world. The trade-off is that it is harder to test specific hypotheses about an interface because many environmental factors that influence the interaction cannot be controlled. Therefore,

it is not possible to account, with the same degree of certainty, for how people react to or use a product as can be done in controlled settings like laboratories. This makes it more difficult to determine what causes a particular type of behavior or what is problematic about the usability of a product. Instead, qualitative accounts and descriptions of people's behavior and activities are obtained that reveal how they used the product and reacted to its design.

Field studies can range in time from just a few minutes to a period of several months or even years. Data is collected primarily by observing and interviewing people, such as by collecting video, audio, field notes, and photos to record what occurs in the chosen setting. In addition, participants may be asked to fill out paper-based or electronic diaries, which run on smartphones, tablets, or other handheld devices, at particular points during the day. The kinds of reports that can be of interest include being interrupted during an ongoing activity or when they encounter a problem when interacting with a product or when they are in a particular location, as well as how, when, and if they return to the task that was interrupted. This technique is based on the experience sampling method (ESM), discussed in Chapter 8, which is often used in healthcare (Price et al., 2018). Data on the frequency and patterns of certain daily activities, such as the monitoring of eating and drinking habits, or social interactions like phone and face-to-face conversations, are often recorded. Software running on the smartphones triggers messages to study participants at certain intervals, requesting them to answer questions or fill out dynamic forms and checklists. These might include recording what they are doing, what they are feeling like at a particular time, where they are, or how many conversations they have had in the last hour.

As in any kind of evaluation, when conducting a field study, deciding whether to tell the people being observed, or asked to record information, that they are being studied and how long the study or session will last is more difficult than in a laboratory situation. For example, when studying people's interactions with an ambient display, or the displays in a shopping mall described earlier (Dalton et al. 2016), telling them that they are part of a study will likely change the way they behave. Similarly, if people are using an online street map while walking in a city, their interactions may take only a few seconds, so informing them that they are being studied would disrupt their behavior. It is also important to ensure the privacy of participants in field studies. For example, participants in field studies that run over a period of weeks or months should be informed about the study and asked to sign an informed consent form in the usual way, as mentioned in Chapter 14. In studies that last for a long time, such as those in people's homes, the designers will need to work out and agree with the participants what part of the activity is to be recorded and how. For example, if the designers want to set up cameras, they need to be situated unobtrusively, and participants need to be informed in advance about where the cameras will be and when they will be recording their activities. The designers will also need to work out in advance what to do if the prototype or product breaks down. Can the participants be instructed to fix the problem themselves, or will the designers need to be called in? Security arrangements will also need to be made if expensive or precious equipment is being evaluated in a public place. Other practical issues may also need to be considered depending on the location, product being evaluated, and the participants in the study.

The study in which the Ethnobot (Tallyn et al., 2018) was used to collect information about what users did and how they felt while walking around at the Royal Highland Show in Scotland (discussed in Chapter 14) was an example of a field study. A wide range of other studies have explored how new technologies have been used and adopted by people in their own cultures and settings. By adopted, we mean how the participants use, integrate, and adapt the technology to suit their needs, desires, and ways of living. The findings from studies

in natural settings are typically reported in the form of vignettes, excerpts, critical incidents, patterns of behavior, and narratives to show how the products are being used, adopted, and integrated into their surroundings.

15.4.1 In-the-Wild Studies

For several years now, it has become increasingly popular to conduct in-the-wild studies to determine how people use and persist in using a range of new technologies or prototypes *in situ*. The term *in-the-wild* reflects the context of the study, in which new technologies are deployed and evaluated in natural settings (Rogers, 2011). Instead of developing solutions that fit in with existing practices and settings, researchers often explore new technological possibilities that can change and even disrupt participants' behavior. Opportunities are created, interventions are installed, and different ways of behaving are encouraged. A key concern is how people react, change, and integrate the technology into their everyday lives. The outcome of conducting in-the-wild studies for different periods and at different intervals can be revealing, demonstrating quite different results from those arising out of lab studies. Comparisons of findings from lab studies and in-the-wild studies have revealed that while many usability issues can be uncovered in a lab study, the way the technology is actually used can be difficult to discern. These aspects include how users approach the new technology, the kinds of benefits that they can derive from it, how they use it in everyday contexts, and its sustained use over time (Rogers et al, 2013; Kjeldskov and Skov, 2014; Harjuniemi and Häkikila, 2018). The next case study describes a field study in which the researchers evaluated a pain-monitoring device with patients who had just had surgery.

CASE STUDY:

A field study of a pain monitoring device

Monitoring patients' pain and ensuring that the amount of pain experienced by them after surgery is tolerable is an important part of helping patients to recover. However, accurate pain monitoring is a known problem among physicians, nurses, and caregivers. Collecting scheduled pain readings takes time, and it can be difficult because patients may be asleep or may not want to be bothered. Typically, pain is managed in hospitals by nurses asking patients to rate their pain on a 1–10 scale, which is then recorded by the nurse in the patients' records.

Before launching on the field study that is the focus of our case study, Blaine Price and his colleagues (Price et al., 2018) had already spent a considerable amount of time observing patients in hospitals and talking with nurses. They had also carried out usability tests to ensure that the design of Painpad, a pain-monitoring tangible device for patients to report their pain levels, was functioning properly. For example, they checked the usability of the display and appropriateness of the device covering for the hospital environment and whether the LED display was working and was readable. In other words, they ensured that they had a well-functioning prototype for the field study that they planned to carry out.

The goal of the field study was to evaluate the use of Painpad by patients recovering from ambulatory surgery (total hip or knee replacement) in the natural environments of two UK hospitals. Painpad (see Figure 15.4) enables patients to monitor their own pain levels by

pressing the keys on the pad to record their pain rating. The researchers were interested in many aspects related to how patients interacted with Painpad, particularly on how robust and easy it was to use in the hospital environments. They also wanted to see whether the patients rated their pain every two hours as they should do and how the patients' ratings using Painpad compared with the ratings that the nurses collected. They also looked for insights about the preferences and needs of the older patients who used Painpad and for design insights around visibility, customizability, ease of operation, and the contextual factors that affected its usability in hospital environments.



Figure 15.4 Painpad, a tangible device for inpatient self-logging of pain

Source: Price et al. (2018). Reproduced with permission of ACM Publications

Data Collection and Participants

Two studies were conducted that involved 54 people (31 in one study and 23 in another). Data screening excluded participants who did not provide data using Painpad or for whom the nurses did not collect data that could be compared with the Painpad data. Because of the confidential nature of the study, ethical considerations were carefully applied to ensure that the data was stored securely and that the patients' privacy was assured. Thirteen of the patients were male, and 41 were female. They ranged in age from 32–88, with mean and median ages of 64.6 and 64.5. The time they spent in the hospital ranged from 1–7 days, with an average stay of 2–3 days.

After returning from surgery, the patients were each given a Painpad that stayed by the side of their bed. Patients were encouraged to use it at their earliest convenience. The Painpad was programmed to prompt the patients to report their pain levels every two hours. This two-hour interval was based on the hospital's desired clinical target for collecting pain data. Each time a pain rating was due, alternating red and green lights flashed on the Painpad for up to

five minutes, and an audio notification of a few seconds sounded. The patients' pain rating was automatically time-stamped by the Painpad and stored in a secure database. In addition to the pain scores collected using Painpad, the nurses also collected verbal pain scores from the patients every two hours. These scores were entered into the patients' charts and later entered into a database by a senior staff nurse and made available to the researchers for comparison with the Painpad data.

When the patients were ready to leave the second hospital mentioned, they were given a short questionnaire that asked whether Painpad was easy to use, how often they made mistakes using it, and whether they noticed the flashing light and sound notifications. They were also asked to rate how satisfied they were with Painpad on a 1–5 Likert rating scale and to make any other comments that they wanted to share about their experience in a free text field.

Data Analysis and Presentation

Three types of data analysis were used by the researchers. They examined how satisfied the patients were with Painpad based on the questionnaire responses, how the patients complied with the bi-hourly requests to rate their pain on Painpad, and how the data collected with Painpad compared with the data collected by the nurses.

Nineteen fully completed satisfaction questionnaires were collected that indicated that Painpad was well received and easy to use (mean rating 4.63 on a scale 1–5, where 5 was the highest rating) and that it was easy to remember to use it. Sixteen of the respondents commented that they never made an error entering their pain ratings, the aesthetics of Painpad were rated as “good,” and participants were “mostly satisfied” with it. Responses to the flashing lights to draw patients' attention to Painpad were polarized. Most patients noticed the lights most of the time, while others only noticed the lights sometimes, and three patients said they did not notice them at all. The effectiveness of the sound alert received a middle rating; some patients thought it was “too loud and annoying,” and others thought it was too soft. More nuanced reactions and ideas were collected from the free-text response box on the questionnaire. For example, one patient (P49) wrote, “I think it is useful for monitoring the pattern of pain over the day which can be changeable” Patient P52 commented, “A day-to-day chart might be helpful.” Some patients, who had limited dexterity or other challenges, reported how their ability to use Painpad was compromised because Painpad was sometimes hard to reach or to hear.

After removing duplicate entries, there were 824 pain scores provided by the patients using Painpad compared with 645 scores collected by the nurses. This indicated that the patients recorded more pain scores than would typically be collected in the hospital by nurses. To examine how the patients complied with using Painpad every two hours compared with the scores collected by the nurses, the researchers had to define acceptable time ranges of compliance. For example, they accepted all of the time scores that were submitted 15 minutes before and 15 minutes after the bi-hourly time schedule for reporting time scores. This analysis showed that the Painpad scores indicated stronger compliance with the two-hour schedule than with scores collected by the nurses. ■

Overall, the evaluation of Painpad indicated that it was a successful device for collecting patients' pain scores in hospitals. Of course, there are still more questions for Blaine Price and his team to investigate. An obvious one is this: "Why did the patients give more pain scores and adhere more strongly to the scheduled pain recording times with Painpad than with the nurses?"

ACTIVITY 15.3

1. Why do you think Painpad was evaluated in the field rather than in a controlled laboratory setting?
2. Two types of data were collected in the field study: pain ratings and user satisfaction questionnaires. What does each type contribute to our understanding of the design of Painpad?

Comment

1. The researchers wanted to find out how Painpad would be used by patients who had just had ambulatory surgery. They wanted to know whether the patients liked using Painpad and whether they liked its design and what problems they experienced when using it over a period of several days within hospital settings. During the early development of Painpad, the researchers carried out several usability evaluations to check that it was suitable for testing in real hospital environments. It is not possible to do a similar evaluation in a laboratory because it would be difficult, if not impossible, to create realistic and often unpredictable events that happen in hospitals (for example, visitors coming into the ward, conversations with doctors and nurses, and so forth). Furthermore, the kind of pain that patients experience after surgery does not occur, nor can it be simulated, in participants in lab studies. The researchers had already evaluated Painpad's usability, and now they wanted to see how it was used in hospitals.
2. Two kinds of data were collected. Pain data was logged on Painpad and recorded independently by the nurses every two hours. This data enabled the researchers to compare the pain data recorded using Painpad with the data collected by the nurses. A user satisfaction questionnaire was also given to some of the patients. The patients answered questions by selecting a rating from a Likert scale. The patients were also invited to give comments and suggestions in a free text box. These comments helped the researchers to get a more nuanced view of the patients' needs, likes, and dislikes. For example, they learned that some patients were hampered from taking full advantage of Painpad because of other problems, such as poor hearing and restricted movement. ■

15.4.2 Other Perspectives

Field studies may also be conducted where a behavior of interest to the researchers reveals itself only after using a particular type of software for a long time, such as a complex design program or data visualization tool. For example, the expected changes in user problem-solving strategies using a sophisticated visualization tool for knowledge discovery may emerge only after days or weeks of active use because it takes time for users to become familiar,

confident, and competent with the tool (Shneiderman and Plaisant, 2006). To evaluate the efficacy of such tools, users are best studied in realistic settings in their own workplaces so they can deal with their own data and set their own agenda for extracting insights relevant to their professional goals.

These long evaluations of how experts learn and interact with tools for complex tasks typically starts with an initial interview in which the researchers check that the participant has a problem to work on, available data, and a schedule for completion. These are fundamental attributes that have to be present for the evaluation to proceed. Then the participant will get an introductory training session with the tool, followed by 2–4 weeks of novice usage, followed by 2–4 weeks of mature usage, leading to a semistructured exit interview. Additional assistance may be provided by the researcher as needed, thereby reducing the traditional separation between researcher and participant, but this close connection enables the researcher to develop a deeper understanding of the users' struggles and successes with the tools. More data, such as daily diaries, automated logs of usage, structured questionnaires, and interviews can also be used to provide a multidimensional understanding of the weaknesses and strengths of the tool.

Sometimes, a particular conceptual or theoretical framework is adopted to guide how an evaluation is performed or how the data collected from the evaluation is analyzed (see Chapter 9, "Data Analysis"). This enables the data to be explained at a more general level in terms of specific cognitive processes, social practices such as learning, or conversational or linguistic interactions.

BOX 15.1

How Many Participants Are Needed When Carrying Out An Evaluation Study?

The answer to this question depends on the goal of the study, the type of study (such as usability, experiment, field, or another type), and the constraints encountered (for instance, schedules, budgets, recruiting representative participants, and the facilities available). Chapter 8 "Data Gathering," discussed this question more broadly. The focus here is on the types of evaluation studies discussed in this chapter: usability studies, experiments, and field studies.

Usability studies

Many professional usability consultants use to recommend 5–12 participants for studies conducted in controlled or partially controlled settings. However, as the study of the iPad illustrates, six participants generated a lot of useful data. While more participants might have been preferable, Radian Budi and Jakob Nielsen (2010) were constrained in that they needed to complete their study and release their results quickly. Since then, Radian Budi and Jakob Nielsen (2012) has said, "If you want a single number, the answer is simple: test five users in a usability study. Testing with five people lets you find almost as many usability problems as you'd find using many more test participants." Others say that as soon as the same kinds of problems start being revealed and there is nothing new, it is time to stop.

Experiments

Knowing how many participants are needed in an experiment depends on the type of experimental design, the number of dependent variables being examined, and the kinds of statistical tests that will be used. For example, if different participants are being used to test two conditions, more participants will be needed than if the same participants test both conditions. These kinds of differences in experimental design influence the type of statistics used and the number of participants needed. Therefore, consulting with a statistician or referring to books and articles such as those by Caine (2016) and Cairns (2019) is advisable. Fifteen participants is suggested as the minimum for many experiments (Cairns, 2019).

Field studies

The number of participants in a field study will vary, depending on what is of interest: it may be a family at home, a software team in an engineering firm, children in a playground, a whole community in a living lab, or even tens of thousands of people online. Although field studies may not be representative of how other groups would act, the detailed findings gleaned from these studies about how participants learn to use a technology and adapt to it over time can be very revealing. ■

In-Depth Activity

This in-depth activity continues work on the online booking facility introduced at the end of Chapter 11 and continued in Chapter 12. Using any of the prototypes that you have developed to represent the basic structure of your product, follow these instructions to evaluate it:

1. Based on your knowledge of the requirements for this system, develop a standard task (for instance, booking two seats for a particular performance).
2. Consider the relationship between yourself and your participants. Do you need to use an informed consent form? If so, prepare a suitable informed consent form. Justify your decision.
3. Select three typical users, who can be friends or colleagues, and ask them to do the task using your prototype.
4. Note the problems that each user encounters. If possible, time their performance. (If you happen to have a camera or a smartphone with a camera, you could film each participant.)
5. Since the system is not actually implemented, you cannot study it in typical settings of use. However, imagine that you are planning a controlled usability study and a field study. How would you do it? What kinds of things would you need to take into account? What sort of data would you collect, and how would you analyze it?
6. What are the main benefits and problems in this case with doing a controlled study versus studying the product in a natural setting?

Summary

This chapter described evaluation studies in different settings. It focused on controlled laboratory studies, experiments, and field studies in natural settings. A study of the iPad when it first came out and a second study conducted a year later was presented as an example of usability testing. Experimental design was then discussed that involves testing a hypothesis in a controlled research lab. The chapter ended with a discussion of field studies in which participants used prototypes and new technologies in natural settings. The Painpad example involved evaluating how patients in two hospitals, who were recovering from surgery, used a mobile device designed to enable them to self-monitor their pain levels throughout the day.

Key differences between usability testing, experiments, and field studies include the location of the study—usability lab or makeshift usability lab (and living lab or online as discussed in Chapter 14), research lab, or natural environment—and how much control is imposed. At one end of the spectrum are experiments and laboratory testing, and at the other are in-the-wild field studies. Most studies use a combination of different methods, and designers often have to adapt their methods to cope with unusual new circumstances created when evaluating the new systems being developed.

Key points

- Usability testing usually takes place in usability labs or temporary makeshift labs. These labs enable designers and researchers to control the test setting. Versions of usability testing are also conducted remotely, online, and in living labs.
- Usability testing focuses on performance measures, such as how long and how many errors are made, when completing a set of predefined tasks. Direct and indirect observation (video and keystroke logging) is conducted and supplemented by user satisfaction questionnaires and interviews.
- Mobile and remote testing systems have been developed that are more portable and affordable than usability labs. Many contain mobile eye-tracking and face recognition systems and other devices. Many companies continue to use usability labs because they provide a venue for the whole team to come together to observe and discuss how users are responding to the systems being developed.
- Experiments seek to test a hypothesis by manipulating certain variables while keeping others constant.
- The researcher controls independent variable(s) to measure dependent variable(s).
- Field studies are carried out in natural settings. They seek to discover how people interact with technology in the real world.
- Field studies that involve the deployment of prototypes or technologies in natural settings may also be referred to as in-the-wild studies.
- Sometimes the findings of a field study are unexpected, especially for in-the-wild studies in which the goal is typically to explore how novel technologies are used by participants in their own homes, places of work, or outside.

Further Reading

KELLY CAINE (2016). Local Standards for Sample Size at CHI. *Chi4good*, CHI 2016, May 7–12, 2016, San Jose, CA, USA DOI: <https://doi.org/10.1145/2858036.2858498>. In this paper, Kelly Caine points out that the CHI community is composed of researchers from a wide range of disciplines (also mentioned in Chapter 1), who use a variety of methods. Furthermore, CHI researchers often deal with constraints (for instance, access to participants for an accessibility study). Therefore, the number of participants involved in a study may be different from the number suggested in standard stats texts. The discussion in this paper is based on an analysis of papers accepted at CHI, one of the premier conferences in the field.

PAUL CAIRNS (2019). *Doing Better Statistics in Human-Computer Interaction*, Cambridge University Press. This practical book is primarily for HCI researchers when planning or completing the analysis of their data.

ANDY CRABTREE, ALAN CHAMBERLAIN, REBECCA GRINTER, MATT JONES, TOM RODDEN, and YVONNE ROGERS (2013). Introduction to the special issue of “The Turn to The Wild” *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20 (3). This collection of articles provides in-depth case studies of projects that were conducted in the wild over many years, from the widespread uptake of children’s storytelling mobile apps to the adoption of online community technologies.

JONATHON LAZAR, HEIDI J. FENG, and HARRY HOCHHEISER, (2017). *Research Methods in Human-Computer Interaction*. (2nd edition). Cambridge, MA: Elsevier/Morgan Kaufmann Publishers. Chapters 2–4 describe how to design experiments and how to perform basic statistical tests.

JAKOB NIELSEN and RALUCA BUDIU (2012). *Mobile Usability*. New Riders Press. This book asks and attempts to answer the question of how we create usability and a satisfying user experience on smartphones, tablets, and other mobile devices. There is also a wide range of recent papers available on the NN/G website: nngroup.com.

COLIN ROBSON (1994, 2011). *Experimental Design and Statistics in Psychology*. Penguin Psychology. Though now quite old, this book provides a useful introduction to experimental design and basic statistics. Another useful book by the same author is *Real World Research* (3rd ed.), published in 2011 by Blackwell Publishing.



INTERVIEW with danah boyd

danah boyd is a principal researcher at Microsoft Research, the founder and president of the Data & Society Research Institute, and a visiting professor at New York University. In her research, danah examines the intersection of technology and society with an eye to limiting how technology can be abused to reinforce inequity. danah wrote *It's Complicated: The Social Lives of Networked Teens* (Yale University Press, 2014), which examines teens' engagement with social media. She blogs at www.zephoria.org/thoughts and tweets at @zephoria.

danah, can you tell us a bit about your research and what motivates you?

I am an ethnographer who examines the interplay between technology and society. For almost a decade, I researched different aspects of social media, most notably how American teens integrate social media into their daily practices. Because of this, I've followed the rise of many popular social media services—MySpace, Facebook, YouTube, Twitter, Instagram, Snapchat, and so on. I examined what teens do on these services, but I also consider how these technologies fit into teens' lives more generally. Thus, I spent a lot of time driving around the United States talking to teens and their parents, educators and youth ministers, law enforcement, and social workers, trying to get a sense of what teens' lives look like and where technology fits in.

More recently, I've been focused on how data-driven technologies are playing a central role in many facets of society. Techniques like machine learning and other forms of artificial intelligence rely heavily on data infrastructure. But what happens when data is manipulated, abused, or biased? My goal is to examine sociotechnical vulnerabilities and imagine ways of minimizing how technology can be used to reinforce inequities or cause harm. As Melvin Kranzberg once said, "Technology is neither good nor bad; nor is it neutral." I'm trying to figure out how technological decisions intersect with cultural practices, who is affected and in what ways, and what the right points of intervention are to help construct a society that we want to live in. To do this requires moving between disciplines, sectors, and frames to get at the complexity that we've created.

Fundamentally, I'm a social scientist invested in understanding the social world. Technology shapes social dynamics, providing a fascinating vantage point for understanding cultural practices.

How would you characterize good ethnography? (Please include example(s) from your own work.)

Ethnography is about mapping cultural logics and practices. To do this successfully, it's important to dive deep into the everyday practices of a particular community

and try to understand them on their own terms. The next stage is to try to ground what one observes in a broader discourse of theory and ideas to provide a framework for understanding cultural dynamics.

Many people ask me why I bothered driving around the United States, talking to teens when I can see everything that they do online. What's visible online is only a small fraction of what people do, and it's easy to misinterpret why teens do something simply by looking at the traces of their actions. Getting into their lives, understanding their logic, and seeing how technology connects with daily practice is critically important, especially because teens don't have distinct "online" versus "offline" lives. It's all intertwined, so it's necessary to see what's going on from different angles.

Of course, this is just the data collection process. I'm also a firm believer that analysis is iterative and that it's important to include other stakeholders in that process. For over two decades, I've blogged my in-process thinking in part to enable a powerful feedback loop that I've deeply relished.

I know you have encountered some surprises—or maybe even a revelation—in your work on Facebook and MySpace. Would you tell us about it, please?

From 2006–2007, I was talking with teens in different parts of the country, and I started noticing that some teens were talking about MySpace, and some teens were talking about Facebook. In Massachusetts, I met a young woman who uncomfortably told me that the black kids in her school were on MySpace, while the white kids were on Facebook. She described MySpace as "like ghetto." I didn't enter into this project expecting to analyze race

and class dynamics in the United States, but, after her comments, I couldn't avoid them. I started diving into my data, realizing that race and class could explain the difference between which teens preferred which sites. Uncomfortable with this and totally afar from my intellectual strengths, I wrote a really awkward blog post about what I was observing. For better or worse, the BBC picked this up as a "formal report from UC Berkeley," and I received more than 10,000 messages over the next week. Some were hugely critical, with some making assumptions about me and my intentions. But the teens who wrote consistently agreed. And then two teens started pointing out to me that it wasn't just an issue of choice but an issue of movement, with some teens moving from MySpace to Facebook because MySpace was less desirable and Facebook was "safe." Anyhow, recognizing the racist and classist roots of this, I spent a lot of time trying to unpack the different language that teens used when talking about these sites in a paper called "White Flight in Networked Publics? How Race and Class Shaped American Teen Engagement with MySpace and Facebook."

This might all seem antiquated these days, but the patterns I witnessed in MySpace and Facebook continue to repeat themselves. The tensions between Snapchat and Instagram have similar patterns, as does WhatsApp versus iMessage. Moreover, the network dynamics that underpin all adoption and usage of social media are increasingly being manipulated to reinforce social divisions within society. I never imagined that the teens that I watched trying to hack the attention economy in 2004 would create a template that could be used to undermine democratic conversations around the world only a decade later.

I know you are doing a lot of work on big data and that some of that is focused on social media. What are you learning and what are your concerns for the future?

To be honest, what concerns me the most about social media and data analytics is that these technologies operate within a particular formation of financialized capitalism that prioritizes short-term profits and cancerous levels of growth over other social values, including democracy, climate sustainability, and community cohesion.

Even when data-analytics projects start from ideal places, it's hard for those ideals to stay intact as companies grow and face different kinds of financial pressure. As a result, the same technologies that could leverage data to empower communities are quickly used for exploitative purposes. I genuinely struggle to balance my love of technology with my concern that these tools will be used to magnify inequality, spread disinformation, increase climate risks, and polarize society for political purposes. ■

Chapter 16

EVALUATION: INSPECTIONS, ANALYTICS, AND MODELS

16.1 Introduction

16.2 Inspections: Heuristic Evaluation and Walk-Throughs

16.3 Analytics and A/B testing

16.4 Predictive Models

Objectives

The main goals of this chapter are to accomplish the following:

- Describe the key concepts associated with inspection methods.
- Explain how to do heuristic evaluation and walk-throughs.
- Explain the role of analytics in evaluation.
- Describe how A/B testing is used in evaluation.
- Describe how to use Fitts' law—a predictive model.

16.1 Introduction

The evaluation methods described in this book so far have involved interaction with, or direct observation of, users. In this chapter, we introduce methods that are based on understanding users through one of the following:

- Knowledge codified in heuristics
- Data collected remotely
- Models that predict users' performance

None of these methods requires users to be present during the evaluation. Inspection methods often involve a researcher, sometimes known as an *expert*, role-playing the users for whom the product is designed, analyzing aspects of an interface, and identifying potential usability problems. The most well-known methods are *heuristic evaluation* and *walk-throughs*. *Analytics* involves user interaction logging, and *A/B testing* is an experimental method. Both analytics and A/B testing are usually carried out remotely. *Predictive modeling*

involves analyzing the various physical and mental operations that are needed to perform particular tasks at the interface and operationalizing them as quantitative measures. One of the most commonly used predictive models is *Fitts' law*.

16.2 Inspections: Heuristic Evaluation and Walk-Throughs

Sometimes, it is not practical to involve users in an evaluation because they are not available, there is insufficient time, or it is difficult to find people. In such circumstances, other people, often referred to as *experts* or *researchers*, can provide feedback. These are people who are knowledgeable about both interaction design and the needs and typical behavior of users. Various inspection methods were developed as alternatives to usability testing in the early 1990s, drawing on software engineering practice where code and other types of inspections are commonly used. Inspection methods for interaction design include heuristic evaluations and walk-throughs, in which researchers examine the interface of an interactive product, often role-playing typical users, and suggest problems that users would likely have when interacting with the product. One of the attractions of these methods is that they can be used at any stage of a design project. They can also be used to complement user testing.

16.2.1 Heuristic Evaluation

In *heuristic evaluation*, researchers, guided by a set of usability principles known as *heuristics*, evaluate whether user-interface elements, such as dialog boxes, menus, navigation structure, online help, and so on, conform to tried-and-tested principles. These heuristics closely resemble high-level design principles (such as making designs consistent, reducing memory load, and using terms that users understand). Heuristic evaluation was developed by Jakob Nielsen and his colleagues (Nielsen and Mohlich, 1990; Nielsen, 1994a) and later modified by other researchers for evaluating the web and other types of systems (see Hollingshead and Novick, 2007; Budd, 2007; Pinelle et al., 2009; Harley, 2018). In addition, many researchers and practitioners have converted design guidelines into heuristics that are then applied in heuristic evaluation.

The original set of heuristics for HCI evaluation were empirically derived from the analysis of 249 usability problems (Nielsen, 1994b); a revised version of these heuristics follows (Nielsen, 2014: useit.com):

Visibility of System Status The system should always keep users informed about what is going on, through appropriate feedback and within reasonable time.

Match Between System and the Real World The system should speak the users' language, with words, phrases, and concepts familiar to the user, rather than system-oriented terms. It should follow real-world conventions, making information appear in a natural and logical order.

User Control and Freedom Users often choose system functions by mistake and will need a clearly marked emergency exit to leave the unwanted state without having to go through an extended dialog. The system should support undo and redo.

Consistency and Standards Users should not have to wonder whether different words, situations, or actions mean the same thing. The system should follow platform conventions.

Error Prevention Rather than just good error messages, the system should incorporate careful design that prevents a problem from occurring in the first place. Either eliminate error-prone conditions or check for them and present users with a confirmation option before they commit to the action.

Recognition Rather Than Recall Minimize the user's memory load by making objects, actions, and options visible. The user should not have to remember information from one part of the dialog to another. Instructions for use of the system should be visible or easily retrievable whenever appropriate.

Flexibility and Efficiency of Use Accelerators—unseen by the novice user—may often speed up the interaction for the expert user such that the system can cater to both inexperienced and experienced users. Allow users to tailor frequent actions.

Aesthetic and Minimalist Design Dialogs should not contain information that is irrelevant or rarely needed. Every extra unit of information in a dialog competes with the relevant units of information and diminishes their relative visibility.

Help Users Recognize, Diagnose, and Recover from Errors Error messages should be expressed in plain language (not codes), precisely indicate the problem, and constructively suggest a solution.

Help and Documentation Even though it is better if the system can be used without documentation, it may be necessary to provide help and documentation. Any such information should be easy to search, focused on the user's task, list concrete steps to be carried out, and not be too large.

More information about heuristic evaluation is provided at
www.nngroup.com/articles/ux-expert-reviews/

This site shows how a researcher, Wendy Bravo, used heuristics to evaluate two travel websites, Travelocity and Expedia:
<https://medium.com/@WendyBravo/heuristic-evaluation-of-two-travel-websites-13f830cf0111>

This video, developed by David Lazarus and published on May 9, 2011, provides insights into Jakob Nielsen's 10 Usability Heuristics for Interface Design. The video is still useful even though the heuristics have been updated slightly since it was made.
<http://youtu.be/hWc0Fd2AS3s>

Designers and researchers evaluate aspects of the interface against the appropriate heuristics. For example, if a new social media system is being evaluated, the designer might consider how users would add friends to their networks. Those doing the heuristic evaluation go through the interface several times, inspecting the various interaction elements and comparing them with the list of usability heuristics. During each iteration, usability problems will be identified and ways of fixing them may be suggested.

Although many heuristics apply to most products (for example, be consistent and provide meaningful feedback, especially if an error occurs), some of the core heuristics are too general for evaluating products that have come onto the market more recently, such as mobile devices, digital toys, social media, ambient devices, web services, and IoT. Many designers and researchers have therefore developed their own heuristics by tailoring Nielsen's heuristics with other design guidelines, market research, results from research studies, and requirements documents. The Nielsen/Norman Group has also taken a more detailed look at particular heuristics, such as the first heuristic listed above, "visibility of system status," (Harley, 2018a), which focuses on communication and transparency.

Exactly which heuristics are appropriate and how many are needed for different products is debatable and depends on the goals of the evaluation. However, most sets have between 5 and 10 items. This number provides a good range of usability criteria by which to judge the various aspects of a product's design. More than 10 items become difficult for those doing the evaluation to manage, while fewer than 5 items tend not to be sufficiently discriminating.

Another concern is the number of researchers needed to carry out a thorough heuristic evaluation that identifies the majority of usability problems. Empirical tests were conducted suggesting that 3–5 can typically identify up to 75 percent of the total usability problems, as shown in Figure 16.1 (Nielsen, 1994a). However, employing several researchers can be resource intensive. Therefore, the overall conclusion is that while more researchers might be better, fewer can be used—especially if the researchers are experienced and knowledgeable about the product and its intended users.




Figure 16.1 Curve showing the proportion of usability problems in an interface found by heuristic evaluation using various numbers of evaluators

Source: Nielsen and Mack (1994). Used courtesy of John Wiley & Sons, Inc.

Heuristic Evaluation for Websites

A number of different heuristic sets for evaluating websites have been developed based on Nielsen's original 10 heuristics. One of these was developed by Andy Budd after discovering that Nielsen's heuristics did not address the problems of the continuously evolving web. He also found that there was overlap between several of the guidelines and that they varied widely in terms of their scope and specificity, which made them difficult to use. An extract from these heuristics is shown in Box 16.1. Notice that a difference between these and Nielsen's original heuristics is that they place more emphasis on information content.

BOX 16.1

Extract from the Heuristics Developed by Budd (2007) That Emphasize Web Design Issues

Clarity

Make the system as clear, concise, and meaningful as possible for the intended audience.

- Write clear, concise copy.
- Only use technical language for a technical audience.
- Write clear and meaningful labels.
- Use meaningful icons.

Minimize Unnecessary Complexity and Cognitive Load

Make the system as simple as possible for users to accomplish their tasks.

- Remove unnecessary functionality, process steps, and visual clutter.
- Use progressive disclosure to hide advanced features.
- Break down complicated processes into multiple steps.
- Prioritize using size, shape, color, alignment, and proximity.

Provide Users with Context

Interfaces should provide users with a sense of context in time and space.

- Provide a clear site name and purpose.
- Highlight the current section in the navigation.

- Provide a breadcrumb trail (that is, show where the user has been in a website).
- Use appropriate feedback messages.
- Show the number of steps in a process.
- Reduce perception of latency by providing visual cues (for instance, a progress indicator) or by allowing users to complete other tasks while waiting.

Promote a Pleasurable and Positive User Experience

The user should be treated with respect, and the design should be aesthetically pleasing and promote a pleasurable and rewarding experience.

- Create a pleasurable and attractive design.
- Provide easily attainable goals.
- Provide rewards for usage and progression. ■

A similar approach to Budd's is also taken by Leigh Howells in her article entitled "A guide to heuristic website reviews" (Howells, 2011). In this article and in a more recent one by Toni Granollers (2018), techniques for making the results of heuristic evaluation more objective are proposed. This can be done either to show the occurrence of different heuristics from an evaluation or to compare the results of different researchers' evaluations, as shown in Figure 16.2. First, a calculation is done to estimate the percentage of usability problems identified by each researcher, which is then displayed around the diagram (in this case there were seven researchers). Then a single value representing the mean of all of the researchers' individual means is calculated and displayed in the center of the diagram. In addition to being able to compare the relative performance of different researchers and the overall usability of the design, a version of this procedure can be used to compare the usability of different prototypes or for comparisons with competitors' products.

Doing Heuristic Evaluations

Doing a heuristic evaluation can be broken down into three main stages (Nielsen and Mack, 1994; Muniz, 2016).

- A *briefing session*, in which the user researchers are briefed about the goal of the evaluation. If there is more than one researcher, a prepared script may be used to ensure that each person receives the same briefing.
- The *evaluation period*, in which the user researchers typically spend 1–2 hours independently inspecting the product, using the heuristics for guidance.

Typically, the researchers will take at least two passes through the interface. The first pass gives a feel for the flow of the interaction and the product's scope. The second pass allows them to focus on specific interface elements in the context of the whole product and to identify potential usability problems.

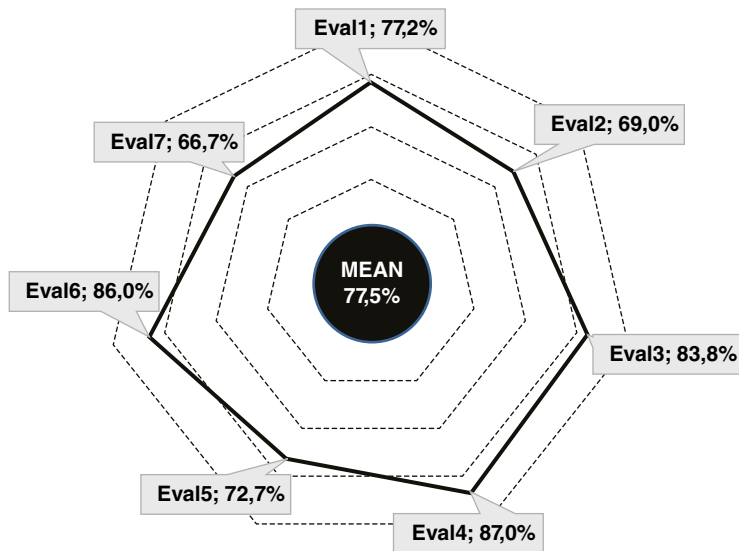


Figure 16.2 Radar diagram showing the mean number of problems identified by each of the seven researchers and the overall mean shown in the center of the diagram

Source: Granollers (2018). Used courtesy of Springer Nature

ACTIVITY 16.1

1. Use some of Budd's heuristics (Box 16.1) to evaluate a website that you visit regularly. Do these heuristics help you to identify important usability and user experience issues? If so, how?
2. How does being aware of the heuristics influence how you interact with the website?
3. Was it difficult to use these heuristics?

Comment

1. The heuristics focus on key usability criteria, such as whether the interface seems unnecessarily complex and how color is used. Budd's heuristics also encourage consideration of how the user feels about the experience of interacting with a website.
2. Being aware of the heuristics may lead to a stronger focus on the design and interaction, and it can raise awareness of what the user is trying to do and how the website is responding.
3. When applied at a high level, these guidelines can be tricky to use. For example, what exactly does "clarity" mean in regard to a website? Although the detailed list (write clear, concise copy; only use technical language for a technical audience, and so on) provides some guidance, making the evaluation task a bit easier, it may still seem quite difficult, particularly for those not used to doing heuristic evaluations. ■

If the evaluation is for a functioning product, the researchers will typically have some specific user tasks in mind so that their exploration is focused. Suggesting tasks may be helpful, but many UX researchers suggest their own tasks. However, this approach is more difficult if the evaluation is done early in design when there are only screen mock-ups or a specification. Therefore, the approach needs to be adapted for the evaluation circumstances. While working through the interface, specification, or mock-ups, a second researcher may record the problems identified, while the other researcher may think aloud, which can be video recorded. Alternatively, each researcher may take notes.

- The *debriefing session*, in which the researchers come together to discuss their findings with designers and to prioritize the problems they found and give suggestions for solutions.

The heuristics focus the researchers' attention on particular issues, so selecting appropriate heuristics is critically important. Even so, sometimes there is disagreement among researchers, as discussed in the next "Dilemma."

DILEMMA

Classic Problems or False Alarms?

Some researchers and designers may have the impression that heuristic evaluation is a panacea that can reveal all that is wrong with a design with little demand on a design team's resources. However, in addition to being quite difficult to use as just discussed, heuristic evaluation has other problems, such as sometimes missing key problems that would likely be found by testing the product with real users.

Shortly after heuristic evaluation was developed, several independent studies compared it with other methods, particularly user testing. They found that the different approaches often identify different problems and that sometimes heuristic evaluation misses severe problems (Karat, 1994). In addition, its efficacy can be influenced both by the number of experts and by the nature of the problems, as mentioned earlier (Cockton and Woolrych, 2001; Woolrych and Cockton, 2001). Heuristic evaluation, therefore, should not be viewed as a replacement for user testing.

Another issue concerns researchers reporting problems that don't exist. In other words, some of the researchers' predictions are wrong. Bill Bailey (2001) cites analyses from three published sources showing that about 33 percent of the problems reported were real usability problems, some of which were serious, while others were trivial. However, the researchers missed about 21 percent of users' problems. Furthermore, about 43 percent of the problems identified by the researchers were not problems at all; they were false alarms. He points out that this means that only about half of the problems identified were true problems: "More specifically, for every true usability problem identified, there will be a little over one false alarm (1.2) and about one-half of one missed problem (0.6). If this analysis is true, the experts tend to identify more false alarms and miss more problems than they have true hits."

How can the number of false alarms or missed serious problems be reduced? Checking that researchers really have the expertise that is required could help, particularly that they have a good understanding of the target user population. But how can this be done? One way to overcome these problems is to have several researchers. This helps to reduce the impact of

one person's experience or poor performance. Using heuristic evaluation along with user testing and other methods is also a good idea. Providing support for researchers and designers to use heuristics effectively is yet another way to reduce these shortcomings. For example, Bruce Tognazzini (2014) now includes short case studies to illustrate some of the principles that he advocates using as heuristics. Analyzing the meaning of each heuristic and developing a set of questions can also be helpful, as mentioned previously. ■

Another important issue when designing and evaluating web pages, mobile apps, and other types of products is their accessibility to a broad range of users, for example, people with sight, hearing, and mobility challenges. Many countries now have web content accessibility guidelines (WCAG) to which designers must pay attention, as discussed in Box 16.2.

BOX 16.2

Evaluating for Accessibility Using the Web Content Accessibility Guidelines

Web Content Accessibility Guidelines (WCAG) are a detailed set of standards about how to ensure that web page content is accessible for users with various disabilities (Lazar et al., 2015). While heuristics such as Ben Shneiderman's eight golden rules (Shneiderman et al., 2016) and Nielsen and Mohlich's heuristic evaluation are well-known within the HCI community, the WCAG is probably the best-known set of interface guidelines or standards outside of the HCI community. Why? Because many countries around the world have laws that require that government websites, and websites of public accommodations (such as hotels, libraries, and retail stores), are accessible for people with disabilities. A majority of those laws, including the Disability Discrimination Act in Australia, Stanca Act in Italy, Equality Act in the United Kingdom, and Section 508 of the Rehabilitation Act in the United States, as well as policies such as Canada's Policy on Communications and Federal Identity and India's Guidelines for Indian Government Websites, use WCAG as the benchmark for web accessibility.

For more information about the web accessibility guidelines, laws, and policies, see <https://www.w3.org/WAI/>

The concept of web accessibility is as old as the web itself. Tim Berners-Lee said, "The power of the Web is in its universality. Access by everyone, regardless of disability, is an essential aspect" (<https://www.w3.org/Press/IPO-announce>). To fulfill this mission, the WCAG were created, approved, and released in 1999. The WCAG were created by committee members from 475 member organizations, including leading tech companies such as Microsoft, Google, and Apple. The process for developing them was transparent and open, and all of the stakeholders, including many members of the HCI community, were encouraged to contribute and comment.

WCAG 2.0 was released in 2008. WCAG 2.1 was released in 2018, with a modification to improve accessibility further for low-vision users and for web content presented on mobile devices. In addition, when designers follow these guidelines, there are often benefits for all users, such as improved readability and search results that are presented in more meaningful ways.

While all of the various WCAG documents online would add up to hundreds of printed pages, the key concepts and core requirements are summarized in “WCAG 2.1 at a Glance,” (www.w3.org/WAI/standards-guidelines/wcag/glance) a document that could be considered to be a set of HCI heuristics.

The key concepts of web accessibility, according to WCAG, are summarized as POUR—Perceivable, Operable, Understandable, and Robust.

1. **Perceivable**

- 1.1 Provide text alternatives for non-text content.
- 1.2 Provide captions and other alternatives for multimedia.
- 1.3 Create content that can be presented in different ways, including by assistive technologies, without losing meaning.
- 1.4 Make it easier for users to see and hear content.

2. **Operable**

- 2.1 Make all functionality available from a keyboard.
- 2.2 Give users enough time to read and use content.
- 2.3 Do not use content that causes seizures or physical reactions.
- 2.4 Help users navigate and find content.
- 2.5 Make it easy to use inputs other than keyboard.

3. **Understandable**

- 3.1 Make text readable and understandable.
- 3.2 Make content appear and operate in predictable ways.
- 3.3 Help users avoid and correct mistakes.

4. **Robust**

- 4.1 Maximize compatibility with current and future user tools.

Source: <https://www.w3.org/WAI/standards-guidelines/wcag/glance/>.

These guidelines can be used as heuristics to evaluate basic web page accessibility. For example, they can be converted into specific questions such as: Is there ALT text on graphics? Is the entire page usable if a pointing device cannot be used? Is there any flashing content that will trigger seizures? Is there captioning on videos? While some of these issues can be addressed directly by designers, captioning is typically contracted out to organizations that specialize in developing and inserting captions. Governments and large organizations have to make their websites accessible to avoid possible legal action in the United States and some other countries. However, tools and advice to enable small companies and individuals to develop appropriate captions help to make captioning more universal. ■

Some researchers have created heuristics specifically to ensure that websites and other products are accessible to users with disabilities. For example, Jenn Mankoff et al. (2005) discovered that developers who did heuristic evaluation using a screen reader found 50 percent of known usability problems. Although, admirably, much research focuses on accessibility for people with sight problems, research to support other types of disabilities is also needed. An example is the research by Alexandros Yeratziotis and Panayiotis Zaphiris (2018), who created a method comprising 12 heuristics for evaluating deaf users' experiences with websites.

While automated software testing tools have been developed in an attempt to apply WCAG guidelines to web pages, this approach had limited success because there are so many accessibility requirements that are not currently machine-testable. Human inspection using the WCAG, or user testing involving people with disabilities, are still the superior methods for evaluating web compliance with WCAG 2.1 standards.

Turning Design Guidelines, Principles, and Golden Rules into Heuristics

An approach to developing heuristics for evaluating the many different types of digital technologies is to convert design guidelines into heuristics. Often this is done by just using guidelines as though they are heuristics, so guidelines and heuristics are assumed to be interchangeable. A more principled approach is for designers and researchers to translate the design guidelines into questions. For example, Kaisa Väänänen-Vainio-Mattila and Minna Wäljas (2009) adopted this approach when developing heuristics for evaluating user experience with a web service. They identified what they called *hedonic heuristics*, which directly addressed how users feel about their interactions. These were based on design guidelines concerning whether the user feels that the web service provides a lively place where it is enjoyable to spend time and whether it satisfies a user's curiosity by frequently offering interesting content. When stated as questions, these become: Is the service a lively place where it is enjoyable to spend time? Does the service satisfy users' curiosity by frequently offering interesting content?

In a critique of Nielsen's *Heuristics* (1994) and a similar set of heuristics proposed by Bruce Tognazzini's known as "First Principles of HCI Design and Usability" (Tognazzini, 2014), Toni Granollers points out the need for revising these heuristics. She claims that there is considerable overlap both within each of the two sets of heuristics and between them. Furthermore, she stresses the need for more guidance in using heuristics and advocates for developing questions as a way to provide this support. Granollers suggests first converting the heuristics into principles, and then, as was suggested earlier, identifying pertinent questions to ground the principles so that they are useful. For example, consider the heuristic "visibility and system state," which is a composite between Nielsen's and Tognazzini's heuristics. Granollers suggests the following questions:

Does the application include a visible title page, section or site? Does the user always know where they are located? Does the user always know what the system or application is doing? Are the links clearly defined? Can all actions be visualized directly (i.e., no other actions are required)?

Granollers, 2018, p. 62

Each heuristic is therefore decomposed into a set of questions like these, which could be further adapted for evaluating specific products.

Heuristics (some of which may be guidelines or rules) have been created for designing and evaluating a wide range of products including shared groupware (Baker et al., 2002),

video games (Pinelle et al., 2008), multiplayer games (Pinelle et al., 2009), online communities (Preece and Shneiderman, 2009), information visualization (Forsell and Johansson, 2010), captchas (Reynaga et al., 2015), and e-commerce sites (Hartley, 2018b). David Travis (2016), a consultant from Userfocus, has compiled 247 guidelines that are used in evaluations. These include 20 guidelines for home page usability, 20 for search usability, 29 for navigation and information architecture, 23 for trust and credibility, and more.

To access more information about these guidelines, check out David Travis’s website at <https://www.userfocus.co.uk/resources/guidelines.html>

In the mid-1980s Ben Shneiderman also proposed design guidelines that are frequently used as heuristics for evaluation. These are called the “eight golden rules.” They were slightly revised recently (Shneiderman et al., 2016) and are now stated as follows:

1. Strive for consistency.
2. Seek universal usability.
3. Offer informative feedback.
4. Design dialogs to yield closure.
5. Prevent errors.
6. Permit easy reversal of actions.
7. Keep users in control.
8. Reduce short-term memory load

ACTIVITY 16.2

COMPARING HEURISTICS

1. Compare Nielsen’s usability heuristics with Shneiderman’s eight golden rules. Which are similar, and which are different?
2. Then select another set of heuristics or guidelines for evaluating a system in which you are particularly interested and add them to the comparison.

Comment

1. Only a few heuristics and golden rules nearly match, for instance, Nielsen’s guidelines for “consistency and standards,” “error prevention,” and “user control and freedom” match up with Shneiderman’s rules of “striving for consistency,” “prevent errors,” and “keep users in control.” Looking deeper, Nielsen’s “help users recognize, diagnose and recover from errors” and “help and documentation” map with Shneiderman’s “offer informative feedback.” It is harder to find heuristics and golden rules that are unique to each researcher’s set; “aesthetic and minimalist design” appears only in Nielsen’s list, whereas “seek universal usability” appears only in Shneiderman’s list. However, with even deeper analysis, it could be argued that there is considerable overlap between the two sets. Without examining

and considering each heuristic and guideline in detail, making comparisons like this is not straightforward. It is therefore difficult to judge when faced with choosing between these and/or other heuristics. In the end, perhaps the best way forward is for researchers to select the set of heuristics that seem most appropriate for their own evaluation context.

2. We selected the web accessibility guidelines listed in Box 16.2. Unlike the Nielsen heuristics and Shneiderman's eight golden rules, these guidelines specifically target the accessibility of websites for users with disabilities, particularly those who are blind or have limited vision. The ones under "perceivable," "operable," and "robust" do not appear in either of the other two lists. The guidelines listed for "understandable" are more like those in Nielsen's and Shneiderman's lists. They focus on reminding designers to make content appear in consistent and predictable ways and to help users to avoid making mistakes. ■

16.2.2 Walk-Throughs

Walk-throughs offer an alternative approach to heuristic evaluation for predicting users' problems without doing user testing. As the name suggests, *walk-throughs* involve walking through a task with the product and noting problematic usability features. While most walk-through methods do not involve users, others, such as pluralistic walk-throughs, involve a team that may include users, as well as developers and usability specialists.

In this section, we consider cognitive and pluralistic walk-throughs. Both were originally developed for evaluating desktop systems, but, as with heuristic evaluation, they can be adapted for other kinds of interfaces.

Cognitive Walk-Throughs

Cognitive walk-throughs involve simulating how users go about problem-solving at each step in a human-computer interaction. A cognitive walk-through, as the name implies, takes a cognitive perspective in which the focus is on evaluating designs for ease of learning—a focus that is motivated by observations that users learn by exploration. This well-established method (Wharton et al., 1994) is now often integrated with a range of other evaluation and design processes. See, for example, the Jared Spool blog at <https://medium.com/@jmspool>, (Spool 2018).

The main steps involved in cognitive walk-throughs are as follows:

1. The characteristics of typical users are identified and documented, and sample tasks are developed that focus on the aspects of the design to be evaluated. A description, mock-up, or prototype of the interface to be developed is also produced, along with a clear sequence of the actions needed for the users to complete the task.
2. A designer and one or more UX researchers come together to do the analysis.
3. The UX researchers walk through the action sequences for each task, placing it within the context of a typical scenario. As they do this, they try to answer the following questions:
 - a. Will the correct action be sufficiently evident to the user?
(*Will the user know what to do to achieve the task?*)
 - b. Will the user notice that the correct action is available?
(*Can users see the button or menu item that they should use for the next action? Is it apparent when it is needed?*)

- c. Will the user associate and interpret the response from the action correctly?
(Will users know from the feedback that they have made a correct or incorrect choice of action?)
In other words, will users know what to do, see how to do it, and understand from feedback whether the action was completed correctly or not?
- 4. As the walk-through is being done, a record of critical information is compiled.
 - a. The assumptions about what would cause problems and why are identified.
 - b. Notes about side issues and design changes are made.
 - c. A summary of the results is compiled.
- 5. The design is then revised to fix the problems presented. Before making the fix, insights derived from the walk-through are often checked by testing them with real users.

When doing a cognitive walk-through, it is important to document the process, keeping account of what works and what doesn't. A standardized feedback form can be used in which answers are recorded to each question. Any negative answers are carefully documented on a separate form, along with details of the product, its version number, and the date of the evaluation. It is also useful to document the severity of the problems. For example, how likely a problem is to occur, and how serious it will be for users. The form can also be used to record the process details outlined in steps 1–4.

Brad Dalrymple (2017) describes doing a walk-through with himself as the user in three steps. Notice that there are fewer steps and they are a bit different from those previously listed.

1. Identify the user goal you want to examine.
2. Identify the tasks you must complete to accomplish that goal.
3. Document the experience while completing the tasks.

Dalrymple provides an example of the actions that he needs to go through to create a Spotify playlist (the task) of music for guests who will attend his dinner party (the goal).

Check out this link for the Dalrymple cognitive walk-through to create a Spotify playlist:
<https://medium.com/user-research/cognitive-walk-throughs-b84c4f0a14d4>

Compared with heuristic evaluation, walk-throughs focus more closely on identifying specific user problems at a detailed level. Another type of walk-through that takes a semiotic engineering perspective is described in Box 16.3.

BOX 16.3

A Semiotic Engineering Inspection Technique

Humans use a variety of signs and symbols to communicate and encode information. These include everyday things like road signs, written or spoken words, mathematical symbols, gestures, and icons. The study of how signs and symbols are constituted, interpreted, and produced is known as *semiotics*.

UX designs use a variety of signs to communicate meanings to users. Some of these are well established, such as the trashcan for deleting files, while others are created for or used only in particular types of applications, such as a row of birds in a bird identification app (see Figure 16.3). The goal of UX designers is that users of their designs understand what they mean to communicate with familiar and unfamiliar signs alike.




Figure 16.3 (a) Icons for a trashcan and (b) bird symbols

Source: (a) University of Maryland; (b) Merlin Bird ID app, Cornell Lab of Ornithology

An important aspect of UX design is how to get the designers' *message* across to the users by means of interaction *signs* alone. Knowledge of semiotic and engineering concepts—brought together in the *semiotic engineering* of human interaction with and through digital technologies (de Souza, 2005)—contributes to improving the communication of principles, features, and values of UX design.

The primary method used to evaluate the quality of semiotic engineering is *SigniFYing Message*—an inspection procedure (de Souza et al., 2016) that focuses on the communicative power of signs that UX designers can choose in order to communicate their message to users. These are the very signs through which users, in turn, will be able to express what they want to do, explore, or experience during interaction. The method is suitable for evaluating small portions of a UX design in detail. When carrying out this kind of semiotic evaluation, inspectors are guided by specific questions about three types of interaction signs.

- *Static signs*, which communicate what they mean instantly and do not require further interaction for a user to make sense of them.
- *Dynamic signs*, which only communicate meaning over time and through interaction. In other words, the user can only make sense of them if they engage in interaction.
- *Metalinguistic signs*, which can be static or dynamic. Their distinctive feature is that their meaning is an explanation, a description, some information, warning, or commentary about another interface sign.

Figure 16.4 shows examples of how these signs achieve communication within four screens of a smartphone app for arranging meetings. To help users avoid time zone errors when meeting participants who are in different time zones, UX designers may elect to communicate times using Greenwich mean time (GMT) and also expose their rationale to users.




Figure 16.4 Examples of static, dynamic, and metalinguistic signs used in UX design sketches for a meeting arrangement app

Source: de Souza et al. (2016). Used courtesy of Springer Nature

The outcome of a *SigniFYing Message* inspection is an assessment of the quality of the messages and the strategies of communication that a piece of UX design offers to the users. Using this information, designers may choose to modify the signs to clarify the communication. ■

ACTIVITY 16.3

Conduct a cognitive walk-through of typical users who want to buy a copy of this book as an ebook at www.amazon.com or www.wiley.com. Follow the steps outlined earlier by Cathleen Wharton (Wharton et al., 2009).

Comment

Step 1

Typical Users Students and professional designers who use the web regularly.

Task To buy an ebook version of this book from www.amazon.com or www.wiley.com.

Step 2

You will play the role of the expert evaluator.

Step 3

(Note that the interface for www.amazon.com or www.wiley.com may have changed since the authors did this evaluation.)

The first action will probably be to select the search box on the home page of the website selected and then type in the title or names of the author(s) of the book.

Q: Will users know what to do?

A: Yes. They know that they must find books, and the search box is a good place to start.

Q: Will users see how to do it?

A: Yes. They have seen a search box before, will type in the appropriate term, and will click the Go or Search icon.

Q: Will users understand from the feedback provided whether the action was correct or not?

A: Yes. Their action should take them to a page that shows them the cover of this book. They need to click this or a Buy icon next to the cover of the book.

Q: Will users understand from the feedback provided whether the action was correct or not?

A: Yes. They have probably done this before, and they will be able to continue to purchase the book. ■

ACTIVITY 16.4

From your experience of reading about and trying a heuristic evaluation and cognitive walk-through, how do you think they compare for evaluating a website in terms of the following?

1. The time typically needed to do each kind of evaluation
2. The suitability of each method for evaluating a whole website

Comment

1. A cognitive walk-through would typically take longer because it is a more detailed process than a heuristic evaluation.
2. A cognitive walk-through would typically not be used to evaluate a whole website unless it was a small one. A cognitive walk-through is a detailed process, whereas a heuristic evaluation is more holistic. ■

A variation of a cognitive walk-through was developed by Rick Spencer (2000) to overcome some problems that he encountered when using the original form of a cognitive walk-through for a design team. The first problem was that answering the questions and discussing the answers took too long. Second, the designers tended to be defensive, often

invoking long explanations of cognitive theory to justify their designs. This was particularly difficult because it undermined the efficacy of the method and the social relationships of team members. To cope with these problems, he adapted the method by asking fewer detailed questions and curtailing discussion. This meant that the analysis was more coarse-grained but could normally be completed in about 2.5 hours, depending on the task being evaluated by the cognitive walk-through. He also identified a leader and set strong ground rules for the session, including a ban on defending a design, debating cognitive theory, or doing designs on the fly.

More recently, Valentina Grigoreanu and Manal Mohanna (2013) modified the cognitive walk-through so that it could be used effectively within an agile design process in which a quick turnaround in design-evaluate-design cycles is needed. Their method involves an informal, simplified *streamlined cognitive walk-through* (SSCW) followed by an informal pluralistic walk-through (discussed next). When compared to a traditional user study on the same user interface, they found that approximately 80 percent of the findings from the user study were also revealed by the SSCW.

A discussion of the value of the cognitive walk-through method for evaluating various devices can be found at www.userfocus.co.uk/articles/cogwalk.html

Pluralistic Walk-Throughs

Pluralistic walk-throughs are another type of well-established walk-through in which users, developers, and usability researchers work together to step through a task scenario. As they do this, they discuss usability issues associated with dialog elements involved in the scenario steps (Nielsen and Mack, 1994). In a pluralistic walk-through, each person is asked to assume the role of a typical user. Scenarios of use, consisting of a few prototype screens, are given to each person who writes down the sequence of actions that they would take to move from one screen to another, without conferring with each other. Then they all discuss the actions they each suggested before moving on to the next round of screens. This process continues until all of the scenarios have been evaluated (Bias, 1994).

The benefits of pluralistic walk-throughs include a strong focus on users' tasks at a detailed level, that is, looking at the steps taken. This level of analysis can be invaluable for certain kinds of systems, such as safety-critical ones, where a usability problem identified for a single step could be critical to its safety or efficiency. The approach lends itself well to participatory design practices, as discussed in Chapter 12, "Design, Prototyping, and Construction," by involving a multidisciplinary team in which users play a key role. Furthermore, the researchers bring a variety of expertise and opinions for interpreting each stage of the interaction. The limitations with this approach include having to get the researchers together at one time and then proceed at the rate of the slowest. Furthermore, only a limited number of scenarios, and hence paths through the interface, can usually be explored because of time constraints.

For an overview of walk-throughs and an example of a cognitive walk-through of iTunes, see the following site:

<http://team17-cs3240.blogspot.com/2012/03/cognitive-walkthrough-and-pluralistic.html>

Note: The link to pluralistic walk-throughs may not work correctly on all browsers.

16.3 Analytics and A/B Testing

A variety of users' actions can be recorded by software automatically, including key presses, mouse or other pointing device movements, time spent searching a web page, looking at help systems, and task flow through software modules. A key advantage of logging activity automatically is that it is unobtrusive provided the system's performance is not affected, but it also raises ethical concerns about observing participants if this is done without their knowledge, as discussed in Chapter 10, "Data at Scale." Another advantage is that large volumes of data can be logged automatically and then explored and analyzed using visualization and other tools.

16.3.1 Web Analytics

Web analytics is a form of interaction logging that was specifically created to analyze users' activity on websites so that designers could modify their designs to attract and retain customers. For example, if a website promises users information about how to plant a wildflower garden but the home page is unattractive and it only shows gardens in arid and tropical regions, then users from more temperate zones will not look any further because the information they see isn't relevant to them. These users become one-time visitors and leave to look for other websites that contain the information they need to create their gardens. If the website is used by thousands of users and a small number of users do not return, this loss of users may not be noticed by the web designers and web owners unless they track users' activities.

Using web analytics, web designers and developers can trace the activity of the users who visit their website. They can see how many people came to the site, how many stayed and for how long, and which pages they visited. They can also find out about where the users came from and much more. Web analytics is therefore a powerful evaluation tool for web designers that can be used on its own or in conjunction with other types of evaluations, particularly user testing. For instance, web analytics can provide a "big-picture" overview of user interaction on a website, whereas user testing with a few typical users can reveal details about UX design problems that need to be fixed.

Because the goal of using web analytics is to enable designers to optimize users' usage of the website, web analytics is especially valued by businesses and market research organizations. For example, web analytics can be used to evaluate the effectiveness of a print or media advertising campaign by showing how traffic to a website changes during and after the campaign.

Web analytics are also used in evaluating non-transactional products such as information and entertainment websites, including hobby, music, games, blogs, and personal websites (refer to Sleeper et al., 2014), and for learning. When analytics are used in learning, they are often referred to as *learning analytics* (for example, Oviatt et al., 2013; Educause, 2016). Learning analytics play a strong role in evaluating learners' activities in massive open online courses (MOOCs) and with Open Education Resources (OERs). The designers of these systems are interested in questions such as at what point do learners tend to drop out and why?

Other types of specialist analytics have also been developed that can be used in evaluation studies, such as *visual analytics* (discussed in Chapter 10, "Data at Scale, in which thousands and often millions of data points are displayed and can be manipulated visually, as in social network analysis (Hansen et al., 2019).

Box 16.5 and Box 16.6 contain two short case examples of web analytics being used in different evaluation contexts. The first is an early example designed to evaluate visitor traffic to a website for Mountain Wines of California. The second shows the use of Google Analytics for evaluating the use of a community website for air monitoring.

A video of Simon Buckingham Shum's 2014 keynote presentation at the EdMedia 2014 Conference can be found at <http://people.kmi.open.ac.uk/sbs/2014/06/edmedia2014-keynote/> The video introduces learning analytics and how analytics are used to answer key questions in a world where people are dealing with large volumes of digital data.

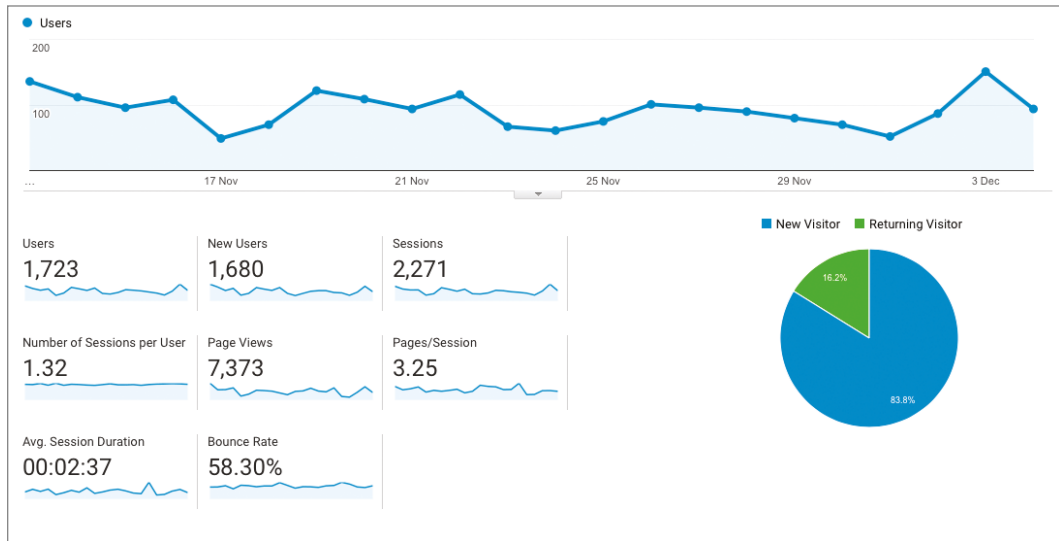
Using Web Analytics

There are two types of web analytics: on-site and off-site analytics. *On-site analytics* are used by website owners to measure visitor behavior. *Off-site analytics* measure a website's visibility and potential to acquire an audience on the Internet regardless of who owns the website. In recent years, however, the difference between off-site and on-site analytics has blurred but some people still use these terms. Additional sources may also be used to augment the data collected about a website, such as email, direct mail campaign data, sales, and history data, which can be paired with web traffic data to provide further insights into users' behavior.

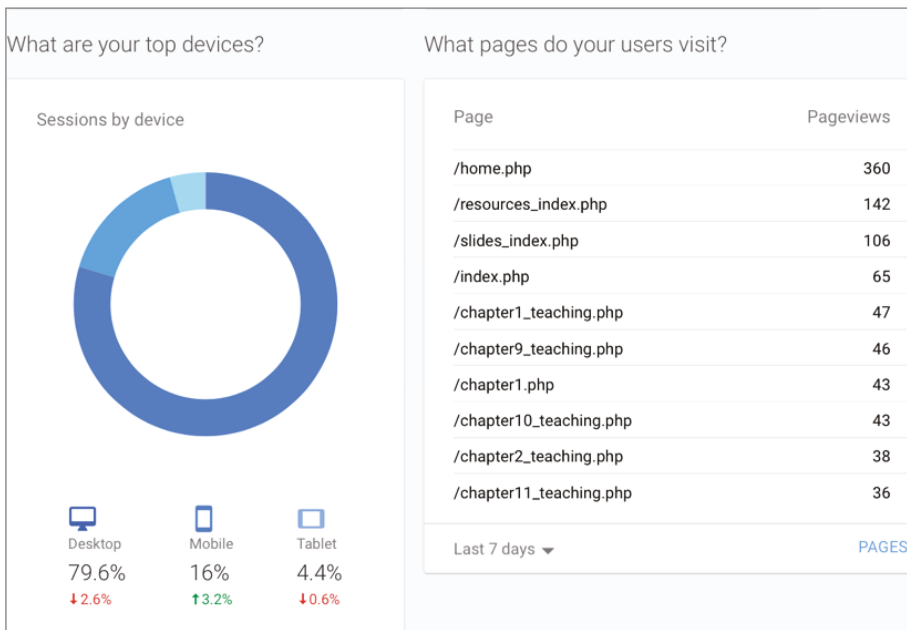
Google Analytics

Even as early as 2012, Google Analytics was the most widely used on-site web analytics and statistics service. More than 50 percent of the 10,000 most popular websites at that time (Empson, 2012) used Google Analytics, and its popularity continues to soar.

Figure 16.5 shows parts of the Google Analytics dashboard for the accompanying website for the previous edition of this book, id-book.com, for the week starting at the end of November 2018 until the beginning of December 2018. The first segment (a) shows information about who accessed the site and how long they stayed, the second segment (b) shows the devices used to view the website and the pages visited, and the third segment (c) shows the languages spoken by the users.



(a)



(b)

| Language ? | Acquisition | | | Behavior | | |
|------------|-------------------------------------|-------------------------------------|-------------------------------------|--|--|--|
| | Users ? ↓ | New Users ? | Sessions ? | Bounce Rate ? | Pages / Session ? | Avg. Session Duration ? |
| | 529 % of Total: 100.00% (529) | 462 % of Total: 100.22% (461) | 642 % of Total: 100.00% (642) | 60.28% Avg for View: 60.28% (0.00%) | 3.26 Avg for View: 3.26 (0.00%) | 00:02:31 Avg for View: 00:02:31 (0.00%) |
| 1. en-us | 317 (59.81%) | 279 (60.39%) | 391 (60.90%) | 55.50% | 3.80 | 00:03:02 |
| 2. en-gb | 44 (8.30%) | 34 (7.36%) | 52 (8.10%) | 63.46% | 2.44 | 00:01:21 |
| 3. zh-cn | 27 (5.09%) | 21 (4.55%) | 35 (5.45%) | 82.86% | 2.40 | 00:01:31 |
| 4. es-es | 12 (2.26%) | 11 (2.38%) | 13 (2.02%) | 61.54% | 2.08 | 00:00:32 |
| 5. sv-se | 11 (2.08%) | 9 (1.95%) | 13 (2.02%) | 69.23% | 1.46 | 00:01:36 |
| 6. ko-kr | 9 (1.70%) | 9 (1.95%) | 14 (2.18%) | 35.71% | 6.29 | 00:04:10 |
| 7. de-de | 6 (1.13%) | 6 (1.30%) | 6 (0.93%) | 66.67% | 3.33 | 00:00:25 |
| 8. en | 6 (1.13%) | 6 (1.30%) | 6 (0.93%) | 83.33% | 1.17 | 00:00:06 |
| 9. ar | 5 (0.94%) | 3 (0.65%) | 6 (0.93%) | 66.67% | 4.17 | 00:01:00 |
| 10. nl-nl | 5 (0.94%) | 5 (1.08%) | 5 (0.78%) | 40.00% | 2.80 | 00:01:02 |

(c)

Figure 16.5 Segments of the Google Analytics dashboard for id-book.com in December 2018: (a) audience overview, (b) the devices used to access the site, and (c) the languages of the users

ACTIVITY 16.5

Consider the three screenshot segments shown in Figure 16.5 from the Google Analytics for id-book.com, and then answer the following questions.

1. How many people visited the site during this period?
2. What do you think someone might look at in 2 minutes, 37 seconds (the average time they spent on the site)?
3. *Bounce rate* refers to the percentage of visitors who view just one page of your site. What is the bounce rate for this book, and why do you think this might be a useful metric to capture for any website?
4. Which devices are being used to access the site?
5. Which were the three largest language groups during the period, and what can you say about the bounce rate for each of them.

Comment

1. 1,723 users visited the site over this period. Notice that some users must have had more than one session since the number of users is not the same as the number of sessions, which was 2,271.

2. The number of pages viewed per session on average is about 3.25 in 2 minutes, 37 seconds. This suggests that a user probably won't have played any of the videos on the site nor read any of the case studies in any great detail. From part (b), it appears that they did check out some of the chapters, resources, and slides.
3. The bounce rate is 58.30 percent. This is a useful metric because it represents a simple but significant characteristic of user behavior, which is that after visiting the home page, they did not go anywhere else on the site. Typical bounce rates are 40–60 percent, while greater than 65 percent is high and less than 35 percent is low). If the bounce rate is high, it merits further investigation to see whether there is a problem with the website.
4. 79.6 percent of users accessed the site using a desktop, 16 percent used a mobile device, and 4.4 percent used a laptop. Compared to the previous week, the number of mobile users increased by 3.2 percent.
5. American English speakers were the largest group (317, or 59.81 percent), followed by British English speakers (44, or 8.3 percent), and then Chinese speakers (27, or 5.09 percent). The bounce rate for the Chinese visitors was by far the highest at 82.86 percent, compared with 55.5 percent for the Americans and 63.45 percent for the British visitors. ■

Ian Lurie's "Google Analytics Tutorial—Install" video explains how to install and use Google Analytics on your website. This video can be found at http://youtu.be/P_I4oc6tbYk

Scott Bradley's Google Analytics Tutorial Step-by-Step video describes the statistics included in Google Analytics, and it provides insight into how the analytics may be used to improve user traffic. This video can be found at <http://youtu.be/mm78xlsADgc>

For an overview of different dashboards that can be customized in Google Analytics, see Ned Poulter's website (2013) 6 Google Analytics Custom Dashboards to Save You Time NOW! at <http://www.stateofdigital.com/google-analytics-dashboards/>

You can also study an online course, developed by FutureLearn, on data science with Google Analytics at www.futurelearn.com/courses/data-science-google-analytics/

BOX 16.4

Other Analytics Tools

In addition to Google Analytics, other tools continue to emerge that provide additional layers of information, good access control options, and raw and real-time data collection.

Moz Analytics Tracks search marketing, social media marketing, brand activity, links, and content marketing, and it is particularly useful for link management and analysis: www.moz.com

TruSocialMetrics Tracks social media metrics, and it helps calculate social media marketing return on investment: www.truesocialmetrics.com

Clicky Comprehensive and real-time analytics tool that shows individual visitors and the actions they take, and it helps define what people from different demographics find interesting: www.clicky.com

KISSmetrics Detailed analytics tool that displays what website visitors are doing on your website before, during, and after they buy: www.kissmetrics.com

Crazy Egg Tracks visitor clicks based on where they are specifically clicking, and it creates click heat maps useful for website design, usability, and conversion: www.crazyegg.com

ClickTale Records website visitor actions and uses meta-statistics to create visual heat map reports on customer mouse movement, scrolling, and other visitor behaviors: www.clicktale.com ■

There are many sites on the web that provide lists of analytics tools. One of these, which includes some tools in addition to those mentioned in Box 16.4, is the following:

<https://www.computerworlduk.com/galleries/data/best-web-analytics-tools-alternatives-google-analytics-3628473/>

BOX 16.5

Tracking Visitors to Mountain Wines Website

In this study, Mountain Wines of California hired VisiStat to do an early study of the traffic to its website. Mountain Wines wanted to find ways to encourage more visitors to come to its website with the hope of enticing them to visit the winery. The first step to achieving this goal was to discover how many visitors were currently visiting the website, what they did there, and where they came from. Obtaining analytics about the website enabled Mountain Wines to start to understand what was happening and how to increase the number of visitors (VisiStat, 2010). Part of the results of this early analysis are shown in Figure 16.6, which provides an overview of the number of page views provided by VisiStat. Figure 16.7 shows where some of the IP addresses are located.

Using this and other data provided by VisiStat, the Mountain Wines founders could see visitor totals, traffic averages, traffic sources, visitor activity, and more. They discovered the importance of visibility for their top search words; they could pinpoint where their visitors were going on their website; and they could see where their visitors were geographically located. ■

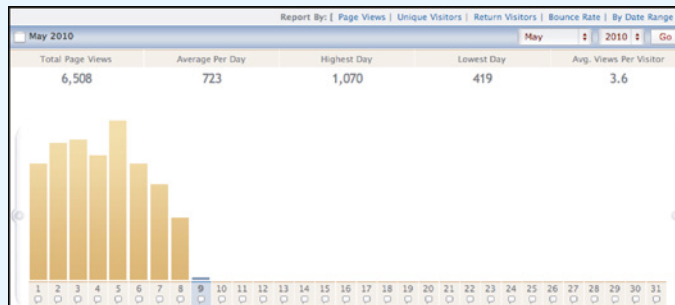


Figure 16.6 A general view of the kind of data provided by VisiStat

Source: <http://www.visistat.com/tracking/monthly-page-views.php>

The screenshot shows a table titled "Display By: Geographic Location". The table lists the top 5 unique visitors by geographic location, ranked by the number of views.

| | Unique Visitor | Views | Detail |
|----|-------------------------|-------|----------|
| 1. | Los Angeles, California | 6 | [Detail] |
| 2. | Sharpsburg, Maryland | 1 | [Detail] |
| 3. | Phoenix, Arizona | 3 | [Detail] |
| 4. | Lemesos, Limassol | 2 | [Detail] |
| 5. | Targu-mures, Mures | 1 | [Detail] |

Figure 16.7 Where the 13 visitors to the website are located by the IP address

Source: <http://www.visistat.com/tracking/monthly-page-views.php>

BOX 16.6

Using Google Analytics for Air Quality Monitoring

Many parts of the world suffer from poor air quality caused by pollution from industry, traffic congestion, and forest fires. More recently, fires in California, the northwest United States, Canada, and parts of Europe have created severe air quality problems. Consequently, communities are developing devices to crowdsource air quality readings for monitoring the quality of the air that they breathe. In one of these community-empowered air quality monitoring project, Yen-Chia Hsu and her colleagues (2017) developed a website that integrates animated smoke images, data from sensors, and crowdsourced smell reports and wind data.

Having enabled the community to monitor its own air quality and to collect reliable data to advocate for change, these researchers were eager to track users' activity on their website. They carried out a Google Analytics evaluation of the website from August 2015 to July 2016, which showed that there were 542 unique users who visited the website on 1,480 occasions for an average of 3 minutes each.

This study was innovative because, like many other local communities, this community was not technically savvy. Furthermore, developing information technology to democratize scientific knowledge and support citizen empowerment is a challenging task. However, Google Analytics, along with user testing, enabled these researchers to modify the design of the website and the associated system so that it was easier for the community to use. ■

16.3.2 A/B Testing

Another way to evaluate a website, part of a website, an application, or an app running on a mobile device is by carrying out a large-scale experiment to evaluate how two groups of users perform using two different designs—one of which acts as the control and the other as the experimental condition, that is, the new design being tested. This approach is known as *A/B testing*, and it is basically a controlled experiment but one that often involves hundreds or thousands of participants. Like the experimental design discussed in Chapter 15, “Evaluation Studies: From Controlled to Natural Settings,” A/B testing involves a “between subjects” experimental design in which two similar groups of participants are randomly selected from a single large user population (Kohavi and Longbotham, 2015), for instance, from users of social media sites such as Twitter, Facebook, or Instagram. The main differences between A/B testing and the experiments discussed in Chapter 15 is one of scale and that typically A/B testing is done online.

To do A/B testing, a variable of interest is identified, such as the design of an advertisement. Group A is served design A, the existing design, and group B is served design B, the new design. A dependent measure is then identified, such as how many times participants in each group, A and B, click the advertisement that they are presented over a particular period of time, such as a day, week, or a month. Because this is a controlled experiment, the results can be analyzed statistically to establish the probability that if a difference is observed, it is because of the treatment (in this case, the design) and not because of chance.

As Ron Kohavi (2012) mentions, A/B testing provides a valuable data-driven approach for evaluating the impact of small or large differences in the designs of web and social media sites. From front-end user-interface changes to backend algorithms, from search engines (such as Google, Bing, and Yahoo!) to retailers (for example, Amazon, eBay, and Etsy) to social networking services (such as Facebook, LinkedIn, and Twitter) to travel services (for instance, Expedia, Airbnb, and Booking.com) to many startups, online controlled experiments are now utilized to make data-driven decisions at a wide range of companies (Deng et al., 2017).

To get the most benefit from running an online A/B test, Ron Kohavi and Roger Longbotham (2015) recommend first running an *A/A test*. This is a test in which both populations of participants see the same design and should have the same experience. The results of the

A/A test are then examined, and they should show no statistically significant difference. Following this procedure ensures that the two randomly selected populations are indeed random and that the conditions under which the experiment is running are indeed similar. This is important because the Internet is complex, and users' interactions can be influenced in ways that researchers do not expect (for example by bots or the way browsers refresh or redirect), which could reduce the value of the A/B test, possibly even invalidating it.

Powerful though A/B testing may be, researchers are advised to check their plans in detail to ensure that they are testing what they expect to test. For example, Ron Kohabi and Roger Longbottom carried out an A/B test on two versions of a design for early versions of the Microsoft Office 2007 home page. The idea was to test the effectiveness of a new and more modern-looking home page with the primary objective of increasing the number of download clicks. However, instead of the number of download clicks going up as expected, it actually decreased by 64 percent. The researchers wondered what caused such an unexpected result. Upon closer examination of the two designs, they noticed the words in the new design were “Buy now” with a \$149.95 price, whereas the old design said, “Try 2007 for free” and “Buy now”. The impact of being asked to pay \$149.95 distorted the experiment, even though the new design might have actually been better. Microsoft Office has gone through many revisions since testing the 2007 version, but this example is included because it demonstrates the care that is needed when setting up A/B testing to ensure that it is actually testing the intended design features. Other design features, particularly ones that involve payments by users, can have powerful unexpected consequences that even experienced researchers like Ron Kohabi and Roger Longbottom may overlook when setting up the test.

ACTIVITY 16.6

From your knowledge of web analytics and A/B testing:

1. What would you be able to find out by using each method to evaluate a website?
2. What skills would you need to use each successfully?

Comment

1. Analytics would most likely be used to get an overview of how users are using the website. It would show who is using the website, when and for how long, where the users' IP addresses are located, bounce rates, and more. In contrast, A/B testing is a controlled experiment that enables researchers to evaluate and compare the impact of two or more UX designs. Typically, A/B testing is used to look at one or two features rather than a whole website.
2. There are many tools for evaluating websites using analytics. These tools are typically fairly straightforward to use, and with just a little knowledge designers can embed prewritten code into their designs to obtain analytics. Alternatively, there are many consultancy companies that can be hired to perform this service. In contrast, knowledge of experimental design and statistics is needed to do A/B testing. ■

16.4 Predictive Models

Like inspection methods and analytics, *predictive models* can be used to evaluate a product without users being present. Rather than user researchers being involved in role-playing during inspections, or tracking their behavior using analytics, predictive models use formulas to derive various measures of user performance. *Predictive modeling* provides estimates of the efficiency of different systems for various kinds of tasks. For example, a smartphone designer may choose to use a predictive model because it enables them to determine accurately, which is the optimal sequence of keys for performing a particular operation.

16.4.1 Fitts' Law

One kind of predictive model that has been influential in HCI and interaction design over the years is Fitts' law. *Fitts' law* (Fitts, 1954) predicts the time it takes to reach a target using a pointing device. It was originally used in human factors research to model the relationship between speed and accuracy when moving toward a target on a display. In interaction design, it has been used to model the time it takes to point at a target (for example, an icon on a screen), based on the size of the object and the distance to the object (Mackenzie, 1992). One of its main benefits is that it can help designers decide where to locate physical or digital buttons, what size to make them, and how close together to put them on a touch display or a physical device. In the early days, it was most useful for designing physical laptop/PC keyboard layouts and the placement of physical keys on mobile devices, such as smartphones, watches, and remote controls. It has also been used for designing the layout of digital displays for input on touchscreen interfaces.

Fitts' law states that:

$$T = k \log_2 (D/S + 1.0)$$

where

T = time to move the pointer to a target

D = distance between the pointer and the target

S = size of the target

k is a constant of approximately 200 ms/bit.

In a nutshell, the bigger the target, the easier and quicker it is to reach it. This is why interfaces that have big buttons are easier to use than interfaces that present lots of tiny buttons crammed together. Fitts' law also predicts that the targets accessed most quickly on any computer display are positioned at the four corners of the screen. This is because of their *pinning action*; in other words, the sides of the display constrain the user from over-stepping the target.

Fitts' law can be useful for evaluating systems where the time to locate an object physically is critical to the task at hand. In particular, it can help designers think about where to locate objects on the screen in relation to each other. This is especially useful for mobile devices, where there is limited space for placing icons and buttons on the screen. For example, in an early study carried out by Nokia, Fitts' law was used to predict text entry rates for several input methods on a 12-key cell phone keypad (Silverberg et al., 2000). The study helped the designers make decisions involving trade-offs about the size of keys, their positioning, and the sequences of keypresses to perform common tasks.

Scott MacKenzie and Robert Teather (2012) used Fitts' law in several studies including one designed to evaluate tilt as an input method for devices with built-in accelerometers, such

as touchscreen phones and tablet computers. It was also used to examine the effect of the size of the physical gap between displays and the proximity of targets in multiple-display environments (Hutchings, 2012). In addition, Fitts' law has been used to compare eye-tracking input with manual input for visual targets (Vertegaal, 2008), different ways of mapping Chinese characters to the keypad of cell phones (Liu and R ih , 2010); and gestural, touch, and mouse interaction (Sambrooks and Wilkinson, 2013). More recently, Fitts' law has been used for considering the effectiveness of new ways of input such as different game controllers (Ramcharitar and Teather, 2017), cursor positions for 3D selections in VR (Li et al., 2018), and gaze input on large displays with touch and mouse input (Rajanna and Hammond, 2018). Another creative use of Fitts' law is to evaluate the efficacy of simulating users with motor impairments interacting with a head-controlled mouse pointer system (Ritzvi et al., 2018). This application of Fitts' law was especially useful because it can be difficult to recruit participants with motor impairments to take part in user tests.

ACTIVITY 16.7

Microsoft toolbars provide the user with the option of displaying a label below each tool. Give a reason why labeled tools may be accessed more quickly. (Assume that the user knows the tool and does not need the label to identify it.)

Comment

The label becomes part of the target and hence the target gets bigger. As mentioned earlier, bigger targets can be accessed more quickly.

Furthermore, tool icons that don't have labels are likely to be placed closer together so that they are more crowded. Spreading the icons farther apart creates buffer zones of space around the icons so that if users accidentally go past the target, they will be less likely to select the wrong icon. When the icons are crowded together, the user is at greater risk of accidentally overshooting and selecting the wrong icon. The same is true of menus where the items are closely bunched together. ■

In-Depth Activity

This in-depth activity continues the work you did on the new interactive product for booking tickets at the end of Chapters 11, 12, and 15. The goal of this assignment is to evaluate the prototypes produced in the assignment from Chapter 12 by using heuristic evaluation.

1. Decide on an appropriate set of heuristics and perform a heuristic evaluation of one of the prototypes that you designed in Chapter 12.
2. Based on this evaluation, redesign the prototype to overcome the problems that you encountered.
3. Compare the findings from this evaluation with those from the usability testing in the previous chapter. What differences do you observe? Which evaluation approach do you prefer and why?

Summary

This chapter presented inspection evaluation methods, focusing on heuristic evaluation and walk-throughs, which are usually done by specialists (often referred to as experts), who role-play users' interactions with designs, prototypes, and specifications. They use their knowledge of the kinds of problems that users typically encounter, and then they offer their opinions. Heuristic evaluation and walk-throughs offer a structure to guide the evaluation process.

Analytics, in which users' interactions are logged, are often performed remotely and without users being aware that their interactions are being tracked. Large volumes of data are collected, anonymized, and statistically analyzed using specially developed software services, such as Google Analytics. The analysis provides information about how a product is used, for instance, how different versions of a website or prototype perform, or which parts of a website are seldom used—possibly because of poor usability design or lack of appeal. Data are often presented visually so that it is easier to see trends and interpret the results.

A/B testing is another form of remote testing. Fundamentally, A/B testing is a controlled experiment in which two or more dependent variables are investigated using large numbers of participants who are randomly allocated to the different experimental conditions. Small differences in the UX design of a home page can, for example, be tested using A/B testing. For sites with very large populations of users, such as popular social media sites, even small differences in design can strongly impact the number of users who use the application.

Fitts' law is an example of an evaluation method that can be used to predict user performance by determining whether a proposed interface design or keypad layout will be optimal. Typically, Fitts' law is used to compare different design layouts for virtual or physical objects, such as buttons on a device or screen.

Designers and researchers often find that they have to modify these methods, as they do for those described in the previous chapter, for use with the wide range of products that have come onto the market since they were originally developed.

Key Points

- Inspections can be used for evaluating a range of representations including requirements, mockups, prototypes, or products.
- User testing and heuristic evaluation often reveal different usability problems.
- Other types of inspections used in UX design include pluralistic and cognitive walk-throughs.
- Walk-throughs are a fine-grained, focused methods that are suitable for evaluating small parts of a product.
- Analytics involve collecting data about user interactions to identify how users use a website or product and which parts are underused.
- When applied to websites, analytics are often referred to as *web analytics*. Similarly, when applied to learning systems, they are referred to as *learning analytics*.
- Fitts' law is a predictive model that has been used in HCI to evaluate keypress sequences for handheld devices.

Further Reading

BUDI, R. and NIELSEN, J. (2012) *Mobile Usability*. New Riders Press. This book discusses why designing for mobile devices is different than designing for other systems. It describes how to evaluate these systems, including doing expert reviews, and it provides many examples.

FUTURELEARN, (2018) offers a course entitled *Data Science with Google Analytics* www.futurelearn.com/courses/data-science-google-analytics/ This course provides a good introduction for those who are new to Google Analytics. The course is regularly updated, and it is free. There is a small cost, however, if you want to buy the course materials.

GRANOLLERS, T. (2018) Usability Evaluation with Heuristics, Beyond Nielsen's list. *ACHI 2018: The Eighth International Conference on Advances in Human Computer Interaction*. 60–65. This paper provides a detailed comparison of sets of heuristics, and it suggests ways to improve heuristic evaluation.

KOHA, R. and LONGBOTHAM, R. (2015) Unexpected Results in Online Controlled Experiments. *SIGKDD Explorations* Volume 12, Issue 2, 31–35. This paper describes some of the things to look out for when doing A/B testing.

MACKENZIE, S.I. and SOUKOREFF, R. W. (2002) Text entry for mobile computing: models and methods, theory and practice. *Human-Computer Interaction*, 17, 147–198. This paper provides a useful survey of mobile text-entry techniques and discusses how Fitts' law can inform their design.

References

- Abdelnour-Nocéra, J., Clemmensen, T., and Kurosu, M. (2013b) Reframing HCI Through Local and Indigenous Perspectives. *International Journal of Human-Computer Interaction*, 29: 201–204.
- Abdul, A., Vermeulen, J., Wang, D., Lim, B.Y., and Kankanhalli, M. (2018) Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 582, 18 pages.
- Abelein, U., Sharp, H., and Paech, B. (2013) Does Involving Users in Software Development Really Influence System Success?, *IEEE Software*, Nov/Dec 2013, 13–19.
- Abowd, G., and Mynatt, E. (2000) Charting Past, Present, and Future Research in Ubiquitous Computing. *ACM Transactions on Computer–Human Interaction*, 7(1), 29–58.
- Abowd, G. D. (2012) What Next, UbiComp?: Celebrating an Intellectual Disappearing Act. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp'12)*. ACM, New York, NY, pp. 31–40.
- Abowd, G. D., Atkeson, C. G., Bobick, A. F., Essa, I. A., MacIntyre, B., Mynatt, E. D., and Starner, T. E. (2000) Living Laboratories: The Future Computing Environments Group at the Georgia Institute of Technology. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '00)*. ACM, New York, NY, pp. 215–216.
- Adams, A., and Sasses, M.A. (1999) Users Are Not The Enemy. *Communications of the ACM*, 42(12), 41–46.
- Adib, F., Mao, H., Kabelac Z., Katabi, D., and Miller, R. C. (2015) Smart Homes That Monitor Breathing and Heart Rate. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, pp. 837–846.
- Adlin, T., and Pruitt, J. (2010) *The Essential Persona Lifecycle: Your Guide to Building and Using Personas*. Morgan Kaufmann.
- Alavi, H.S., Churchill, E.F., Wiberg, M., Lalanne, D., Dalsgaard, P., Schieck, A.F., and Rogers, Y. (2019) Introduction to Human-Building Interaction (HBI)—Interfacing HCI with Architecture and Urban Design. To appear in *ACM ToCHI*.
- Alexander I., and Robertson, S. (2004) Understanding Project Sociology by Modeling Stakeholders, *IEEE Software*, 21(1), 23–27.
- Alexander, C. (1979) *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
- Al-Humairi, A., Al-Kindi, O., and Jabeur, N. (2018) Automated Musical Instruments. In *Proceedings of ICMRE 2018 Proceedings of the 2018 4th International Conference on Mechatronics and Robotics Engineering*, pp. 163–169.
- Ali, R., Arden-Close, E., and McAlaney, J. (2018) Digital Addiction: How Technology Keeps Us Hooked. *The Conversation*. Downloaded from <http://theconversation.com/digital-addiction-how-technology-keeps-us-hooked-97499>.

- Allanwood, G., and Beare, P. (2014) *User Experience Design*. Fairchild Books.
- Allison, D., Wills, B., Bowman, D., Wineman, J., and Hodges, L. (1997) The Virtual Reality Gorilla Exhibit, *IEEE Computer Graphics and Applications*, 30–38.
- Ambler, S. (2002) *Agile modeling*. John Wiley. Also at <http://www.agilemodeling.com/essays/agileDocumentationBestPractices.htm>.
- Anderson, C. (2013) *Makers*. Random House Business Books.
- Anderson, D.J. (2010) *Kanban: Successful Evolutionary Change for Your Technology Business*. Blue Hole Press.
- Antle, A. N., Corness, G., and Droumeva, M. (2009) Human–Computer–Intuition? Exploring the Cognitive Basis for Intuition in Embodied Interaction. *International Journal of Arts and Technology*, 2, 3, 235–254.
- Ardito, C., Buono, P., Costabile, D. C. M. F., and Lanzilotti, R. (2014) Investigating and Promoting UX Practice in Industry: An Experimental Study. *International Journal of Human-Computer Studies*, 72, 542–551.
- Armitage, U. (2004) Navigation and Learning in Electronic Texts. PhD thesis, Centre for HCI Design, City University London.
- Aronson-Rath, R., Milward, J., Owen, T., and Pitt, F. (2016). *Virtual Reality Journalism*. New York, NY: Columbia Journalism School.
- Ayobi, A., Sonne, T., Marshall, P., and Cox, A. L. (2018) Flexible and Mindful Self-Tracking: Design Implications from Paper Bullet Journals. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18). ACM, New York, NY, Paper 28, 14 pages.
- Babich, N. (2016) Designing Card-Based User Interfaces. Downloaded from <https://www.smashingmagazine.com/2016/10/designing-card-based-user-interfaces/>.
- Babich, N. (2018) The Do's and Don'ts of Mobile UX Design. Downloaded from <https://theblog.adobe.com/10-dos-donts-mobile-ux-design/>.
- Bachour, K., Kaplan, F., and Dillenbourg, P. (2008) Reflect: An Interactive Table for Regulating Face-to-Face Collaborative Learning. In *Proceedings of the 3rd European Conference on Technology Enhanced Learning: Times of Convergence: Technologies Across Learning Contexts*. In P. Dillenbourg and M. Specht (eds) Lecture Notes in Computer Science, 5192. Springer-Verlag, Berlin, Heidelberg, pp. 39–48.
- Bachour, K., Seied Alavi, H., Kaplan, F., and Dillenbourg, P. (2010) Low-Resolution Ambient Awareness Tools for Educational Support. In *Proceedings of CHI 2010 Workshop: The Future of HCI and Education*.
- Bailey, B. (2000) How to Improve Design Decisions by Reducing Reliance on Superstition. Let's Start with Miller's 'Magic 7.' Human Factors International, Inc. www.humanfactors.com (accessed December 16, 2010).
- Bailey, R. W. (2001) Insights from Human Factors International Inc. (HFI). Providing Consulting and Training in Software Ergonomics. January (www.humanfactors.com/home).
- Bainbridge, D. (2014) *Information Technology and Intellectual Property Law* (6th edn). Bloomsbury Professional.
- Baker, K., Greenberg, S., and Gutwin, C. (2002) Empirical Development of a Heuristic Evaluation Methodology for Shared Workspace Groupware. In *ACM Proceedings of CSCW'02 Conference*.

- Baker, M., Casey, R., Keyes, B., and Yanco, H. A. (2004) Improved Interfaces for Human–Robot Interaction in Urban Search and Rescue. In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, October.
- Balakrishnan, A. D., Kiesler, S., Cummings, J. N., and Zadeh, Reza. (2011) Research Team Integration: What it is and Why it Matters. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*, ACM Press, pp. 523–532.
- Balestrini, M., Diez, T., Marshall, P., Gluhak, A., and Rogers, Y. (2015) IoT Community Technologies: Leaving Users to Their Own Devices or Orchestration of Engagement? In *Endorsed Transactions on Internet of Things, EAI*, Vol. 15 (1).
- Bano, M., and Zowghi, D. (2015) A Systematic Review on the Relationship Between User Involvement and System Success, *Information and Software Technology*, 58, pp.148–169.
- Bano, M., Zowghi, D., and Rimini, F. (2017) User Satisfaction and System Success: An Empirical Exploration of User Involvement in Software Development. *Empirical Software Engineering*, 22, pp. 2339–2372.
- Banzy, M. (2009) *Getting Started with Arduino*. O’Reilly Media Inc.
- Barnard, P. J., Hammond, N., Maclean, A., and Morten, J. (1982) Learning and Remembering Interactive Commands in a Text Editing Task, *Behavior and Information Technology*, 1, 347–358.
- Baskinger, M. (2008) Pencils Before Pixels: A Primer in Hand-Generated Sketching, *Interactions*, March–April, 28–36.
- Bastos, J. A. D. M., Afonso, L. M., and de Souza, C.S. (2017) Metacommunication Between Programmers Through an Application Programming Interface: A Semiotic Analysis of Date and Time APIs, in *IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, 213–221.
- Baum, F. L., and Denslow, W. (1900) *The Wizard of Oz*. Random House, New York.
- Baumeister, R.F. Vohs, K.D., DeWall, C.N., and Zhang, L. (2007) How Emotion Shapes Behavior: Feedback, Anticipation, and Reflection, Rather than Direct Causation. *Personality and Social Psychology Review*, 11(2), 167–203.
- Baumer, E. P.S., Berrill, T., Botwinick, S.C., Gonzales, J.L., Ho, K., Kundrik, A., Kwon, L., LaRowe, T., Nguyen, C.P. Ramirez, F., Schaedler, P., Ulrich, W., Wallace, A., Wan, Y., and Weinfeld, B. (2018) What Would You Do? Design Fiction and Ethics. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork (GROUP ’18)*. ACM, New York, NY, pp. 244–256.
- Baumer, E. P. S., and Thomlinson, B. (2011) Comparing Activity Theory with Distributed Cognition for Video Analysis: Beyond Kicking the Tyres. In *ACM Proceedings of CHI ’11*, 133–142.
- Bauwens, V., and Genoud, P. (2014) Lessons Learned: Online Ethnography. A Tool for Creative Dialogue Between State and Citizens, *Interactions*, 60–65.
- Baxter, G., and Sommerville, I. (2011) Socio-Technical Systems: From Design Methods to Systems Engineering, *Interacting with Computers*, 23, (1) 4–17.
- Beck, K., and Andres, C. (2005) *Extreme Programming Explained: Embrace Change* (2nd edn). Addison-Wesley.
- Bell, G., Blythe, M., and Sengers, P. (2005) Making by Making Strange: Defamiliarization and the Design of Domestic Technologies, *ACM Transactions on Computer–Human Interaction*, 12(2), 149–173.

- Bergman, E., and Haitani, R. (2000) Designing the PalmPilot: A Conversation with Rob Haitani. In *Information Appliances*. Morgan Kaufmann, San Francisco.
- Bergman, O., and Whittaker, S. (2016) *The Science of Managing Our Digital Stuff*. MIT Press.
- Bernard, H. R. (2017) Direct and Indirect Observation. Chapter 14 in *Research Methods in Anthropology: Qualitative and Quantitative Approaches*. 6th edition. Rowman & Littlefield Publishers.
- Beyer, H., and Holtzblatt, K. (1998) *Contextual Design: Defining Customer-Centered Systems*. Morgan Kaufmann, San Francisco.
- Bias, R. G. (1994) The Pluralistic Usability Walk-Through – Coordinated Empathoes. In J. Nielsen and R. L. Mack (eds) *Usability Inspection Methods*. John Wiley & Sons Inc., New York.
- Bird, J., and Rogers, Y. (2010) The Pulse of Tidy Street: Measuring and Publicly Displaying Domestic Electricity Consumption. *Workshop on Energy Awareness and Conservation through Pervasive Applications, Pervasive 2010 Conference*. Downloaded from <http://www.changeproject.info/projects.html> (retrieved September 2014).
- Blandford, A., and Furniss, D. (2006) DiCoT: A Methodology for Applying Distributed Cognition to the Design of Team Working Systems. In S. W. Gilroy and M. D. Harrison (eds) *Interactive Systems: 12th International Workshop, DSVIS 2005*, Lecture Notes in Computer Science, 3941 Springer-Verlag, Berlin, Heidelberg, pp. 26–38.
- Blandford, A., Furniss, D., and Makri, S. (2017) *Qualitative HCI Research: Going Behind the Scenes*. Morgan Claypool Publishers.
- Blazevski, B., and Hallewell Haslwanter, J.D. (2017) User-Centered Development of a System to Support Assembly Line Worker. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, Article 57, 7 pages.
- Blythe M., and Cairns, P. (2009) Critical Methods and User Generated Content: The iPhone on YouTube. In *Proceedings of CHI 2009*. ACM, New York, NY, pp. 1467–1476.
- Blythe, M. (2017) Research Fiction: Storytelling, Plot and Design. In *Proceedings of Conference on Human Factors in Computing Systems (CHI'17)*. ACM, New York, NY, pp. 5400–5411.
- Bødker, S. (2000) Scenarios in User-Centered Design—Setting the Stage for Reflection and Action, *Interacting with Computers*, 13(1), 61–76.
- Boehm, B., and Basili, V.R. (2001) Software Defect Reduction Top 10 List, *IEEE Computer*, 34(1), 135–137.
- Boehner, K. Vertesi, J., Sengers, P. Dourish, P. (2007) How HCI Interprets the Probes. In *Proceedings of Conference on Human Factors in Computing Systems (CHI'07)*. ACM, New York, NY, pp. 1077–1086.
- Borchers, J. (2001) *A Pattern Approach to Interaction Design*. Wiley.
- Bornschein, J., and Weber, G. (2017) Digital Drawing Tools for Blind Users: A State-of-the-Art and Requirement Analysis. In *Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments (PETRA'17)*, ACM, New York, NY, pp. 21–28.
- Bostock, M., Ogievetsky, V., and Heer, J. (2011) D3: Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics*, 17, 12, 2301–2309.

- Bouchet, J., and Nigay, L. (2004) ICARE: A Component-Based Approach for the Design and Development of Multimodal Interfaces. In *Proceedings of CHI 2004*. ACM, New York, NY, pp. 1325–1328.
- Bowman, L.L., Levine, L.E., Waite, B.M., and Gendron, M. (2010) Can Students Really Multitask? An Experimental Study of Instant Messaging while Reading. *Computers and Education* 54, 4, 927–931.
- Bowser, A., Shilton, K., Preece, J., and Warrick, E. (2017) Accounting for Privacy in Citizen Science: Ethical Research in a Context of Openness. In *Proceedings of CSCW'17*, ACM, New York, NY, pp. 2124–2136.
- Boyd, D. (2014) *It's Complicated: The Social Lives of Networked Teens*. Yale.
- Braun, V., and Clarke, V. (2006) Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2). pp. 77–101, ISSN1478-0887.
- Brendon, G. (2017) Forget “Best” or “Sincerely,” This Email Closing Gets the Most Replies. Downloaded from <https://blog.boomerangapp.com/author/brendan/>.
- Brereton, M., and McGarry, B. (2000) An Observational Study of How Objects Support Engineering Design Thinking and Communication: Implications for the Design of Tangible Media. In *Proceedings of CHI 2000*. ACM, New York, NY, pp. 217–224.
- Briggs, G. F., Hole, G. J., and Turner, J. A. J. (2018) The Impact of Attentional Set and Situational Awareness on Dual-task Driving Performance. *Transportation Research Part F: Traffic Psychology and Behaviour*, 57, 36–47.
- Brignull, H., and Rogers, Y. (2003) Enticing People to Interact with Large Public Displays in Public Spaces. In *Proceedings of INTERACT 2003*, Zurich, pp. 17–24.
- Brkan, M. (2017) AI-Supported Decision-Making Under the General Data Protection Regulation. In *Proceedings of the 16th Edition of the International Conference on Artificial Intelligence and Law (ICAIL '17)*. ACM, New York, NY, pp. 3–8.
- Brudy, F., Houben, S., Marquardt, N., and Rogers, Y. (2016) CurationSpace: Cross-Device Content Curation Using Instrumental Interaction. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces (ISS '16)*. ACM, New York, NY, pp. 159–168.
- Buchenau, M., and Suri, J. F. (2000) Experience Prototyping. In *Proceedings of DIS 2000, Design Interactive Systems: Processes, Practices, Methods, Techniques*, pp. 17–19.
- Budd, A. (2007) *Web Heuristics*. Available at: www.andybudd.com/archives/2007/01/heuristics_for_modern_web_application_development/ (accessed September 2010).
- Budiu, R., and Nielsen, J. (2010) *Usability of iPad Apps and Websites. First Research Findings*. Nielsen Norman Group. Downloaded from www.nngroup.com/reports/mobile/ipad/ (retrieved August 2010).
- Budiu, R., and Nielsen, J. (2012) *Mobile Usability*. New Riders Press.
- Buechley, L., and Qiu, K. (2014) *Sew Electric. A Collection of DIY Projects that Combine Fabric, Electronics, and Programming*. HLT Press.
- Buolamwini, J., and Gebru, T. (2018) Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, PMLR, 81, pp.77–91.
- Burgess, P.W. (2015) Serial Versus Concurrent Multitasking: From Lab to Life. In J. M. Fawcett, E.F. Risko and A. Kingstone (eds.) *The Handbook of Attention*. 443–462. Cambridge, MA: The MIT Press.

- Buxton, B. (2007) *Sketching User Experiences*. Morgan Kaufmann, San Francisco.
- Caddick, R., and Cable, S. (2011) *Communicating the User Experience: A Practical Guide for Creating Useful UX Documentation*. Wiley.
- Caine, K. (2016) Local Standards for Sample Size at CHI. In *Proceedings of Conference on Human Factors in Computing Systems (CHI'16)*. ACM, New York, NY, pp. 981–992.
- Caird, J.K., Simmons, S.M., Wiley, K., Johnston, K.A., and Horrey, W.J. (2018) Does Talking on a Cell Phone, With a Passenger, or Dialing Affect Driving Performance? An Updated Systematic Review and Meta-Analysis of Experimental Studies. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, **60** (1).
- Cairns, P. (2019) *Doing Better Statistics in Human-Computer Interaction: Short Essays on Some Common Questions*. Oxford University Press: Oxford, UK.
- Campbell, J. L., Brown, J. L., Graving, J. S., Richard, C. M., Lichty, M. G., Sanquist, T., . . . and Morgan, J. L. (2016, December). Human Factors Design Guidance for Driver-Vehicle Interfaces. Report No. DOT HS 812 360. Washington.
- Card, S. K., Mackinley, J. D., and Shneiderman, B. (eds) (1999) *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, San Francisco.
- Card, S. K., Moran, T. P., and Newell, A. (1983) *The Psychology of Human-Computer Interaction*. Lawrence Earlbaum Associates, Hillsdale, NJ.
- Carroll J. M. (2000) Introduction to the Special Issue on Scenario-Based Systems Development, *Interacting with Computers*, **13**(1), 41–42.
- Carroll, J. M. (2004) Beyond Fun, *Interactions*, **11**(5), 38–40.
- Carroll, J. M. (ed.) (2003) *HCI Models, Theories and Frameworks: Towards a Multidisciplinary Science*. Morgan Kaufmann, San Francisco.
- Carter, S., and Mankoff, J. (2005) When Participants Do the Capturing: The Role of Media in Diary Studies. In *Proceedings of CHI 2005*. ACM, New York, NY, pp. 899–908.
- Chang, C., Hinze, A., Bowen, J., Gilbert, L., and Starkey, N. (2018) Mymemory: A Mobile Memory Assistant for People with Traumatic Brain Injury. *International Journal of Human-Computer Studies*, **117**, 4–19.
- Chang, T. (2004) The Results of Student Ratings: Paper vs Online, *Journal of Taiwan Normal University: Education*, **49**(1), 171–186.
- Charmaz, K. (2014) *Constructing Grounded Theory* (2nd edn). SAGE Publications.
- Chen, X., Zou, Q., Fan, B, Zheng, Z., and Luo, X. (2018) Recommending Software Features for Mobile Applications Based on User Interface Comparison, *Requirements Engineering* (online), 1–15, Available: <http://link.springer.com/10.1007/s00766-018-0303-4>
- Chidziwisano, G., and Wyche, S. (2018) M-Kulinda: Using a Sensor-Based Technology Probe to Explore Domestic Security in Rural Kenya. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 10, 13 pages.
- Chuang, L.L., and Pfeil, U. (2018). Transparency and Openness Promotion Guidelines for HCI. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, Paper SIG04, 4 pages.
- Churchill, E. (2018) Data, Design, and Ethnography, *Interactions*, 22–23.
- Churchill, E., Bowser, A., and Preece, J. (2013) Teaching and Learning Human-Computer Interaction. *Interactions*, **20**, 2, 44–53.

- Clegg, T., Preece, J., Warrick, E., Pauw, D., Boston, C., and Cameron, J. (2019). Community-Driven Informal Adult Environmental Learning: Using Theory as a Lens to Identify Steps Toward Concientización. To appear in *Journal of Environmental Education*.
- Clemmensen, T., Hertzum, M., Hornbaek, K., Shi, Q., and Yammiyavar, P. (2008) Cultural Cognition in the Thinking-Aloud Method for Usability Evaluation. In *Proceedings of 29th International Conference on Information Systems*. Paris, 2008, Paper 189.
- Cliffe, A. D. (2017) A Review of the Benefits and Drawbacks to Virtual Field Guides in Today's Geoscience Higher Education Environment. *International Journal of Educational Technology in Higher Education*. 14, 28.
- Cline, D.H. (2012) Six Degrees of Alexander: Social Network Analysis as a Tool for Ancient History. *AHB*, 26, 59–70.
- Cobb, S., Beardon, L., Eastgate, R., Glover, T., Kerr, S., Neale, H., Parsons, S., Benford, S., Hopkins, E., Mitchell, P., Reynard, G., and Wilson, J. (2002) Applied Virtual Environments to Support Learning of Social Interaction Skills in Users with Asperger's Syndrome, *Digital Creativity*, 13(1), N-22.
- Cockton, G., and Woolrych, A. (2001) Understanding Inspection Methods: Lessons from an Assessment of Heuristic Evaluation. In A. Blandford and J. Vanderdonck (eds), *People and Computers XV*. Springer-Verlag, Berlin, pp. 171–191.
- Cohen, J. (1960) A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1), 37–46.
- Cohen, M., Giangola, J. P., and Balogh, J. (2004) *Voice User Interface Design*. Addison-Wesley, Harlow, Essex.
- Constantine, L. L., and Lockwood, L. A. D. (1999) *Software for Use*. Addison-Wesley, Harlow, Essex.
- Cooper, A. (1999) *The Inmates are Running the Asylum*. SAMS, Indianapolis.
- Cooper, A. (2018) When Companies Question the Value of Design. Downloaded from <https://medium.com/s/story/whats-the-roi-of-ux-c47defb033d2>.
- Corbin, J. M., and Strauss, A. (2014) *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. SAGE Publications.
- Costa, N.A., Holder, E., and MacKinnon, S.N. (2017) Implementing Human Centred Design in the Context of a Graphical User Interface Redesign for Ship Manoeuvring. *International Journal of Human-Computer Studies*, 100, 55–65.
- Coyle, A. (1995) Discourse Analysis. In G. M. Breakwell, S. Hammond and C. Fife-Schaw (eds) *Research Methods in Psychology*. SAGE, London.
- Crabtree, A. (2003) *Designing Collaborative Systems: A Practical Guide to Ethnography*. Springer-Verlag, Berlin.
- Craik, K. J. W. (1943) *The Nature of Explanation*. Cambridge University Press, Cambridge.
- Cramer, H., Evers, V., Ramlal, S., Someren, M., Rutledge, L., Stash, N., Aroyo, L., and Wielinga, B. (2008) The Effects of Transparency on Trust in and Acceptance of a Content-Based Art Recommender. *User Model User-Adap Inter*, 18, 455.
- Crampton Smith, G. (1995) The Hand that Rocks the Cradle. *ID Magazine* May/June, 60–65.
- Crowcroft, J., Haddadi, H., and Henderson, T. (2018) Responsible Research on Social Networks: Dilemmas and Solutions. In *The Oxford Handbook of Networked*

- Communication*. Brooke Foucault Welles and Sandra González-Bailón (eds.) Oxford Handbooks online.
- Crumlish, C., and Malone, E. (2009) *Designing Social Interfaces: Principles, Patterns and Practices for Improving the User Experience*, O'Reilly.
- Csikszentmihalyi, M. (1996) Go with the flow. *Wired Interview*. www.wired.com/wired/archive/4.09/czik.html (retrieved May 6, 2005).
- Csikszentmihalyi, M. (1997) *Finding Flow: The Psychology of Engagement with Everyday Life*. Basic Books, New York.
- Cutting, J., Gundry, D., and Cairns, P. (2019) Busy Doing Nothing? What Do Players Do in Idle Games? *International Journal of Human-Computer Studies*, **122**, 133–144.
- Cybil, C., Perry, M., Laurier, E., and Taylor, A. (2013) 'Eyes Free' In-Car Assistance: Parent and Child Passenger Collaboration During Phone Calls. In *Proceedings of MobileHCI'13*, ACM Press, pp. 332–341.
- Da Silva, T., Silveira, M. S., Maurer, F., and Silveria, F.F. (2018) The Evolution of Agile UX. *Information and Software Technology*, **102**, 1–5.
- Dalton, N.S., Collins, E., and Marshall, P. (2015) Display Blindness?: Looking Again at the Visibility of Situated Displays Using Eye-Tracking. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, pp. 3889–3898.
- Davis, K. (2012) *Ethics of Big Data*. O'Reilly.
- de Rada, D., and Dominguez-Alvarez (2014) Response Quality of Self-Administration Questionnaires: A Comparison Between Paper and Web Questionnaires. *Social Science Computer Review*, **32**(2), 256–269.
- de Souza, C. S. (2005) *The Semiotic Engineering of Human-Computer Interaction*. Cambridge, MA. The MIT Press.
- de Souza, C. S., Cerqueira, R., Afonso, L., Brandão, R., Ferreira, J. (2016) *Software Developers as Users: Semiotic Investigations in Human-Centered Software Development*. Springer.
- Deng, A., and Shi, X. (2016) Data-Driven Metric Development for Online Controlled Experiments: Seven Lessons Learned. In *Proceedings of 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*, pp. 77–86.
- Deng, A., Dmitriev, P., Gupta, S., Raff, R. K. P., and Vermeer, L (2017) A/B Testing at Scale: Accelerating Software Innovation. In *Proceedings of Special Interest Group on Information Retrieval '17*, pp. 1397–1397.
- Deng, L., and Huang, X. (2004) Challenges in Adopting Speech Recognition, *Communications of the ACM*, **47**(1), 69–75.
- Denzin, N. (2006). *Sociological Methods: A Sourcebook*. Aldine Transaction (5th edn), ISBN 9780-202308401.
- Denzin, N. K., and Lincoln, Y. S. (2011) *The SAGE Handbook of Qualitative Research*. SAGE Publications.
- Deshpande, A., Sharp, H., Barroca, L., and Gregory, A.J. (2016) Remote Working and Collaboration in Agile Teams. In *Proceedings of International Conference on Information Systems*, pp. 4543–4559.

- Dhillon, B., Banach, P., Hocielnik, R., Emparanza, J.P., Politis, I., Paczewska, A., and Markopoulos, P. (2011) Visual Fidelity of Video Prototypes and User Feedback: A Case Study, In *Proceedings of BCS-HCI*.
- Diaz de Rada, V., and Dominguez-Alvarez, J. A. (2014) Response Quality of Self-Administered Questions: A Comparison Between Paper and Web. *Social Science Computer Review*, 32, 2, SAGE Publications.
- Dietz, P. H., and Leigh, D. L. (2001) DiamondTouch: A Multi-User Touch Technology. In *Symposium on User Interface Software and Technology (UIST)*. ACM, New York, NY, pp. 219–226.
- Diez, T., and Posada, A. (2013) The Fab and the Smart City: The Use of Machines and Technology for the City Production by Its Citizens. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction*, ACM Press, pp. 447–454.
- DiSalvo, C., Sengers, P., and Brynjarsdottir, H. (2010) Mapping the Landscape of Sustainable HCI. In *Proceedings of CHI 2010*. ACM, New York, NY, pp. 1975–1984.
- Dix, A., Finlay, J., Abowd, G., and Beale, R. (2004) *Human–Computer Interaction* (3rd edn). Pearson Education, Harlow, Essex.
- Douglas, H.E., Raban, M. Z., Walter, S. R., and Westbrook, J. I. (2017) Improving Our Understanding of Multi-Tasking in Healthcare: Drawing Together the Cognitive Psychology and Healthcare Literature. *Applied Ergonomics*, Volume 59, Part A, 45–55.
- Dourish, P. (2001) *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press, Cambridge, MA.
- Dourish, P., and Bly, S. (1992) Portholes: Supporting Awareness in a Distributed Work Group. In *Proceedings of CHI '92*. ACM, New York, NY, pp. 541–547.
- Drascic, D., and Milgram, P. (1996) Perceptual Issues in Augmented Reality. In M. T. Bolas, S. S. Fisher and J. O. Merritt (eds) *SPIE Volume 2653: Stereoscopic Displays and Virtual Reality Systems III*. SPIE, San Jose, CA, pp. 123–134.
- Druga, S., Williams, R., Breazeal, C., and Resnick, M. (2017) “Hey Google, is it OK if I eat you?” In *Proc. of the 2017 Conference on Interaction Design and Children (IDC '17)*, ACM Press, pp. 595–600.
- DSDM (2014) *The DSDM Agile Project Framework Handbook*, DSDM Consortium, Kent, UK, ISBN 978-0-9544832-9-6.
- Dumas, B., Lalanne, D., and Oviatt, S. (2009) Multimodal Interfaces: A Survey of Principles, Models and Frameworks. Human Machine Interaction Lecture Notes in *Computer Science*, 5440, 3–26.
- Dumas, J. S., and Redish, J. C. (1999) *A Practical Guide to Usability Testing* (rev. edn). Intellect, Exeter.
- Eason, K. (1987) *Information Technology and Organizational Change*. Taylor and Francis, London.
- Eason, K. (2014) Afterword: The Past, Present, and Future of Sociotechnical Systems Theory. *Applied Ergonomics*, 45, 213–220.
- Ebert, J.F., Huibers, L., Christensen, B., and Christensen M.B. (2018) Paper- or Web-Based Questionnaire Invitations as a Method for Data Collection: Cross-Sectional Comparative Study of Differences in Response Rate, Completeness of Data, and Financial Cost. *Journal of Medical Internet Research*, 20(1): e24.

- Ecker, R. (2016) Automated Detection and Guessing Without Semantics of Sender-Receiver Relations in Computer-Mediated Discourses. In *Proceedings of the 18th International Conference on Information Integration and Web-based Applications and Services (iiWAS '16)*. ACM, New York, NY, pp. 170–178.
- Educause (2016) *The 2016 Horizon Report*. Downloaded from: <https://library.educause.edu/topics/teaching-and-learning/learning-analytics>.
- Eggers, D. (2013) *The Circle*. Knopf.
- Elgon, M. (2018) The Case Against Teaching Kids to be Polite to Alexa. *Mind and Machine*. Downloaded from: <https://www.fastcompany.com/40588020/the-case-against-teaching-kids-to-be-polite-to-alexa>.
- Eliot, C., and Woolf, B. (1994) Reasoning about the User Within a Simulation-Based Real-Time Training System. In *Proceedings of 4th International Conference on User Modeling*, Mitre Corp., Bedford, MA.
- Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tang, J., and Welch, B. (1992) Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations and Remote Collaboration. In *Proceedings of CHI '92*. ACM, New York, NY, pp. 599–607.
- Empson, R. (2012) *Google Biz Chief: Over 10M Websites Now Using Google Analytics*. Retrieved from <http://techcrunch.com/2012/04/12/google-analytics-officially-at-10m/>.
- Erickson, T., and Kellogg, W. A. (2000) Social Translucence: An Approach to Designing Systems that Support Social Processes, *Transactions of Computer-Human Interaction*, 7(1), 59–83.
- Erickson, T. D. (1990) Working with Interface Metaphors. In B. Laurel (ed.) *The Art of Human-Computer Interface Design*. Addison-Wesley, Boston.
- Erickson, T. D., Smith, D. N., Kellogg, W. A., Laff, M., Richards, J. T., and Bradner, E. (1999) Socially Translucent Systems: Social Proxies, Persistent Conversation and the Design of 'Babble'. In *Proceedings of CHI '99*, pp. 72–79.
- Eysenck, M., and Brysbaert, M. (2018) *Fundamentals of Cognition*. 3rd Edition. Routledge.
- Fernaes, Y., and Tholander, J. (2006) Finding Design Qualities in a Tangible Programming Space. In *Proceedings of CHI 2006*. ACM, New York, NY, pp. 447–456.
- Fernández-Luque, F., Zapata, J., Ruiz, R., and Iborra, E. (2009) A Wireless Sensor Network for Assisted Living at Home of Elderly People, *Lecture Notes in Computer Science*, 5602. Springer-Verlag, Berlin, Heidelberg, pp. 65–74.
- Ferreira, J., Sharp, H., and Robinson, H. (2012) Agile Development and User Experience Design Integration as an On-going Achievement in Practice. In *Proceedings of Agile 2012*, Dallas, Texas.
- Ferreira, J., Sharp, H., and Robinson, H.M. (2011) User Experience Design and Agile Development: Managing Cooperation Through Articulation Work. In: *Software Practice and Experience*, 41(9), 963–974.
- Fetterman, D. M. (2010) *Ethnography: Step by Step* (3rd edn). Applied Social Research Methods Series, Vol. 17. SAGE.
- Fialho, P., and Coheur, L. (2015) ChatWoz: Chatting through a Wizard of Oz. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '15)*. ACM, New York, NY, pp. 423–424.

- Fishkin, K. P. (2004) A Taxonomy for and Analysis of Tangible Interfaces, *Personal and Ubiquitous Computing*, 8, 347–358.
- Fiske, J. (1994) Audiencing: Cultural Practice and Cultural Studies. In N. K. Denzin and Y. S. Lincoln (eds) *Handbook of Qualitative Research*. SAGE, Thousand Oaks, CA, pp. 189–198.
- Fitts, P. M. (1954) The Information Capacity of the Human Motor System in Controlling Amplitude of Movement, *Journal of Experimental Psychology*, 47, 381–391.
- Flanagan, J. C. (1954) The Critical Incident Technique, *Psychological Bulletin*, 51, 327–358.
- Fogg, B.J. (2009) A Behavior Model for Persuasive Design. In *Proceedings of the 4th International Conference on Persuasive Technology (Persuasive '09)*. ACM, New York, NY, Article 40, 7 pages.
- Folmer, E., Yuan, B., Carr, D., and Sapre, M. (2009) TextSL: A Command-Based Virtual World Interface for the Visually Impaired. In *Proceedings 11th international ACM SIGACCESS Conference on Computers and Accessibility*, pp. 59–66.
- Fontana, A., and Frey, J. H. (2005). The Interview: From Neutral Stance to Political Involvement. In N. K. Denzin and Y. S. Lincoln (eds) *The SAGE Handbook of Qualitative Research* (3rd edn), pp. 695–727. Thousand Oaks, CA: SAGE.
- Foong, E., Gergle, D., and Gerber, E.M. (2017) Novice and Expert Sensemaking of Crowdsourced Feedback. In *Proceedings of the ACM on Human-Computer Interaction*, 1, 2, Article 45.
- Forsell, C., and Johansson, J. (2010) An Heuristic Set for Evaluation in Information Visualization. In *Proceedings of the International Conference on Advanced Visual Interfaces (AVI '10)*, Giuseppe Santucci (ed.). ACM, New York, NY, pp. 199–206.
- Friedman, N., and Cabral, A. (2018) Using a Telepresence Robot to Improve Self-Efficacy of People with Developmental Disabilities. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '18)*. ACM, New York, NY. 489–491.
- Froehlich, J., Findlater, L., and Landay, J. (2010) The Design of Eco-Feedback Technology. In *Proceedings of CHI '10*, ACM, New York, NY, pp. 1999–2008.
- Furniss, D., and Blandford, A. (2006) Understanding Emergency Dispatch in Terms of Distributed Cognition: A Case Study, *Ergonomics*, 49(12/13), October, pp. 1174–1203.
- Gabrielli, S., Rogers, Y., and Scaife, M. (2000) Young Children's Spatial Representations Developed Through Exploration of a Desktop Virtual Reality Scene, *Education and Information Technologies*, 5(4), 251–262.
- Galitz, W. O. (1997) *The Essential Guide to User Interface Design*. John Wiley & Sons Inc., New York.
- Gallo, A. (2017) A Refresher on A/B Testing, *Harvard Business Review*. Downloaded from <https://hbr.org/2017/06/a-refresher-on-ab-testing>.
- Gardner, H., and Davis, K. (2014) *The App Generation: How Today's Youth Navigate Identity, intimacy, and Imagination in a Digital World*. Yale University Press.
- Garrett, J. J. (2010) *The Elements of User Experience: User-Centered Design for the Web and Beyond* (2nd edn). New Riders Press.
- Gaver, B., Dunne, T., and Pacenti, E. (1999) Cultural Probes, *ACM Interactions Magazine* January/February, 21–29.

- Gibson, J. (2014) *Introduction to Game Design, Prototyping, and Development*. Addison Wesley.
- Gigante, M. A. (1993) Virtual Reality: Enabling Technologies. In R. A. Earnshaw, M. A. Gigante and H. Jones (eds) *Virtual Reality Systems*. Academic Press, London, pp. 15–25.
- Gigerenzer, G., Todd, P., and the ABC Research Group (1999) *Simple Heuristics That Make Us Smart*. Oxford University Press, New York.
- Glaser, B. G. (1992) *Basics of Grounded Theory: Emergence vs Forcing*. Sociology Press.
- Glaser, B. G., and Strauss, A. (1967) *Discovery of Grounded Theory*. Aldine, London.
- Golsteijn, C., Gallacher, S., Koeman, L., Wall, L., Andberg, S., Rogers, Y., and Capra, L. (2015) VoxBox: a Tangible Machine That Gathers Opinions from the Public at Events. In *Proc. of TEI 2015*. ACM.
- Gooch, D., Barker, M., Hudson, L., Kelly, R., Kortuem, G., van der Linden, J., Petre, M., Brown, R., Klis-Davies, A., Forbes, H., Mackinnon, J., Macpherson, R., and Walton, C. (2018) Amplifying Quiet Voices: Challenges and Opportunities for Participatory Design at an Urban Scale. *ACM Transactions on Computer-Human Interaction*. 25, 1, 2.
- Gosper, J., Agathos, J-L., Rutter, R., and Coatta, T. (2011) Case Study: UX Design and Agile: A Natural Fit? *Communications of the ACM*, 54(1), 54–60.
- Gothelf, J and Seiden, J. (2016) *Lean UX: Designing Great Products with Agile Teams* (2nd edition). O'Reilly.
- Gottesdiener, E., and Gorman, M. (2012) *Discover to Deliver: Product Planning and Analysis*. EBG Consulting, Inc.
- Gould, J. D., and Lewis, C. H. (1985) Designing for Usability: Key Principles and What Designers Think, *Communications of the ACM*, 28(3), 300–311.
- Granollers, T. (2018) Usability Evaluation with Heuristics, Beyond Nielsen's list. In *Proceedings of the 8th International Conference on Advances in Human Computer Interaction*. pp. 60–65.
- Gray, C.M., Kou, Y., Battles, B., Hoggatt, J., and Toombs, A.L. (2018) The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 534, 14 pages.
- Greenberg, S., Carpendale, S., Marquardt, N., and Buxton, B. (2012) *Sketching User Experiences*. Morgan Kaufmann.
- Griffiths, A. (2014) How Paro the Robot Seal is Being Used to Help UK Dementia Patients. Downloaded from: <http://www.theguardian.com/society/2014/jul/08/paro-robot-seal-dementia-patients-nhs-japan>.
- Grison, E., Gyselinck, V., and Burkhardt, J-M. (2013) Using the Critical Incidents Technique to Explore Variables Related to Users' Experience of Public Transport Modes. In *Proceedings of ECCE '13 Proceedings of the 31st European Conference on Cognitive Ergonomics*, Article No. 21, ACM.
- Grosse-Hering, B., Mason, J., Aliakseyeu, D., Bakker, C., and Desmet, P. (2013) Slow Design for Meaningful Interactions. In *Proceedings of Conference on Human Factors in Computing Systems (CHI'13)*. ACM, New York, NY, pp.3431–3440.
- Grudin, J. (1989) The Case Against User Interface Consistency, *Communications of the ACM*, 32(10), 1164–1173.

- Gubbels, M., and Froehlich, J. (2014) Physically Computing Physical Computing: Creative Tools for Building with Physical Materials and Computation. *IDC '14 Extended Abstracts*.
- Guha, M. L., Druin, A., and Fails, J.A. (2013) Cooperative Inquiry Revisited: Reflections of the Past and Guidelines for the Future of Intergenerational Co-Design. *International Journal of Child-Computer Interaction*, 1(1), 14–23.
- Gui, X., Chen, Y., Caldeira, C., Xiao, D., and Chen, Y. (2017) When Fitness Meets Social Networks: Investigating Fitness Tracking and Social Practices on WeRun. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, pp. 1647–1659.
- Gun, R. J. A., and Billingham, L. M. (2016) A Comparative Study of Simulated Augmented Reality Displays for Vehicle Navigation. In *Proceedings of the 28th Australian Conference on Computer-Human Interaction (OzCHI '16)*. ACM, New York, NY, pp. 40–48.
- Gunther, V. A., Burns, D. J., and Payne, D. J. (1986) Text Editing Performance as a Function of Training with Command Terms of Differing Lengths and Frequencies, *SIGCHI Bulletin*, 18, 57–59.
- Hampton, K., and Wellman, B. (2003) Neighboring in Netville: How the Internet Supports Community and Social Capital in a Wired Suburb. *City and Community*, 2, 4, 277–311.
- Handel, D., Hochman, J., and Santo, D. (2015) Visualizing Teacher Tweets: Finding Professional Learning Networks in Topical Networks. *ASIST*, 2015, 1–3.
- Hansen, D., Shneiderman, B., Smith, M. A., and Himelboim, I. (2019) *Analyzing Social Media Networks with NodeXL: Insights from a Connected World*: (2nd ed.). Elsevier Publishers.
- Harari, G. M., Gosling, S. D., Wang, R., Chen, F., Chen, Z., Campbell, A. T. (2017) Patterns of Behavior Change in Students Over an Academic Term: A Preliminary Study of Activity and Sociability Behaviors Using Smartphone Sensing Methods, *Computers in Human Behavior*, 67, 129–138.
- Harboe, G., and Huang, E.M. (2015) Real-World Affinity Diagramming Practices: Bridging the Paper-Digital Gap. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, pp. 95–104.
- Harjuniemi E., and Häkkinen, J. (2018) Smart Handbag for Remembering Keys. In *Proceedings of the 22nd International Academic Mindtrek Conference (MindTrek 2018)*. New York, NY: ACM Press, pp. 244–247.
- Harley, A. (2018a) Visibility of System Status. Downloaded from: <https://www.nngroup.com/articles/visibility-system-status/>
- Harley, A. (2018b) UX Guidelines for Ecommerce Homepages, Category Pages, and Product Listing Pages. Downloaded from: <https://www.nngroup.com/articles/ecommerce-homepages-listing-pages/>.
- Harley, A. (2018c) UX Expert Reviews: Downloaded from: <https://www.nngroup.com/articles/ux-expert-reviews/>.
- Harman, M., Jia, Y., and Zhang, Y. (2012) App Store Mining and Analysis: MSR for App Stores. In *Proceedings of the 9th IEEE Working Conference on Mining Software Repositories (MSR 12)*, 108–111.
- Harper, R., Rodden, T., Rogers, Y., and Sellen, A. (2008) *Being Human: HCI in the Year 2020*. Microsoft (free copies from: <http://research.microsoft.com/en-us/um/cambridge/projects/hci2020/>).

- Harrison, E. (ed.) (2009) *Media Space 20 + Years of Mediated Life*. Springer.
- Hartson, H. R., and Hix, D. (1989) Toward Empirically Derived Methodologies and Tools for Human–Computer Interface Development, *International Journal of Man–Machine Studies*, 31, 477–494.
- Hassenzahl, M. (2010) *Experience Design: Technology for All the Right Reasons*. Morgan & Claypool.
- Hatch, M. (2014) *The Maker Movement Manifesto*. McGraw Hill.
- Hayashi, E., Maas, M., and Hong, J.I. (2014) Wave to Me: User Identification Using Body Lengths and Natural Gestures, In *Proceedings of CHI 14*, pp. 3453–3462.
- Hayes, G.R. (2011) The Relationship of Action Research to Human-Computer Interaction. *ACM Transactions on Human-Computer Interaction*, 18, 3, Article 15, 20 pages.
- Hazas, M., Bernheim Brush, A.J., and Scott, J. (2012) Sustainability Does not Begin with the Individual, *Interactions*, 19(5), 14–17.
- Hazlewood, W., Dalton, N. S., Rogers, Y., Marshall, P., and Hertrich, S. (2010) Bricolage and Consultation: A Case Study to Inform the Development of Large-Scale Prototypes for HCI Research. In *Proceedings of Designing Interactive Systems, DIS 2010*. ACM, New York, NY, pp. 380–388.
- Heath, C., and Luff, P. (1992) Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. In *Proceedings of CSCW '92 1(1&2)*, 69–94.
- Heath, C., Hindmarsh, J., and Luff, P. (2010) *Video in Qualitative Research*. SAGE.
- Heer, J., and Bostock, M. (2010) Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design. In *Proceedings of CHI 2010*. ACM, New York, NY, pp. 203–212.
- Hektner, J. M., Schmidt, J. A., and Csikszentmihalyi, M. (2006) *Experience Sampling Method: Measuring the Quality of Everyday Life*. SAGE.
- Hendriks-Jansen, H. (1996) *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought*. MIT Press, Cambridge, MA.
- Henkel, L. A. (2014) Point-and-Shoot Memories The Influence of Taking Photos on Memory for a Museum Tour. *Psychological science*, 25.2, 396–402.
- Henschke, M., Gedeon, T., and Jones, R. (2015) Touchless Gestural Interaction with Wizard-of-Oz: Analysing User Behaviour. In *Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction (OzCHI '15)*, ACM, New York, NY, pp. 207–211.
- Hicks, J. (2012) *The Icon Handbook*. Five Simple Steps Publishing Ltd.
- Hine, C. (2000) *Virtual Ethnography*. SAGE.
- Hodges, S., Hartmann, B., Gellersen, H., and Schmidt, A. (2014) A Revolution in the Making [Guest editors' introduction]. *IEEE Pervasive Computing*, 13(3): 18–21.
- Hodges, S., Scott, J., Sentance, S., Miller, C., Villar, N., Schwiderski-Grosche, S., Hammil, K., and Johnston, S. (2013) .NETGadgeteer: A New Platform for K–12 Computer Science Education. In *Proceedings of SIGCSE '13*, ACM, 391–396.
- Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., and Wood, K. (2006) SenseCam: A Retrospective Memory Aid. In P. Dourish and A. Friday (eds) *Ubicomp 2006*, LNCS 4206. Springer-Verlag, pp. 177–193.

- Hofte H. T., Jensen, K. L., Nurmi, P., and Froehlich, J. (2009) Mobile Living Labs 09: Methods and Tools for Evaluation in the Wild. In *In Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '09)*. ACM, New York, NY, Article 107, 2 pages.
- Hollingshead, T., and Novick, D. G. (2007) Usability Inspection Methods After 15 Years of Research and Practice. *SIGDOC 2007*, pp. 249–255.
- Hollis, V., Konrad, A., Springer, A., Antoun, M., Antoun, C., Martin, R., and Whittaker. S (2017) What Does All This Data Mean for My Future Mood? Actionable Analytics and Targeted Reflection for Emotional Well Being. *Human-Computer Interaction*, 32 (5–6), 208–267.
- Holloway, C., and Dawes, H. (2016) Disrupting the World of Disability: The Next Generation of Assistive Technologies and Rehabilitation Practices. *Healthcare Technology Letters*, 3 (4), 254–256.
- Holtzblatt, K. (2001) *Contextual Design: Experience in Real Life*. Mensch & Computer.
- Holtzblatt, K., and Beyer, H. (2017) *Contextual Design (2nd ed.) Design for life*. Morgan Kaufmann.
- Holtzblatt, K., and Jones, S. (1993) Contextual Inquiry: A Participatory Technique for Systems Design. In D. Schuler and A. Namioka (eds) *Participatory Design: Principles and practice*. Lawrence Earlbaum Associates, Hillsdale, NJ, pp. 177–210.
- Höök, K. (2018) *Designing with the Body Somaesthetic Interaction Design*. MIT
- Hornbæk, K., and Hertzum, M. (2017) Technology Acceptance and User Experience: A Review of the Experiential Component in HCI. *Transactions on Human-Computer Interaction*, 24, 5, Article 33, 30 pages.
- Hornecker, E. (2005) A Design Theme for Tangible Interaction: Embodied Facilitation. In *Proceedings of the 9th European Conference on Computer Supported Cooperative Work, ECSCW '05*, 18–22 September, Paris. Kluwer/Springer, pp. 23–43.
- Hornecker, E. Marshall, P., and Jörn Hurtienne, J. (2017) Locating Theories of Embodiment Along Three Axes: 1st–3d Person, Body-Context, Practice-Cognition. Workshop position paper for CHI 2017 workshop on Soma-Based Design Theory. Downloaded from: <http://www.ehornecker.de/Papers/SomaestheticWS-embodimentshortie.pdf>
- Horton, S. (2005) *Access by Design: A Guide to Universal Usability for Web Designers*. New Riders Press, Indianapolis, IN.
- Houben, S., Golsteijn, C., Gallacher, S., Johnson, R., Bakker, S., Marquardt, N., Capra, L., and Rogers, Y. (2016) Physikit: Data Engagement Through Physical Ambient Visualizations in the Home. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, pp. 1608–1619.
- Howells, L. (2011). *A Guide to Heuristic Website Reviews* <http://www.smashingmagazine.com/2011/12/16/a-guide-to-heuristic-website-reviews/> (accessed August, 2014).
- Hsu, Y., Dille, P., Cross, J., Dias, B., Sargent, R., and Nourbakhsh, I. (2017) Community-Empowered Air Quality Monitoring System. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, pp. 1607–1619.
- Huff, D. (1991) *How to Lie with Statistics*. Penguin.
- Hutchings, D. (2012) An Investigation of Fitts' Law in a Multiple-Display Environment. In *ACM Proceedings of CHI'12*.

- Hutchings, D., Smith, G., Meyers, B., Czerwinski, M., and Robertson, G. (2004) Display Space Usage and Window Management Operation Comparisons Between Single Monitor and Multiple Monitor Users. In *Proceedings of the Working Conference on Advanced Visual Interfaces, AVI 2004*, pp. 32–39.
- Hutchins, E. (1995) *Cognition in the Wild*. MIT Press, Cambridge, MA.
- Hutchins, E., Holan, J. D., and Norman, D. (1986) Direct manipulation interfaces. In D. Norman and S. W. Draper (eds) *User Centered System Design*. Lawrence Earlbaum Associates, Hillsdale, NJ, pp. 87–124.
- Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B. B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H., Roussel, N., and Eiderbäck, B. (2003) Technology Probes: Inspiring Design for and with Families, In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, New York, NY, pp. 17–24.
- IEEE Ethically Aligned Design (2018), Downloaded from: <https://ethicsinaction.ieee.org/>.
- Isaacs, E., Konrad, A., Walendowski, A., Lennig, T., Hollis, V., and Whittaker, S. (2013) Echoes from the Past: How Technology Mediated Reflection Improves Well-Being. In *Proceedings of CHI '13*, ACM, 1071–1080.
- Ishii, H., and Ullmer, B. (1997) Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms. In *Proceedings CHI 1997*. ACM, New York, NY, pp. 234–241.
- Ishii, H., Kobayashi, M., and Grudin, J. (1993) Integration of Interpersonal Space and Shared Work-Space: Clearboard Design and Experiments. *ACM Transactions on Information Systems*, 11(4), 349–375.
- Jadhav, D., Bhutkar, G., and Mehta, V. (2013) Usability Evaluation of Messenger Applications for Android Phones Using Cognitive Walkthrough. In *Proceedings of APCHI'2013*, pp. 9–18.
- Jaidka, S., Reeves, S., and Bowen, J. (2017) Modelling Safety-Critical Devices: Coloured Petri Nets and Z. In *Proceedings of ECIS, the ACM SIGCHI Symposium on Engineering Interactive Computing Systems*, pp. 51–56.
- Jang, J., Zhao, D., Hong, W., Park, Y., and Yong Yi, M. (2016) Uncovering the Underlying Factors of Smart TV UX over Time: A Multi-study, Mixed-method Approach, In *Proceedings of the ACM International Conference on Interactive Experiences for TV and Online Video (TVX '16)*. ACM, New York, NY, pp. 3–12.
- Jaques, N., Rudovic, O., Taylor, S., Sano, A., and Picard, R. (2017) Predicting Tomorrow's Mood, Health, and Stress Level using Personalized Multitask Learning and Domain Adaptation. In *Proceedings of Machine Learning Research*, 48, 17–33.
- Javornik, A., Freeman, R., and Moutinho, A. (2017) Retail Experience With Face Application. Free to download from https://www.amazon.co.uk/dp/B076C41L31/ref=rdr_ext_sb_ti_hist_2.
- Javornik, A., Rogers, Y., Gander, D., and Moutinho, A. (2017) MagicFace: Stepping into Character through an Augmented Reality Mirror. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI'17)*. ACM, New York, NY, pp. 4838–4849.
- Johnson, J. (2014) *Designing with the Mind in Mind: Simple Guide to Understanding User Interface Design Rules*. Morgan Kaufmann.
- Johnson, J., and Henderson, A. (2002) Conceptual Models: Begin by Designing What to Design, *Interactions* January/February, 25–32.

- Johnson, J., and Henderson, A. (2012) *Conceptual Models: Core to Good Design*. Morgan & Claypool Publishers.
- Johnson, J., and Finn, K. (2017) *Designing User Interfaces for an Aging Population: Towards Universal Design*. Morgan Kaufmann.
- Johnson, R., Van der Linden, J., and Rogers, Y. (2010). To Buzz or Not to Buzz: Improving Awareness of Posture Through Vibrotactile Feedback. In *Whole Body Interaction Workshop, CHI 2010*. ACM.
- Johnson-Laird, P. N. (1983) *Mental Models*. Cambridge University Press, Cambridge.
- Jokela, T., Ojala, J., and Olsson, T. (2015) A Diary Study on Combining Multiple Information Devices in Everyday Activities and Tasks. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, pp. 3–12.
- Jones, L. A., and Sarter, N. B. (2008) Tactile Displays: Guidance for their Design and Application, *Human Factors: The Journal of the Human Factors and Ergonomics Society* 50, 90–111.
- Joonhwan, L., Forlizzi, J., and Hudson, S. E. (2005) Studying the Effectiveness of MOVE: A Contextually Optimized In-Vehicle Navigation System. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. ACM, New York, NY, pp. 571–580.
- Jordan, B., and Henderson, A. (1995) Interaction Analysis: Foundations and Practice. *Journal of the Learning Sciences*, 4(1), 39–103.
- Joshi, S., Nistala, P.V., Jani, H., Sakhardande, P., and Dsouza, T. (2017) User-Centered Design Journey for Pattern Development. In *Proceedings of the 22nd European Conference on Pattern Languages of Programs (EuroPLoP '17)*. ACM, New York, NY, Article 23, 19 pages.
- Jupp, V. (ed.) (2006) *The SAGE Dictionary of Social Research Methods*. SAGE.
- Kahn, R., and Cannell, C. (1957) *The Dynamics of Interviewing*. John Wiley & Sons Inc., New York.
- Kahneman, D. (2011) *Thinking, Fast and Slow*. Penguin.
- Kammer, D., Schmidt, D., Keck, M., and Groh, R. (2013) Developing Mobile Interface Metaphors and Gestures. In *Proceedings of MobileHCI*, ACM Press, pp. 516–521.
- Kang, S., Clegg, T., Norooz, L., Froehlich, J., and Byrne V. (2018) Prototyping and Simulating Complex Systems with Paper Craft and Augmented Reality: An Initial Investigation. In *Proceedings of the 12th International Conference on Tangible, Embedded, and Embodied Interaction (TEI '18)*. ACM, New York, NY, pp. 320–328.
- Kaninsky, M., Gallacher, S., and Rogers, Y. (2018) Confronting People's Fears about Bats: Combining Multi-modal and Environmentally Sensed Data to Promote Curiosity and Discovery. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. ACM, New York, NY, pp. 931–943.
- Karapanos, E., Martensi, J.-B., and Hassenzahl, M. (2009) Accounting for Diversity in Subjective Judgments. In *Proceedings of CHI 2009*. ACM, New York, NY, pp. 639–648.
- Karat, C.-M. (1994) A Comparison of User Interface Evaluation Methods. In J. Nielsen and R. L. Mack (eds) *Usability Inspection Methods*. John Wiley & Sons Inc., New York.
- Kari, T., Arjoranta, J., and Salo, M. (2017) Behavior Change Types with Pokémon GO. In *Proceedings of the 12th International Conference on the Foundations of Digital Games (FDG '17)*. ACM, New York, NY, Article 33, 10 pages.

- Kawa, L. (2018) Two Major Apple Shareholders Push for Study of iPhone Addiction in Children. Available at: <https://www.bloomberg.com/news/articles/2018-01-08/jana-calpers-push-apple-to-study-iphone-addiction-in-children>.
- Kazemitabaar, M., McPeak, J., Jiao, A., He, L., Outing, T., and Froehlich, J.E. (2017) Maker-Wear: A Tangible Approach to Interactive Wearable Creation for Children. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, pp. 133–145.
- Keirnan, A., Ahmadpour, N., Pedell, S., and Mayasari, A. (2015) Lights, Camera, Action: Using Animations to Co-Evaluate User Experience Scenarios. In *Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction (OzCHI '15)*, ACM, New York, NY, pp. 492–496.
- Kelley, T. with Littman, J. (2016) *The Art of Innovation*. Profile Books, Croydon, Surrey.
- Kempton, W. (1986) Two Theories of Home Heat Control, *Cognitive Science* 10, 75–90.
- Kerr, S. J., Tan, O., Chua, J. C. (2014) Cooking Personas: Goal-Directed Design Requirements in the Kitchen, *International Journal of Human-Computer Studies*, 72, 255–274.
- Khalid, H., Shihab, E., Nagappan, M., and Hassan, A.E. (2015) What Do Mobile App Users Complain About? *IEEE Software*, May/June, 70–77.
- Kim, H., Coutrix, C., and Roudaut, A. (2018) KnobSlider: Design of a Shape-Changing UI for Parameter Control. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 339, 13 pages.
- Kim, J., and Hastak M. (2018) Social Network Analysis: Characteristics of Online Social Networks after a Disaster. *International Journal of Information Management*, 38, 86–96.
- Kim, S. (1990) Interdisciplinary Cooperation. In B. Laurel (ed.) *The Art of Human-Computer Interface Design*. Addison-Wesley, Reading, MA.
- Kim, Y. (2015). Libero: On-the-Go Crowdsourcing for Package Delivery. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15)*. ACM, New York, NY, pp. 121–126.
- Kinshumann, K., Glerum, K., Greenberg, S., Aul, G., Orgovan, V., Nichols, G., Grant, D., Loihle, G., and Hunt, G. (2011) Debugging in the (Very) Large: Ten Years of Implementation and Experience, *CACM*, 54(7), 111–116.
- Kirk, D. S., Durrant, A., Wood, G., Leong, T.W., and Wright, P. (2016) Understanding the Sociality of Experience in Mobile Music Listening with Pocketsong. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, pp. 50–61.
- Kirkham, R., Mellor, S., Green, D., Lin, J-S., Ladha, K., Ladha, C., Jackson, D., Olivier, P., Wright, P., and Ploetz, T. (2013) The Break-Time Barometer: An Exploratory System for Workplace Break-Time Social Awareness. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, 73–82.
- Kirsh, D. (2010) Thinking with External Representations, *AI & Society*. Online version: downloaded from www.springerlink.com/content/5913082573146k68/ (retrieved May 1, 2010).
- Kirsh, D. (2013) Embodied Cognition and the Magical Future of Interaction Design, *ACM Transactions on Computer-Human Interaction*, Vol. 20, No. 1, Article 3, 30 pages.
- Kjeldskov, J., and Skov, M. (2014) Was it Worth the Hassle? Ten Years of Mobile HCI Research Discussions on Lab and Field Evaluations. In *ACM Proceedings of MobileHCI*. Toronto, Canada, September 23–26.

- Klein, L. (2014) What Do We Actually Mean by ‘Sociotechnical’? On Values, Boundaries and the Problems of Language. *Applied Ergonomics*, 45, 137–142.
- Klemmer, S. R., Hartmann, B., and Takayama, L. (2006) How Bodies Matter: Five Themes for Interaction Design. In *Proceedings of the 6th Conference on Designing Interactive Systems, DIS 2006*. ACM, New York, NY, pp. 140–149.
- Klemmer, S. R., Newman, M. W., Farrell, R., Bilezikjian, M., and Landay, J. A. (2001) The Designer’s Outpost: A Tangible Interface for Collaborative Website Design. In *Symposium on User Interface Software and Technology*. ACM, New York, NY, pp. 1–10.
- Knapp, J. with Zeratsky, J., and Kowitz, B. (2016) *Sprint: How to Solve Big Problems and Test New Ideas in Just Five Days*. Bantam Press, UK.
- Knowles, B., and Hanson, V. L. (2018) The Wisdom of Older Technology (Non-Users). In *Communications of the ACM*, 61, 3, 72–77.
- Kohavi, R. (2012) Online Controlled Experiments: Introduction, Learnings, and Humbling Statistics. In *Proceedings of the 6th ACM conference on Recommender systems (RecSys ’12)*. ACM, New York, NY, pp. 1–2.
- Kohavi, R., and Longbotham, R. (2015) Unexpected Results in Online Controlled Experiments. *SIGKDD Explorations*, 12, 2, 31–35.
- Kollman, J., Sharp, H., and Blandford, A. (2009) The Importance of Identify and Vision To User Experience Designers On Agile Projects. In *Proceedings of the 2009 Agile Conference*, IEEE Computer Society, Washington DC.
- Komninos, A. (2017) How Emotions Impact Cognition. Downloaded from: <https://www.interaction-design.org/literature/article/how-emotions-impact-cognition>.
- Kozinets, V. (2010) *Netnography*. SAGE.
- Kraut, R., Fish, R., Root, R., and Chalfonte, B. (1990) Informal Communications in Organizations: Form, Function and Technology. In S. Oskamp and S. Krug (eds) *Don’t Make Me Think*. New Riders/Peachpit.
- Krippendorff, K. (2013) *Content Analysis: An Introduction to Its Methodology* (3rd edn). SAGE Publications.
- Krug, S. (2014) *Don’t Make Me Think, Revisited: A Common Sense Approach to Web Usability* (3rd edn). Pearson.
- Kuhn, T. S. (1972/1962) *The Structure of Scientific Revolutions* (2nd edn). University of Chicago Press, Chicago.
- Kushniruk A., Monkman H., Borycki E., and Kannry J. (2015) User-Centered Design and Evaluation of Clinical Information Systems: A Usability Engineering Perspective. In: Patel V., Kannampallil T., and Kaufman D. (eds.) *Cognitive Informatics for Biomedicine. Health Informatics*. Springer.
- Lakoff, G., and Johnson, M. (1980) *Metaphors We Live By*. University of Chicago Press, Chicago.
- Lane, N. D., and Georgiev, P. (2015). Can Deep Learning Revolutionize Mobile Sensing? In *Proceedings of the 16th International Workshop on Mobile Computing Systems and Application*. ACM, New York, NY, pp. 117–122.
- Law, E. L., Roto, V., Hassenzahl, M., Vermeeren, A. P., and Kort, J. (2009) Understanding, Scoping and Defining User Experience: A Survey Approach. In *Proceedings of the 27th*

- International Conference on Human Factors in Computing Systems, CHI 2009*. ACM, New York, NY, pp. 719–728.
- Lazar, J., Feng, H. J., and Hochheiser, H. (2017) *Research Methods in Human-Computer Interaction*. (2nd ed.). Cambridge, MA: Elsevier/Morgan Kaufmann Publishers.
- Lazar, J., Goldstein, D., and Taylor, A. (2015). *Ensuring Digital Accessibility Through Process and Policy*. Waltham, MA: Elsevier/Morgan Kaufmann Publishers.
- Lazar, J., Jaeger, P., Adams, A., Angelozzi, A., Manohar, J., Marciniak, J., Murphy, J. Norasteh, P., Olsen, C., Ponerer, E., Scott, T., Vaidya, N., and Walsh, J. (2010) Up in the Air: Are Airlines Following the New DOT Rules on Equal Pricing for People with Disabilities When Websites are Inaccessible? *Government Information Quarterly*, 27(4), 329–336.
- Lechelt, Z., Rogers, Y., Yuill, N., Nagl, L., Ragone, G., and Marquardt, N. (2018) Inclusive Computing in Special Needs Classrooms: Designing for All. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 517, 12 pages.
- Ledgard, H., Singer, A., and Whiteside, J. (1981) Directions in Human Factors for Interactive Systems. In G. Goos and J. Hartmanis (eds) *Lecture Notes in Computer Science*, 103. Springer-Verlag, Berlin.
- Ley, B., Ogonowski, C., Jan Hess, M. M., Race, N., Randall, D., Rouncefield, M., and Wulf, V. (2015) At Home with Users: A Comparative View of Living Labs. *Interacting with Computers*, Volume 27, Issue 1, pp. 21–35
- Li, J., Cho, I., and Wartell I. (2018) Evaluation of Cursor Offset on 3D Selection in VR. In *Proceedings of the Symposium on Spatial User Interaction (SUV'18)*, pp.120–129.
- Lim, B.Y., Dey, A.K., and Avrahami, D. (2009) Why and Why Not Explanations Improve the Intelligibility of Context-Aware Intelligent Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, pp. 2119–2128.
- Lim, S. L., and Finkelstein, A. (2012) StakeRare: Using Social Networks and Collaborative Filtering for Large-Scale Requirements Elicitation, *IEEE Transactions on Software Engineering*, 38(2), 707–735.
- Lim, Y.-K., Stolterman, E., and Tenenbunrg, J. (2008) The Anatomy of Prototypes: Prototypes as Filters, Prototypes as Manifestations of Design Ideas, *ACM Transactions on Computer-Human Interaction*, 15(2).
- Lin, J., Newman, M.W., Hong, J.I., and Landay, J.A. (2000) DENIM: Finding a Tighter Fit Between Tools and Practice for Website Design. In *Proceedings of CHI '00*, 510–517.
- Lingel, J. (2012) Ethics and Dilemmas of Online Ethnography. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems (CHI EA '12)*. ACM, New York, NY, pp. 41–50.
- Liu, X., Feng, X., Pan, S., Peng, J., and Zhao, X. (2018) Skeleton Tracking Based on Kinect Camera and the Application in Virtual Reality System. In *Proceedings of the 4th International Conference on Virtual Reality*, pp. 21–25.
- Liu, Y., and R  ih  , K.-J. (2010) Predicting Chinese Text Entry Speeds on Mobile Phones. In *Proceedings of CHI 2010: HCI in China*, April 10–15, Atlanta, GA, pp. 2183–2192.
- Loma, N. (2018) WTF is Dark Pattern Design? Downloaded from: <https://techcrunch.com/2018/07/01/wtf-is-dark-pattern-design/>.

- Loranger, H., and Laubheimer, P. (2017) The State of UX Agile Development, downloaded from <https://www.nngroup.com/articles/state-ux-agile-development/>.
- Lottridge, D.M., Rosakranse, C., Oh, C.S., Westwood, S.J., Baldoni, K.A., Mann, A.S., and Nass, C.I. (2015) The Effects of Chronic Multitasking on Analytical Writing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, pp. 2967–2970.
- Lotz, N., Sharp, H., Woodroffe, M., Blyth., Rajah, D., and Ranganai, T. (2014) Framing Behaviours in Novice Interaction Designers, In *Proceedings of DRS 2014 Design's Big Debates*. pp. 1178–1190.
- Luce, K. H., Winzelberg, A. J., Das, S., Osborn, M. I., Bryson, S. W., and Taylor, C. B. (2003) Reliability of Self-Report: Paper Versus Online Administration, *Computers in Human Behavior* (accessed online January 20, 2005).
- Lucero, A. (2015) Using Affinity Diagrams to Evaluate Interactive Prototypes. In *Proceedings of INTERACT 2015, Part II, LNCS 9297*, pp. 231–248.
- Ludwig, T., Korthaus, C., and Pipek, V. (2016) Situated and Ubiquitous Crowdsourcing with Volunteers During Disasters. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, pp. 1441–1447.
- Lueg, C., Banks, B., Michalek, J., Dimsey, J., and Oswin, D. (2019) Close Encounters of the 5th Kind: Recognizing System-Initiated Engagement as Interaction Type. *JASIST*. Wiley Online.
- Lugmayr, A., Grenfeld, A., and Zhang, D.J. (2017) Selected Advanced Data Visualizations: The UX-Machine, Cultural Visualisation, Cognitive Big Data, and Communication of Health and Wellness Data. In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW '17 Companion)*, pp. 247–251.
- Mackay, W., and Fayard, A.-L. (1997) HCI, Natural Science and Design: A Framework for Triangulation Across Disciplines. In *Proceedings of the 2nd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS'97)* pp. 223–234.
- MacKenzie, I. S. (1992) Fitts' Law as a Research and Design Tool in Human–Computer Interaction, *Human–Computer Interaction* 7, 91–139.
- MacKenzie, I. S. (1995). Movement Time Prediction in Human-Computer Interfaces. In R. M. Baecker, W. A. S. Buxton, J. Grudin, and S. Greenberg (eds.), *Readings in Human-Computer Interaction* (2nd edn) (pp. 483–493). Los Altos, CA: Kaufmann
- Maglio, P. P., Matlock, T., Raphaely, D., Chernicky, B., and Kirsh, D. (1999) Interactive Skill in Scrabble. In *Proceedings of Twenty-first Annual Conference of the Cognitive Science Society*. Lawrence Earlbaum Associates, Mahwah, NJ.
- Maguire, M. (2014) Socio-Technical Systems and Interaction Design—21st Century Relevance. *Applied Ergonomics*, 45, 162–170.
- Maher, M. L., Preece, J., Yeh, T., Boston, C., Grace, K., Pasupuleti, A., and Stangl, A. (2014) NatureNet: A Model for Crowdsourcing the Design of Citizen Science Systems. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW Companion '14)* pp. 201–204. New York: ACM.
- Mahyar, N., James, M. R., Ng, M. M., Wu, R. A., and Dow, S. P. (2018) CommunityCrit: Inviting the Public to Improve and Evaluate Urban Design Ideas through Micro-Activities.

- In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 195, 14 pages.
- Maiden, N. A. M., Ncube, C., and Robertson, S. (2007) Can Requirements Be Creative? Experiences with an Enhanced Air Space Management System. In *Proceedings of ICSE '07*.
- Malamed, C. (2009) *Visual Language for Designers: Principles for Creating Graphics that People Understand*. Rockport Publishers.
- Mancini, C., Rogers, Y., Bandara, A.K., Coe, T., Joinson, A.N., Jedrzejczyk, L. Price, B. A. Thomas, K., and Nuseibeh, B. (2010) ContraVision: Exploring Users' Reactions to Futuristic Technology. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, pp.153–162.
- Mancini, C., Thomas, K., Rogers, Y., Price, B. A., Jedrzejczyk, L., Bandara, A. K., Joinson, A. N., and Nuseibeh, B. (2009) From Spaces to Places: Emerging Contexts in Mobile Privacy, *UbiComp 2009*, September 30–October 3.
- Mankoff, J., Fait, H., and Tran, T. (2005) Is Your Web Page Accessible? a Comparative Study of Methods for Assessing Web Page Accessibility for the Blind. In *Proceedings of CHI 2003*. ACM, New York, NY, pp. 41–50.
- Mankoff, J., Kravets, R., and Blevins, E. (2008) Some Computer Science Issues in Creating a Sustainable World, *Computer*, 41(8), 102–105.
- Mann, S. (1997) An Historical Account of the 'WearComp' and 'WearCam' Inventions Developed for Applications in Personal Imaging. In *The First International Symposium on Wearable Computers: Digest of Papers*. IEEE Computer Society, pp. 66–73.
- Manuel, D., Moore, D., and Charissis, V. (2012) An Investigation into Immersion in Games Through Motion Control and Stereo Audio Reproduction, September 26–28 AM '12: In *Proceedings of the 7th Audio Mostly Conference: A Conference on Interaction with Sound* pp. 124–129.
- Marquardt, N., Brudy, F., Liu, C., Bengler, B., and Holz, C. (2018) SurfaceConstellations: A Modular Hardware Platform for Ad-Hoc Reconfigurable Cross-Device Workspaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 354, 14 pages.
- Marsden, G., Maunder, A., and Parker, M. (2008) People Are People, but Technology is not Technology, *Philosophical Transactions of the Royal Society* 366, 3795–3804.
- Marshall, P., and Hornecker, E. (2013) Theories of Embodiment in HCI. In *The SAGE Handbook of Digital Technology Research*. 144–158.
- Marshall, P., Price, S., and Rogers, Y. (2003) Conceptualizing Tangibles to Support Learning. In *Proceedings of Interaction Design and Children, IDC 2003*. ACM, New York, p. 101–109.
- Martin, A., Biddle, R., and Noble, J. (2009) XP Customer Practice: A Grounded Theory. In *Proceedings of the 2009 Agile Conference*, IEEE Computer Society, Washington DC.
- Mason, B. (2017) Virtual Reality has a Motion Sickness Problem. Downloaded from <https://www.sciencenews.org/article/virtual-reality-has-motion-sickness-problem>
- Maurya, A. (2018) IEEE Big Data 2017 Panel Discussion on Bias and Transparency. *AI Matters*, 4 (2).
- McCarthy, J., and Wright, P. (2004) *Technology as Experience*. MIT Press, Cambridge, MA.
- McCullough, M. (2004) *Digital Ground: Architecture, Pervasive Computing and Environmental Knowing*. MIT Press, Cambridge, MA.

- McInerney, P. (2017) UX in Agile projects: Taking Stock After 12 Years, *Interactions*, March–April, 58–61.
- Mekler, E.D., Tuch, A.N., Martig, A.L., and Opwis, K. (2014) A Diary Study Exploring Game Completion and Player Experience. In *Proceedings of the First ACM SIGGHI Annual Symposium on Computer-Human Interaction In Play (CHI PLAY '14)*. ACM, New York, NY, pp. 433–434.
- Mifsud, J. (2011) 12 Effective Guidelines for Breadcrumb Usability and SEO. Downloaded from <https://usabilitygeek.com/12-effective-guidelines-for-breadcrumb-usability-and-seo/>.
- Miller, G. (1956) The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information, *Psychological Review* **63**, 81–97.
- Miller, L. (2006) Interaction Designers and Agile Development: A Partnership. In *Proceedings of UPA 2006*. Denver/Broomfield: Usability Professionals' Association.
- Miller, L. H., and Johnson, J. (1996) The Xerox Star: An Influential User Interface Design. In M. Rudisill, C. Lewis, P. G. Polson and T. D. McKay (eds) *Human-Computer Interface Design*. Morgan Kaufmann, San Francisco.
- Mittelstadt, S., Behrisch, M., Weber, S., Schreck, T., Stoffel, A., Pompl, R., Keim, D., Last, H., and Zhang, L. (2012) Visual Analytics for The Big Data Era—A Comparative Review of State-of-The-Art Commercial Systems. In *Proceedings of the 2012 IEEE Conference on Visual Analytics Science and Technology (VAST) (VAST '12)*. IEEE Computer Society, Washington, DC, pp. 173–182.
- Miyake, N. (1986) Constructive Interaction and the Iterative Process of Understanding, *Cognitive Science*, **10**(2) pp. 151–177.
- Molich, R., Laurel, B., Snyder, C., Quesenbery, W., and Wilson, C.E. (2001) Ethics in HCI, In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, NY, pp. 217–218.
- Morrison, C., Villar, N., Thieme, A., Ashktorab, Z., Taysom, E., Salandin, O., Cletheroe, D., Saul, G., Blackwell, A.F., Edge, D., Grayson, M., and Zhang, H. (2018) Torino: A Tangible Programming Language Inclusive of Children with Visual Disabilities. *Human-Computer Interaction*, 1–49.
- Morville, P. (2005) *Ambient Findability*. O'Reilly Media Inc.
- Müller, J., Oulasvirta, A., and Murray-Smith, R. (2017) Control Theoretic Models of Pointing. *ACM Transactions on Computer-Human Interaction*, **24**, 4, Article 27, 36 pages.
- Müller-Tomfelde, C. (ed.) (2010) *Tabletops: Horizontal Interactive Displays*. Springer.
- Mullet, K., and Sano, D. (1995) *Designing Visual Interfaces*. Prentice Hall, Mountain View, CA.
- Mumford, E. (2006) The Story of Socio-technical Design: Reflections in its Successes, Failures, and Potential, *Information Systems Journal*, **16**, 317–342.
- Muniz, F. (2016) An Introduction to Heuristic Evaluation. Downloaded from: <https://usabilitygeek.com/heuristic-evaluation-introduction/>.
- Nario-Redmond, M. R., Gospodinov, D., and Cobb, A. (2017) Crip for a Day: The Unintended Negative Consequences of Disability Simulations. *Rehabilitation Psychology*, **62**(3), 324–333.
- Ncube, C., Oberndorf, P., and Kark, A. W. (2008) Opportunistic Software Development: Making Systems from What's Available, *IEEE Software*, **25**(6), 38–41.
- Neil, T. (2014) *Mobile Design Pattern Gallery* (2nd edn). O'Reilly.

- Neustaedter, C., Venolia, G., Procyk, J., and Hawkins, D. (2016) To Beam or Not to Beam: A Study of Remote Telepresence Attendance at an Academic Conference. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '16)*. ACM, New York. pp. 418–431.
- Nevo, D., and Wade, M. R. (2007) How to avoid Disappointment by Design, *Communications of the ACM*, 50(4), 43–48.
- Nielsen (2014) www.useit.com.
- Nielsen, J. (1993) *Usability Engineering*. Morgan Kaufmann, San Francisco.
- Nielsen, J. (1994a) Heuristic Evaluation. In J. Nielsen and R. L. Mack (eds) *Usability Inspection Methods*. John Wiley & Sons Inc., New York.
- Nielsen, J. (1994b) Enhancing the Explanatory Power of Usability Heuristics. In *Proceedings of CHI '94*. ACM, New York, NY, pp. 152–158.
- Nielsen, J. (1999) *Designing Web Usability: The Practice of Simplicity*. New Riders Publishing Thousand Oaks, CA.
- Nielsen, J. (2000) *Designing Web Usability*. New Riders Press, Indianapolis, IN.
- Nielsen, J., and Li, A. (2017) Mega Menu Work Well for Site Navigation. Downloaded from <https://www.nngroup.com/articles/mega-menus-work-well/>.
- Nielsen, J., and Loranger, H. (2006) *Prioritizing Web Usability*. New Riders Press.
- Nielsen, J., and Mack, R.L. (eds) (1994) *Usability Inspection Methods*. John Wiley & Sons Inc., New York.
- Nielsen, J., and Mohlich, R. (1990) Heuristic Evaluation of User Interfaces. In *Proceedings of CHI '90*. ACM, New York.
- Nielsen, J., and Norman, D. (2014) *The Definition of User Experience*, www.nngroup.com/articles/definition-user-experience/ (accessed July 2, 2014).
- Nielsen, J., and Tahir, M. (2002) Homepage Usability: 50 Websites Deconstructed. New Riders Press.
- Norman, D. (1983) Some Observations on Mental Models. In D. Gentner and A. L. Stevens (eds) *Mental Models*. Lawrence Earlbaum Associates, Hillsdale, NJ.
- Norman, D. (1986) Cognitive Engineering. In D. Norman and S. W. Draper (eds) *User Centered System Design*. Lawrence Earlbaum Associates, Hillsdale, NJ, pp. 31–62.
- Norman, D. (1988) *The Design of Everyday Things*. Basic Books, New York.
- Norman, D. (1993) *Things That Make Us Smart*. Addison-Wesley, Reading, MA.
- Norman, D. (1999) Affordances, Conventions and Design, *ACM Interactions Magazine*, May/June, 38–42.
- Norman, D. (2004) Beauty, Goodness, and Usability/Change Blindness. *Human-Computer Interaction*, 19(4), 311–318.
- Norman, D. (2005) *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books, New York.
- Norman, D. (2006) Why Doing User Observations First is Wrong, *Interactions*, July/Aug, 50.
- Norman, D. (2010) Natural Interfaces are Not Natural, *Interactions*, May/June, 6–10.
- Norman, D. (2013) *The Design of Everyday Things*. The MIT Press, Cambridge, Massachusetts.

- North, S. (2017) Hey, Where's My Hay?: Design Fictions in Horse-Computer Interaction. In *Proceedings of the Fourth International Conference on Animal-Computer Interaction (ACI2017)*. ACM, New York, NY, Article 17, 5 pages.
- Nudelman, G. (2013) *Android Design Patterns*. John Wiley.
- O'Kaine, A.A., Rogers, Y., and Blandford, A.E. (2015) Concealing or Revealing Mobile Medical Devices? Designing for Onstage and Offstage Presentation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, pp. 1689–1698.
- O'Connell, B., Whittaker, S., and Wilbur, S. (1993) Conversations Over Video Conferences: An Evaluation of the Spoken Aspects of Video-Mediated Communication, *Human-Computer Interaction* 8, 389–428.
- Ofcom Report (2018) A Decade of Digital Dependency. Downloaded from <https://www.ofcom.org.uk/about-ofcom/latest/features-and-news/decade-of-digital-dependency>.
- O'Hara, K., Gonzalez, G., Sellen, A., Penney, G., Varnavas, A., Mentis, H., Criminisi, A., Corish, R., Rouncefield, M., Dastur, N., and Carrell, T. (2013) Touchless Interaction in Surgery, *Communications of the ACM*, 57(1)70–77.
- Oliveira, N., Jun, E., and Reinecke, K. (2017) Citizen Science Opportunities in Volunteer-Based Online Experiments. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, pp. 6800–6812.
- Oliver, J.L., Brereton, M., Watson, D.M., and Roe, P. (2018) Visualisations Elicit Knowledge to Refine Citizen Science Technology Design: Spectrograms Resonate with Birders. In *Proceedings of the 30th Australian Conference on Computer-Human Interaction (OzCHI '18)*. ACM, New York, NY, pp. 133–144.
- Ophir, E., Nass, C. I., and Wagner, A. D. (2009) Cognitive Control in Media Multitaskers, *Proceedings of the National Academy of Sciences USA* 106:15583–15587.
- Oppenheim, A.N. (2000) *Questionnaire Design, Interviewing and Attitude Measurement*. 2nd edition. Pinter Publishers.
- Ortony, A., Norman, D. A., and Revelle, W. (2005) Affect and proto-affect in effective functioning. In J. M. Fellous and M. A. Arbib (eds) *Who Needs Emotions? The Brain Meets the Machine*. New York: Oxford University Press, pp. 173–202.
- Oviatt, S., Cohen, A., and Weibel, N. (2013) Multimodal Learning Analytics: Description of Math Data Corpus of ICMI Grand Challenge Workshop. *ICMI '13: In Proceedings of the 15th ACM International Conference on Multimodal Interaction*.
- Oviatt, S., Schuller, B., Cohen, P.R., Sonntag, D., Potamianos, G., and Krüger, A. (eds.). (2017) *The Handbook of Multimodal-Multisensor Interfaces: Foundations, User Modeling, and Common Modality Combinations—Volume 1*. Association for Computing Machinery and Morgan & Claypool, New York, NY.
- Park, S.Y., and Chen, Y. (2015) Individual and Social Recognition: Challenges and Opportunities in Migraine Management. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '15)*. ACM, New York, NY, pp. 1540–1551.
- Parker, C., Fredericks, J., Tomitsch, M., and Yoo, S. (2017) Towards Adaptive Height-Aware Public Interactive Displays. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*, ACM, New York, NY, pp. 257–260.

- Paterson, B., Winschiers-Theophilus, H., Dunne, T. T., Schinzel, B., and Underhill, L. G. (2011) Interpretation of a Cross-Cultural Usability Evaluation—A Case Study Based on a Hypermedia System for Rare Species Management in Namibia. *Interacting with Computers*, 23, 3, 239–246.
- Pearl, C. (2016) *Designing Voice User Interfaces*. O'Reilly.
- Peatt, K. (2014) Off the Beaten Canvas: Exploring the Potential of the Off-Canvas Pattern downloaded from <http://www.smashingmagazine.com/2014/02/24/off-the-beaten-canvas-exploring-the-potential-of-the-off-canvas-pattern/> Sept 2014.
- Pêcher, C., Lemercier, C., and Cellier, J.-M. (2009) Emotions Drive Attention: Effects on Driver's Behavior. *Safety Science*, 47, 1254–1259.
- Pêcher, C., Lemercier, C., and Cellier, J.-M. (2011) The Influence of Emotions on Driving Behavior. In *Traffic Psychology: An International Perspective*. Chapter IX. Publisher: Nova Science Publishers, Editors: Dwight Hennessy.
- Perterer, N., Sundström, P., Meschtscherjakov, A., Wilfinger, D., and Tscheligi, M. (2013) Come Drive with Me: An Ethnographic Study of Driver-Passenger Pairs to Inform Future In-Car Assistance. In *Proceedings of the Conference on Computer supported cooperative work (CSCW'13)*. ACM, New York, NY, pp. 1539–1548.
- Petrie, H., Hamilton, F., King, N., and Pavan, P. (2006) Remote Usability Evaluations with Disabled People. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, ACM, New York, NY, pp. 1133–1141.
- Picard, R. W. (1998) *Affective Computing*. MIT Press, Cambridge, MA.
- Pinelle, D., Wong, N., and Stach, T. (2008) Heuristic Evaluation for Games: Usability Principles for Video Games. In *Proceedings of SIGCHI 2008*, Florence, Italy, pp. 1453–1462.
- Pinelle, D., Wong, N., Stach, T., and Gutwin, C. (2009) Usability Heuristics for Networked Multiplayer Games. In *ACM Proceedings of GROUP'09*.
- Porcheron, M., Fischer J.E., Reeves, S., and Sharples, S. (2018) Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 640, 12 pages.
- Porter, C., Letier, E., and Sasse, M.A. (2014) Building a National E-Service Using Sentire: Experience Report on the Use of Sentire: A Volere-Based Requirements Framework Driven by Calibrated Personas and Simulated User Feedback. In *Proceedings of RE'2014*, 374–383.
- Poulter, N. (2013) 6 Google Analytics Custom Dashboards to Save You Time NOW! Retrieved from <http://www.stateofdigital.com/google-analytics-dashboards/>.
- Preece, J. (2016) Citizen Science: New Research Challenges in HCI. *International Journal of Human-Computer Interaction*, 32, 8, 585–612.
- Preece, J. (2017) How Two Billion Smartphone Users Can Save Species, *Interactions*. Vol. XXIV.2, 27–33.
- Preece, J., and Shneiderman, B. (2009) The Reader to Leader Framework: Motivating Technology-Mediated Social Participation, *AIS Transactions on Human-Computer Interaction*, 1(1), 13–32.
- Preece, J., Pauw, D., and Clegg, T. (2018) Interaction Design of Community-Driven Environmental Projects (CDEPs): A Case Study from the Anacostia Watershed. *Proceedings of the National Academy of Sciences*, USA.

- Pressman, R.S., and Maxim, B.R. (2014) *Software Engineering: A Practitioner's Approach (Int'l Ed)*. McGraw-Hill Education.
- Price, B., Kelly, R., Mehta, V., McCormick, C., Ahmed, H., and Pearce, O. (2018) Feel My Pain; Design and Evaluation of Painpad, a Tangible Device for Supporting Inpatient Self-Logging of Pain. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI'18)*. ACM, New York, NY, Paper 169, 13 pages.
- Primak, R. (2014) *Walden Warming: Climate Change Comes to Thoreau's Woods*. University of Chicago Press.
- Prohaska, T.R., Anderson, L.A., and Binstock, R.H. (2012) *Public Health for an Aging Society*. JHU Press. 249–252.
- Purkiss, B., and Khaliq, I. (2015) A Study of Interaction in Idle Games and Perceptions on the Definition of a Game. In *Proceedings of IEEE Games Entertainment Media Conference (GEM'2015)*, 1–6.
- Putnam, C., Bungum, M., Spinner, D., Parelkar, A.N., Vipparti, S., and Cass, P. (2018) How User Experience Is Practiced: Two Case Studies from the Field. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, Paper LBW092, 6 pages.
- Putnam, R. D. (2000) *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.
- Rader, E., Cotter, K., and Cho, J. (2018) Explanations as Mechanisms for Supporting Algorithmic Transparency. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 103, 13 pages.
- Rae, I., Mutlu, B., Olson, G.M., Olson, J.S., Takayama, L., and Venolia, G. (2015) Everyday Telepresence: Emerging Practices and Future Research Directions. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15)*. ACM, New York. pp. 2409–2412.
- Rajanna, V., and Hammond, T. (2018) A Fitts' law evaluation of gaze input on large displays compared to touch and mouse inputs. In *Proceedings of the Workshop on Communication by Gaze Interaction (COGAIN '18)*. ACM, New York, NY, Article 8, 5 pages.
- Rajkomar, A., and Blandford, A. (2012) Understanding Infusion Administration in the ICU Through Distributed Cognition. *Journal of Biomedical Informatics*, 45(3), 580–590.
- Rajkomar, A., Mayer, A., and Blandford, A. (2015) Understanding Safety-Critical Interactions with A Home Medical Device Through Distributed Cognition. *Journal of Biomedical Informatics*, 56, 179–194.
- Ramcharitar, A., and Teather, R. J. (2017) A Fitts' Law Evaluation of Video Game Controllers: Thumbstick, Touchpad and Gyrosensor. In *Proceedings of the CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, NY, pp. 2860–2866.
- Raptis, D., Jensen, R. H., Kjeldskov, J., and Skov, M.B. (2017) Aesthetic, Functional and Conceptual Provocation in Research Through Design. In *Proceedings of DIS'2017*, ACM, New York, NY, 29–41.
- Raskin, J. (2000) *The Humane Interface*. Addison-Wesley, Harlow, Essex.
- Ratcliffe, L., and M. McNeill (2012) *Agile Experience Design*. New Riders.
- Rau, P.P., Plocher, T., and Choong, Y. (2013) *Cross-Cultural Design for IT Products and Services*. CRC Press.

- Read, J., Macfarlane, S., and Casey, C. (2002) Endurability, Engagement and Expectations: Measuring Children's Fun. In *Proceedings of Interaction Design and Children 2002*, Eindhoven, Amsterdam. ACM, New York, NY, pp. 189–198.
- Redish, G. (2012) *Letting Go of the Words: Writing Web Content That Works* (2nd edn). Morgan Kaufmann.
- Reeves, B., and Nass, C. (1996) *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, Cambridge.
- Reeves, L., Lai, J. C., Larson, J. A., Oviatt, S. L., Balaji, T. S., Buisine, S., Collings, P., Cohen, P. R., Kraal, B., Martin, J.-C., McTear, M. F., Raman, T. V., Stanney, K. M., Su, H., and Wang, Q. Y. (2004) Guidelines for Multimodal User Interface Design, *Communications of the ACM*, 47(1), 57–59.
- Reynaga, G., Chiasson, S., and van Oorschot, P.C. (2015) Heuristics for the Evaluation of Captchas on Smartphones. In *Proceedings of the 2015 British HCI Conference (British HCI '15)*. ACM, New York, NY, 126–135.
- Richard, P., Burkhardt, J.-M., and Lubart, T. (2014). Users' Participation to Creative Design of New Solutions for Mobility: An Exploratory Study. In *Proceedings of the 2014 European Conference on Cognitive Ergonomics (ECCE '14)*. ACM, New York, NY, Article 21, 7 pages.
- Richards, M., and Woodthorpe J. (2009) Introducing TU100 “My Digital Life”: Ubiquitous Computing in a Distance Learning Environment. In *Proceedings of UbiComp 2009*, Walt Disney Beach Club Resort, Orlando, FL.
- Rideout, V.J., Foehr, U.G., and Roberts, D.F. (2010) *Generation M2: Media in the Lives of 8- to 18-year-Olds*. Menlo Park, CA: Henry J Kaiser Family Foundation.
- Ries, E. (2011) *The Lean Startup: How Constant Innovation Creates Radically Successful Businesses*. Portfolio Penguin.
- Righia, V., Sayagob S., and Blat, J. (2017) When We Talk about Older People in HCI, Who Are We Talking About? Towards A ‘Turn To Community’ In The Design of Technologies For a Growing Aging Population. *International Journal of Human-Computer Studies*, 108, Issue C, 15–31.
- Rizvi, S. A., Tuson, E., Desrochers, B., and Magee, J. (2018) Simulation of Motor Impairment in Head-Controlled Pointer Fitts' Law Task. *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '18)*. ACM, New York, NY, pp. 376–378.
- Robertson, S., and Robertson, J. (2013) *Mastering the Requirements Process* (3rd edn). Pearson Education, New Jersey.
- Robinson, R., Rubin, Z., Márquez Segura, E., and Isbister, K. (2017) All the Feels: Designing A Tool that Reveals Streamers' Biometrics to Spectators. In *Proceedings of the 12th International Conference on the Foundations of Digital Games (FDG '17)*. ACM, New York, NY, Article 36, 6 pages.
- Robson, C., and McCartan, K. (2016) *Real World Research*. John Wiley & Sons.
- Rodden, K., Hutchinson, H., and Fu, X. (2010) Measuring the User Experience on a Large Scale: User-Centered Metrics for Web Applications, In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, pp. 2395–2398.

- Rogers, Y. (1989) Icons at The Interface: Their Usefulness, *Interacting with Computers*, 1(1), 105–117.
- Rogers, Y. (2006) Moving on from Weiser’s Vision of Calm Computing: Engaging UbiComp Experiences. In *Proceedings of UbiComp 2006*, LNCS 4206, Springer-Verlag, Berlin, Heidelberg, pp. 404–421.
- Rogers, Y. (2011) Interaction Design Gone Wild: Striving for Wild Theory, *Interactions*, 18(4): 58–62.
- Rogers, Y. (2012) *HCI Theory: Classical, Modern and Contemporary*. Morgan & Claypool.
- Rogers, Y., and Aldrich, F. (1996) In Search of Clickable Dons: Learning About HCI Through Interacting with Norman’s CD-ROM, *SIGCHI Bulletin*, 28(3).
- Rogers, Y., and Lindley, S. (2004) Collaborating Around Vertical and Horizontal Displays: Which Way is Best? *Interacting With Computers* 16, N33–N52.
- Rogers, Y., and Marsden, G. (2013) Does He Take Sugar? Moving Beyond the Rhetoric of Compassion, *Interactions* XX.4 July–August 2013.
- Rogers, Y., and Marshall, P. (2017) *Research in the Wild*. Morgan & Claypool.
- Rogers, Y., Hazlewood, W., Marshall, P., Dalton, N. S., and Hertrich, S. (2010a) Ambient Influence: Can Twinkly Lights Lure and Abstract Representations Trigger Behavioral Change? In *Proceedings of Ubicomp 2010*, pp. 261–270.
- Rogers, Y., Lim, Y., and Hazlewood, W. (2006) Extending Tabletops to Support Flexible Collaborative Interactions. In *Proceedings of Tabletop 2006*, Adelaide, Australia, January 5–7. IEEE, pp. 71–78.
- Rogers, Y., Lim, Y., Hazlewood, W., and Marshall, P. (2009) Equal Opportunities: Do Shareable Interfaces Promote More Group Participation than Single Users Displays? *Human-Computer Interaction*, 24(2), 79–116.
- Rogers, Y., Paay, J., Brereton, M., Vaisutis, K., Marsden, G., and Vetere, F. (2014) Never Too Old: Engaging Retired People Inventing the Future with MaKey MaKey. In *Proceedings of CHI 2014*, ACM, 2675–2684.
- Rogers, Y., Payne, S., and Todd, P. (2010b) Projecting Instant Information in Situ: Can it Help us Make More Informed Decisions? In *Ubiprojection 2010: Workshop Proceedings, Pervasive 2010*.
- Rogers, Y., Price, S., Randell, C., Fraser, D.S., Weal, M., and Fitzpatrick, G. (2005) Ubi-Learning Integrates Indoor and Outdoor Experiences, *Communications of the ACM*, 48(1), 55–59.
- Rogers, Y., Yuill, N., and Marshall, P. (2013) Contrasting Lab-Based and in-the-wild Studies for Evaluating Multi-User Technologies. In Price, S., Jewitt, C., and Brown, B. (eds.) *SAGE Handbook of Technology Research*. 359–173.
- Rønby-Pedersen, E., McCall, K., Moran, T. P., and Halasz, F. G. (1993) Tivoli: An Electronic Whiteboard for Informal Workgroup Meetings. In *Proceedings of CHI ’93*. ACM, New York, NY, pp. 391–398.
- Rooksby, J., Rost, M., Morrison, A., and Chalmers, M. (2014) Personal Tracking as Lived Informatics. In *Proceedings of CHI’14*, ACM, 1163–1172.
- Rose, D. (2018) Why Gesture is the Next Big Thing in Design. Downloaded from: <https://www.ideo.com/blog/why-gesture-is-the-next-big-thing-in-design>.

- Roth, I. (1986) An Introduction to Object Perception. In I. Roth and J. B. Frisby (eds) *Perception and Representation: A Cognitive Approach*. The Open University Press, Milton Keynes, UK.
- Rotman, D., Hammock, J., Preece, J., Boston, C. L., Hansen, D. L., Bowser, A., and He, Y. (2014). Does Motivation in Citizen Science Change with Time and Culture? In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW Companion '14)*. ACM, New York, NY, 229–232.
- Rotman, D., He, Y., Preece, J., and Druin, A. (2013) Understanding Large Scale Online Environments with Qualitative Methods. *iConference*, February 2012, Texas.
- Rotman, D., Preece, J., He, Y., and Druin, A. (2012) Extreme Ethnography: Challenges for Research in Large Scale Online Environments *iConference*, February 7–10, 2012, Toronto, Ontario, Canada.
- Russell, D. M., and Yarosh, S. (2018). Can We Look to Science Fiction for Innovation in HCI? *Interactions* 25, 2, 36–40.
- Ryall, K., Forlines, C., Shen, C., and Ringel-Morris, M. (2004) Exploring the Effects of Group Size and Table Size on Interactions with Tabletop Shared-Display Groupware. In *Proceedings of Conference on Computer Supported Cooperative Work (CSCW)*. ACM, New York.
- Sacks, H., Schegloff, E., and Jefferson, G. (1978) A Simplest Systematics for the Organization of Turn-taking for Conversation, *Language* 50, 696–735.
- Saffer, D. (2010) *Designing for Interaction: Creating Smart Applications and Clever Devices* (2nd edn). New Riders Press, Indianapolis, IN.
- Saffer, D. (2014) *Microinteractions: Designing with Details*. O'Reilly.
- Sakr, S., Bajaber, F., Barnawi, A., Altalhi, A., Elshawi, R., and Batarfi, O. (2015) Big Data Processing Systems: State-of-the-Art and Open Challenges. In *Proceedings of ICCS 2015*.
- Sambrooks, L., and Wilkinson, B. (2013) Comparison of Gestural, Touch, and Mouse Interaction With Fitts' Law. In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration (OzCHI '13)*, ACM, New York, NY, pp. 119–122.
- Sarikaya, A., Correll, M., Bartram, L., Tory, M., and Fisher, D. (2018) What Do We Talk About When We Talk About Dashboards? *IEEE Trans Vis Comput. Graph*, 1.
- Sas, C., and Whittaker, S. (2013) Design for Forgetting: Disposing of Digital Possessions After a Breakup. In *Proceedings of CHI '13*. ACM, pp.1823–1832.
- Scaife, M., and Rogers, Y. (1996) External Cognition: How Do Graphical Representations Work? *International Journal of Human-Computer Studies* 45, 185–213.
- Scapin, D. L. (1981) Computer Commands in Restricted Natural Language: Some aspects of Memory of Experience, *Human Factors* 23, 365–375.
- Schaffer, E. (2009) Beyond Usability: Designing Web Sites for Persuasion, Emotion and Trust. Downloaded from www.uxmatters.com/mt/archives/2009/01/beyond-usability-designing-web-sites-for-persuasion-emotion-and-trust.php (retrieved 8 March, 2019).
- Schank, R. C. (1982) *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press, Cambridge.

- Schegloff, E. (1981) Discourse as an Interactional Achievement: Some Uses of ‘Uh-Huh’ and Other Things that Come Between Sentences. In D. Tannen (ed.) *Analyzing Discourse: Text and talk*. University Press, Georgetown.
- Schegloff, E. A., and Sacks, H. (1973) Opening Up Closings, *Semiotica* 7, 289–327.
- Schilit, B., Adams, N., Gold, R., Tso, M., and Want, R. (1993) The PARCTAB Mobile Computing System. In *Proceedings of Fourth Workshop on Workstation Operating Systems, WWOS-IV*. IEEE, pp. 34–39.
- Schmidt, A. (2017a) Technologies to Amplify the Mind. *IEEE Computer*, 50(10), 102–106.
- Schmidt, A. (2017b) Augmenting Human Intellect and Amplifying Perception and Cognition. *IEEE Pervasive Computing*, vol. 16, no. 1, pp. 6–10, 2017.
- Schmidt, A., and Herrmann, T. (2017) Intervention User Interfaces: A New Interaction Paradigm for Automated Systems, *Interactions*, vol 24, no 5, pp. 40–45.
- Schmitz, K., Mahapatra, R., and Nerur, S. (2018) User Engagement in the Era of Hybrid Agile Methodology. *IEEE Software* (early access).
- Schnall, R., Cho, H., and Liu, J. (2018) Health Information Technology Usability Evaluation Scale (Health-ITUES) for Usability Assessment of Mobile Health Technology: Validation Study. *JMIR Mhealth Uhealth.*, 6(1):e4.
- Schön, D. (1983) *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, New York.
- Schuler, R. P., Grandhi, S. A., Mayer, J. M., Ricken, S. T., and Jones, Q. (2014) The Doing of Doing Stuff: Understanding the Coordination of Social Group-Activities. In *Proceedings of CHI '14*, ACM, 119–128.
- Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N.J., and Griskevicius, V. (2007) The Constructive, Destructive, and Reconstructive Power of Social Norms, *Psychological Science*, 18(5), 429–434.
- Schwaber, K., and Beedle, M. (2002) *Agile Software Development with Scrum*. Prentice Hall, Englewood Cliffs, NJ.
- Seffah, A., Gulliksen, J., and Desmarais, M. C. (2005) *Human-Centered Software Engineering*. Springer.
- Segura, V.C.B., Barbosa, S.D.J., and Simões, F.P. (2012) UISKEI: A Sketch-Based Prototyping Tool for Defining and Evaluating User Interface Behavior. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, 18–25.
- Sentance, S. Waite, J., Hodges, S., MacLeod, E., and Yeomans, L. (2017) Creating Cool Stuff: Pupils’ Experience of the BBC micro:bit. In *Proceedings of SIGCSE 2017*, 531–536.
- Sethu-Jones, G.R., Rogers, Y., and Marquardt, N. (2017) Data in the Garden: A Framework for Exploring Provocative Prototypes as Part of Research in the Wild. In *Proceedings of the 29th Australian Conference on Computer-Human Interaction (OzCHI '17)*, ACM, New York, NY, pp. 318–327.
- Shaer, O., and Hornecker, E. (2010) Tangible User Interfaces: Past, Present and Future Directions, *Foundations and Trends in HCI (FnT in HCI)* 3(1–2), 1–138.
- Shaer, O., Strait, M., Valdes, C., Wang, H., Feng, T., Lintz, M., Ferreirae, M., Grote, C., Tempel, K., and Liu, S. (2012) The Design, Development, and Deployment of a Tabletop Interface for Collaborative Exploration of Genomic Data, *International Journal of Human-Computer Interaction* 70, 746–764.

- Sharp, H., Biddle, R., Gray, P. G., Miller, L., and Patton, J. (2006) Agile Development: Opportunity or Fad? Addendum to *Proceedings of CHI 2006*, Montreal.
- Sharp, H., Galal, G. H., Finkelstein, A. (1999) Stakeholder Identification in the Requirements Engineering Process. In: *Proceedings of the Database and Expert System Applications Workshop (DEXA)*, pp. 387–391.
- Sharp, H., Robinson, H.M., and Petre, M. (2009) The Role of Physical Artefacts in Agile Software Development: Two Complementary Perspectives, *Interacting with Computers*, 21(1-2) 108–116.
- Shen, C., Everitt, K., and Ryall, K. (2003) UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces. In *Proceedings of Ubicomp 2003*, pp. 281–288.
- Shen, C., Lesh, N. B., Vernier, F., Forlines, C., and Frost, J. (2002) Building and Sharing Digital Group Histories. In *Proceedings CSCW 2002*. ACM, New York, NY, pp. 324–333.
- Shilton, K. (2018) Values and Ethics in Human-Computer Interaction, *Foundations and Trends in Human-Computer Interaction*, 12(2), 107–171.
- Shneiderman, B. (1983) Direct Manipulation: A Step Beyond Programming Languages, *IEEE Computer*, 16(8), 57–69.
- Shneiderman, B. (1992) Tree Visualization with Tree-Maps: 2-d Space-Filling Approach. *ACM Transactions on Graphics*, 11(1), 92–99.
- Shneiderman, B. (1998) *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (3rd edn). Addison-Wesley, Reading, MA.
- Shneiderman, B., and Plaisant, C. (2006) Strategies for Evaluating Information Visualization Tools: Multi-Dimensional In-Depth Long-Term Case Studies. In *Proceedings Beyond Time and Errors: Novel Evaluation Methods for Information Visualization. Workshop of the Advanced Visual Interfaces Conference*.
- Shneiderman, B., Plaisant, C., Cohen, M., Jacobs, S., and Elmquist, N. (2016) *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (6th ed.). Pearson.
- Sidner, C., and Lee, C. (2005) Robots as Laboratory Hosts, *Interactions* 12, 24–26.
- Siek, K. A., Rogers, Y., and Connelly, K. H. (2005) Fat Finger Worries: How Older and Younger Users Physically Interact with PDAs. In *Proceedings of INTERACT '05*, Rome.
- Silver, J., and Rosenbaum, E. (2012). Makey Makey: Improvising Tangible and Nature-Based User Interfaces. In *Adjunct Proceedings of TEI'12*.
- Silverberg, M., MacKenzie, I. S., and Korhonen, P. (2000) Predicting Text Entry Speed on Mobile Phones. In *Proceedings of CHI 2000*. ACM, New York, NY, pp. 9–16.
- Silverio-Fernández, M., Renukappa S., and Suresh, S. (2018) What is a Smart Device?—A Conceptualisation within the Paradigm of the Internet of Things. *Visualization in Engineering*, 6:3.
- Sim, G., Horton, M., and McKnight, L. (2016) iPad vs Paper Prototypes: Does Form Factor Affect Children's Ratings of a Game Concept. In *Proceedings of Interaction Design and Children (IDC '16)*, pp. 190–195.
- Simonsen, J., and Robertson, T. (2012) *Routledge Handbook of Participatory Design*. Routledge, London.
- Singer, L.M., and Alexander, P.A. (2017). Reading on Paper and Digitally: What the Past Decades of Empirical Research Reveal. *Review of Educational Research*, 87(6), 315–343.

- Singer, P. (2011) *The Expanding Circle Ethics, Evolution, and Moral Progress*. Princeton University Press.
- Sitbon, L., and Farhin, S. (2017) Co-Designing Interactive Applications with Adults with Intellectual Disability: A Case Study. In *Proceedings of the 29th Australian Conference on Computer-Human Interaction (OzCHI '17)*. ACM, New York, NY, pp. 487–491.
- Slater, M., and Wilbur, S. (1997) A Framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual Environments, *Presence: Teleoperators and Virtual Environments* 6, 603–616.
- Slater, M., Pertaub, D., and Steed, A. (1999) Public Speaking in Virtual Reality: Facing an Audience of Avatars. *IEEE Computer Graphics and Applications*, 19(2), 6–9.
- Sleeper, M., Consolvo, S., and Staddon, J. (2014) Exploring the Benefits and Uses of Web Analytics Tools for Non-Transactional Websites. In *ACM Proceedings of the 2014 Conference on Designing Interactive Systems (DIS)*.
- Smith, D., Irby, C., Kimball, R., Verplank, B., and Harslem, E. (1982) Designing the Star User Interface, *Byte*, 7(4), 242–282.
- Smith, M.E., Ascenzi, L. Qin, Yingsi and Wetsman, R. (2018) Designing a Video Co-Watching Web App to Support Interpersonal Relationship Maintenance. In *Proceedings of GROUP '18*, 162–165.
- Smyth, J. D., Dillman, D. A., Christian, L. M., and Stern, M. J. (2005) Comparing Check-All and Forced-Choice Question Formats in Web Surveys: The Role of Satisficing, Depth of Processing, and Acquiescence in Explaining Differences. Technical Report #05-029. Washington State University Social and Economic Sciences Research Center, Pullman, 30 pages.
- Sohn, T., Li, K. A., Griswold, W. G., and Hollan, J. D. (2008) A Diary Study of Mobile Information Needs. In *Proceedings of CHI 2008*. ACM, New York.
- Solovey, E.T., Afergan, D., Peck, E., Hincks, S., and Jacob, R.J.K. (2014) Designing Implicit Interfaces for Physiological Computing: Guidelines and Lessons Learned Using fNIRS1. *ACM Transactions on Computer-Human Interaction*, 21(6).
- Sparrow, B., Liu, J., and Wegner, D. M. (2011) Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. *Science*, 333(6043), 308–314.
- Speicher, M., Hell, P., Daiber, F., Simeone, A., and Krüger, A. (2018) A Virtual Reality Shopping Experience Using the Apartment Metaphor. In *AVI '18: 2018 International Conference on Advanced Visual Interfaces*, 9 pages.
- Spelmezan, D., Jacobs, M., Hilgers, A., and Borchers, J. (2009) Tactile Motion Instructions for Physical Activities. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009*. ACM, New York, NY, pp. 2243–2252.
- Spencer, R. (2000) The Streamlined Cognitive Walkthrough Method: Working Around Social Constraints Encountered in a Software Development Company. In *Proceedings of CHI 2000*. ACM, New York, NY, pp. 253–359.
- Spool, J. (2018) <https://medium.com/@jmsspool> (accessed December 2018).
- Starbird, K., Palen, L., Hughes, A. L., and Vieweg, S. (2010) Chatter on the Red: What Hazards Threat Reveals about the Social Life of Microblogged Information. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW 2010*. ACM, New York, NY, pp. 241–250.

- Stavrinos, D., Jones, J.L., Garner, A.A., Griffin, R., Franklin, C.A., Ball, D., Welburn, S.C., Ball, K.K., Sisiopiku, V.P., and Fine, P.R. (2013) Impact of Distracted Driving on Safety and Traffic Flow. *Accident Analysis and Prevention*, **61**, 63–70.
- Steed, A., Ye, P., Zillah, W., and Slater, M. (2018) “We Wait”—The Impact of Character Responsiveness and Self Embodiment on Presence and Interest in an Immersive News Experience. *Frontiers in Robotics and AI*, **5**, 112.
- Stel, M., and Vonk, R. (2010) Mimicry in Social Interaction: Benefits for Mimickers, Mimickees, and their Interaction. *British Journal of Psychology*, **101**(2), 311–323.
- Stephens-Davidowitz, S. (2018) *Everybody Lies: Big Data, New Data, and What the Internet Can Tell Us about Who We Really Are*. Penguin.
- Stobert, E., and Biddle, R. (2018) The Password Life Cycle. *Transactions on Privacy and Security*, **21**(3), Article 13, 32 pages.
- Strauss, A., and Corbin, J. (1998) *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory* (2nd edn). SAGE, London.
- Strommen, E. (1998) When the Interface is a Talking Dinosaur: Learning Across Media with ActiMates Barney. In *Proceedings of CHI '98*. ACM, New York, NY, pp. 288–295.
- Subrayaman, R., Weisstein, F. L., and Krishnan, M. S. (2010) User Participation in Software Development Projects, *Communications of the ACM*, **53**(3), 137–141.
- Suzuki, K., Yokoyama, M., Yoshida, S., Mochizuki, T., Yamada, T., Narumi, T., Tanikawa, T., and Hirose, M. (2017) FaceShare: Mirroring with Pseudo-Smile Enriches Video Chat Communications. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York. pp. 5313–5317.
- Swallow, E. (2013) The U.S. Senate More Divided Than Ever Data Shows, *Forbes*, Downloaded from: <https://www.forbes.com/sites/ericaswallow/2013/11/17/senate-voting-relationships-data/#335036bf4031>.
- Swan, M. (2013) The Quantified Self: Fundamental Disruption in Big Data Science and Biological Discovery. *Big Data*, **1**(2).
- Sy, D. (2007) Adapting Usability Investigations for Development, *Journal of Usability Studies*, **2**(3), May, 112–130.
- Szafir, D. (2018) The Good, the Bad and The Biases: Five Ways Visualizations Can Mislead and How to Fix Them, *Interactions*, Xxv.4.
- Tallyn, E., Fried, H., Gianni, R., Isard, A., Speed, C. (2018) The Ethnobot: Gathering Ethnographies in the Age of IoT. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, Paper 604, 13 pages.
- Tanenbaum, J. (2014) Design Fictional Interactions: Why HCI Should Care About Stories, *Interactions*, XXI.5, 22–23.
- Teixeira, C.R.G., Kurtz, G., Leuck, L.P., Tietzmann, R., Souza, D.R., Lerina, J.M.F., Manssour, I.H., and Silveira, M.S. (2018) Humor, Support, and Criticism: A Taxonomy for Discourse Analysis About Political Crisis on Twitter. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age*. ACM, New York, NY, Article 68, 6 pages.
- Thackara, J. (2001) The Design Challenge of Pervasive Computing. In *Interactions* May/ Jun, 47–52.
- Thaler, R. H., and Sunstein, C. R. (2008) *Nudge: Improving Decisions About Health, Wealth and Happiness*. Penguin.

- Thimbleby, H. (1990) *User Interface Design*. Addison-Wesley, Harlow, Essex.
- Thimbleby, H. (2015) Safer User Interfaces: A Case Study in Improving Number Entry. *IEEE Transactions on Software Engineering* **41** (7), 711–729.
- Tidwell, J. (2006) *Designing Interfaces: Patterns for Effective Interaction Design*. O'Reilly Media Inc.
- Todd, P.M., Rogers, Y., and Payne, S.J. (2011) Nudging the Trolley in the Supermarket: How to Deliver the Right Information to Shoppers. *International Journal of Mobile HCI*, **3**(2), 20–34.
- Toepoel, V. (2016) *Doing Surveys Online*. SAGE Publications Ltd.
- Toeteneel, L., and Rienties, B. (2016) Analyzing 157 Learning Designs Using Learning Analytic Approaches as a Means to Evaluate the Impact of Pedagogical Decision-Making. *British Journal of Educational Technology*, **4**(5), 981–992.
- Tognazzini, B. (2014) First Principles of HCI Design, Revised and Expanded. Downloaded from: asktog.com/atc/principles-of-interaction-design/
- Tomlin, W. C. (2010) Usability, www.usefulnessability.com/ (accessed May 1, 2010).
- Towey, M., Zhang, L., Cottman-Fields, M., Wimmer, J., Zhang, J., and Roe, P. (2014) Visualization of Long Duration Acoustic Recordings of the Environment. *Procedia Computer Science*, **29**, 703–712.
- Tractinsky, N. (2013) Replicating and Extending Research on Relations between Visual Aesthetics and Usability. In *ReplicaCHI 2013*.
- Travis, D. (2016) 247 Web Usability Guidelines. Downloaded from: <https://www.userfocus.co.uk/resources/guidelines.html>.
- Trimble, J., Wales, R., and Gossweiler, R. (2002) *NASA Position Paper for the CSCW 2002 Workshop on Public, Community and Situated Displays: MERBoard*.
- Trist, E.L., and Bamforth, K.W. (1951) Some Social and Psychological Consequences of the Longwall Method of Coal Getting. *Human Relations*, **4**, 3–38.
- Tullis, T., and Albert, B. (2013) *Measuring the User Experience* (2nd ed.). Morgan Kaufmann.
- Tullis, T. S. (1997) Screen Design. In M. Helander, T. K. Landauer and P. Trabhu (eds) *Handbook of Human-Computer Interaction* (2nd edn). Elsevier, New York, pp. 377–411.
- Turkle, S. (2015) *Reclaiming Conversation: The Power of Talk in a Digital Age*. Penguin.
- Ullmar, B., Ishii, H., and Jacob, R. J. K. (2005) Token + Constraint Systems for Tangible Interaction with Digital Information. *TOCHI*, **12**(1), 81–N8.
- Underkoffler, J., and Ishii, H. (1998) Illuminating Light: An Optical Design Tool with a Luminous-Tangible Interface. In *Conference In Proceedings on Human Factors in Computing Systems*. ACM Press/Addison-Wesley, pp. 542–549.
- Unger, R., and Chandler, C. (2012) *A Project Guide to UX Design*. New Riders, Berkeley, CA.
- Väänänen-Vainio-Mattila, K., and Waljas, M. (2009) Development of Evaluation Heuristics for Web Service User Experience. *CHI 2009 Spotlight on Works in Progress, Session 1*, pp. 3679–3684.
- Vaish, R., Snehalakumar, S., Gaikwad, N., Kovacs, G., Veit, A., Krishna, R., Ibarra, I.A., Simoiu, C., Wilber, M., Belongie, S., Goel, S. Davis, J., and Bernstein, M.S. (2017) Crowd Research: Open and Scalable University Laboratories. In *Proceedings of the 30th Annual*

- ACM Symposium on User Interface Software and Technology (UIST '17). ACM, New York, NY, pp. 829–843.
- Valdesolo, P., and DeSteno, D. (2011) Synchrony and The Social Tuning of Compassion. *Emotion*, **11**, 2, 262–266.
- Valentina, G., and Mohanna, M. (2013) Informal Cognitive Walkthroughs (ICW): Paring Down and Pairing Up for an Agile World. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, pp. 3093–3096.
- van Allen, P. (2018) Ways of Prototyping AI, *Interactions*, 47–51.
- van Berkel, N., Ferreira, D., and Kostakos, V. (2017) The Experience Sampling Method on Mobile Devices. *ACM Computing Surveys*, **50**, 6, Article 93.
- van den Broek, E.L. (2013) Ubiquitous Emotion-Aware Computing. *Personal and Ubiquitous Computing*, **17**, 53–67.
- van den Hoven, P., Pieter E. Vermaas, P.E., and van de Poel, I. (eds.) (2015) *Handbook of Ethics, Values, and Technological Design*. Springer.
- van der Linden, J., Schoonderwaldt, E., Bird, J., and Johnson, R. (2011) MusicJacket—Combining Motion Capture and Vibrotactile Feedback to Teach Violin Bowing. *IEEE Transactions on Instrumentation and Measurement*, **60**(1), pp. 104–113.
- van Rens, L. S. (1997) *Usability Problem Classifier*. Unpublished master's thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA.
- Veen, J. (2001) *The Art and Science of Web Design*. New Riders Press, Indianapolis, IN.
- Verma, H., Alavi, H. S., and Lalanne, D. (2017) Studying Space Use: Bringing HCI Tools to Architectural Projects. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI 2017)*. ACM, Press, New York, pp. 3856–3866.
- Veronikha, E. (2016) Sentiment Analysis on Twitter about the Use of City Public Transportation Using Support Vector Machine Method. *International Journal on Information and Communication Technology (IJoICT)*, **2**, 57.
- Vertegaal, R. (2008) A Fitts' Law Comparison of Eye Tracking and Manual Input in the Selection of Visual Targets, *ICMI 2008*, October 20–22, Chania, Crete, Greece, pp. 241–248.
- Villar, N., Scott, J., Hodges, S., Hammil, K., and Miller, C. (2012), NET Gadgeteer: A Platform for Custom Devices. *Pervasive*, 216–233.
- VisiStat (2010) Case Study: The MountainWinery.com, visistat.com/case-study-mountain-winery.php (accessed September 2010).
- von Neumann, J., and Morgenstern, O. (1944) *Theory of Games and Economic Behavior*. Princeton University Press.
- Wallace, J., McCarthy, J., Wright, P. C., and Olivier, P. (2013) Making Design Probes Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, pp. 3441–3450.
- Wang, P., Sibi, S., Mok, B., and Ju, W. (2017) Marionette: Enabling On-Road Wizard-of-Oz Autonomous Driving Studies. In *Proceedings of Human-Robot Interaction*, 234–243.
- Wang, Y., Song, G., Qiao, G., Zhang, Y., Zhang, J., and Wang, W. (2013) Wheeled Robot Control Based on Gesture Recognition Using the Kinect Sensor, In *Proceedings of the IEEE International Conference on Robotics and Biomimetics (ROBIO)*. Shenzhen, China, December 2013, pp. 378–383.

- Warrick, E., Preece, J., Kibutu, J., and Sihanya, B. (2016) Social Media as an Indigenized Information World for Environmental Stewardship. In *Proceedings of the First African Conference on Human Computer Interaction (AfriCHI)*, pp. 126–137.
- Wasserman, S., and Faust, K. (1994) *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK.
- Waterson, P. (2014) Health Information Technology and Sociotechnical Systems: A Report on Recent Developments within the UK National Health Service (NHS), *Applied Ergonomics*, 45, 150–161.
- Watson, N., and Naish, H. (2018) Can Computers Understand Human Emotions? A Sentiment Analysis of Economic News. Downloaded from: <http://blogs.ucl.ac.uk/ippr/can-computers-understand-human-emotions-a-sentiment-analysis-of-economic-news/>.
- Webster, D., and Celik, O. (2014) Systematic Review of Kinect Applications in Elderly Care and Stroke Rehabilitation, *Journal of NeuroEngineering and Rehabilitation*, 11:108.
- Weiser, M. (1991) The Computer for the 21st Century, *Scientific American*, 94–104.
- Weller, D. (2004) The Effects of Contrast and Density on Visual Web Search, *Usability News*, 6(2): <http://psychology.wichita.edu/surl/usabilitynews/62/density.htm> (retrieved July 11, 2005).
- Wellman, B., and Berkovitz, S.D., (1988) *Social Structures: A Network Approach*. Cambridge University Press, Cambridge, UK.
- Wharton, C., Rieman, J., Lewis, C., and Polson, P. (1994) The Cognitive Walkthrough Method: A Practitioner’s Guide. In J. Nielsen and R. L. Mack (eds). *Usability Inspection Methods*. John Wiley & Sons Inc., New York.
- Whitenton, K. (2018) The Two UX Gulfs: Evaluation and Execution. Downloaded from: <https://www.nngroup.com/articles/two-ux-gulfs-evaluation-execution/>.
- Whitenton, K., and Gibbons, S. (2018) Case Study: Iterative Design and Prototype Testing of the NN/g Homepage. Downloaded from: <https://www.nngroup.com/articles/case-study-iterative-design-prototyping>
- Whiteside, J., Bennett, J., and Holtzblatt, K. (1988) Usability Engineering: Our Experience and Evolution. In H. Helander (ed.) *Handbook of Human-Computer Interaction*. Elsevier Science Publishers, Amsterdam, pp. 791–817.
- Whitney, H. (2012) *Data Insights: New Ways to Visualize and Make Sense of Data*. Morgan Kaufmann Publ., San Francisco, CA.
- Williamson, V. (2016) Can Crowdsourcing Be Ethical? Downloaded from: <https://www.brookings.edu/blog/techtank/2016/02/03/can-crowdsourcing-be-ethical-2/>
- Wilson, C., Hargreaves, T., and Hauxwell-Baldwin, R. (2015) Smart Homes and their Users: A Systematic Analysis and Key Challenges, *Personal and Ubiquitous Computing*, 19, 463–476.
- Winograd, T. (1997) From Computing Machinery to Interaction Design. In P. Denning and R. Metcalfe (eds). *Beyond Calculation: The Next Fifty Years of Computing*. Springer-Verlag, Amsterdam, pp. 149–162.
- Winschiers-Theophilus, H., and Bidwell, N.J. (2013) Toward an Afro-Centric Indigenous HCI Paradigm. *International Journal of Human-Computer Interaction*, 29: 243–255.
- Winschiers-Theophilus, H., Bidwell, N. J., and Blake, E. (2012) Community Consensus: Design Beyond Participation, *Design Issues*, 28(3) Summer 2012, 89–100.

- Wixon, D., and Wilson, C. (1997) The Usability Engineering Framework for Product Design and Evaluation. In M. G. Helander, T. K. Landauer and P. V. Prabhu (eds) *Handbook of Human-Computer Interaction*. Elsevier, Amsterdam, pp. 653–688.
- Wong, R.Y., Van Wyk, E., and Pierce, J. (2017) Real-Fictional Entanglements: Using Science Fiction and Design Fiction to Interrogate Sensing Technologies. In *Proceedings of DIS 2017*. ACM, New York, NY, pp. 567–579.
- Woolrych, A., and Cockton, G. (2001) Why and When Five Test Users Aren't Enough. In *Proceedings of IHM-HCI 2001 Conference*, Vol. 2. Cépadèus Éditions, Toulouse, pp. 105–108.
- Xambó, A., Hornecker, E., Marshall, P., Jordà, S., Dobbyn, C., and Laney, R. (2013) Let's Jam The Reactable: Peer Learning During Musical Improvisation with a Tabletop Tangible Interface. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(6), article no. 36.
- Yeratziotis, A., and Zaphiris, P. (2018) A Heuristic Evaluation for Deaf Web User Experience (HEADWUX). *International Journal of Human-Computer Interaction*, 34:3, 195–217.
- Yin R. K. (2013) *Case Study Research*. SAGE Publications.
- Yohanan, S., and MacLean, K. E. (2008) The Haptic Creature Project: Social Human-Robot Interaction Through Affective Touch. In *Proceedings of the AISB 2008 Symposium on the Reign of Cats and Dogs: The Second AISB Symposium on the Role of Virtual Creatures in a Computerized Society*. 1, 7–11.
- Yu, L., Kittur, A., and Kraut, R. E. (2016) Encouraging “Outside-the-box” Thinking in Crowd Innovation Through Identifying Domains of Expertise. In *Proceedings of CSCW '16*. ACM, New York, NY, pp. 1214–1222.
- Yuill, N., and Rogers, Y. (2012) Mechanisms for Collaboration: A Design and Evaluation Framework for Multi-User Interfaces. *ACM Transactions on Computer-Human Interaction*, 19(1), Article 1, 25 pages.
- Zeiliger, R., Reggers, T., Baldewyns, L., and Jans, V. (1997) Facilitating Web Navigation: Integrated Tools for Active and Cooperative Learners. In *Proceedings of the 5th International Conference on Computers in Education, ICCE '97*, Kuching, Sarawak, Malaysia.
- Zhai, W., and Thill, J-C. (2017) Social Media Discourse in Disaster Situations: A Study of the Deadly. In *Proceedings of EM-GIS' 17 Proceedings of the 3rd ACM SIGSPATIAL Workshop on Emergency Management*.
- Zhang, T., and Chan, A.H.S. (2011) The Association Between Driving Anger and Driving Outcomes: A Meta-Analysis of Evidence from the Past Twenty Years. *Accident Analysis and Prevention*, 90, 50–62.
- Zhao, S., Ramos, J., Tao, J., Jiang, Z., Li, S., Wu, Z., Pan, G., and Dey, A.K. (2016) Discovering Different Kinds of Smartphone Users Through Their Application Usage Behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, pp. 498–509.
- Zuckerman, O., and Resnick, M. (2005) Extending Tangible Interfaces for Education: Digital Montessori-Inspired Manipulatives. In *Proceedings of Conference on Human Factors in Computing Systems*. ACM, New York, NY, pp. 859–868.
- Zufferey, G., Jermann, P., Lucchi, A., and Dillenbourg, P. (2009) TinkerSheets: Using Paper Forms to Control and Visualize Tangible Simulations. In *Proceedings of TEI09*. ACM, New York, NY, pp. 377–384.

Index

- 3D printing, 423
- 404 error, 177
- A/B testing, 62, 63, 574–575
- academic disciplines, interaction design and, 10
- accessibility
 - design and, 17–18
 - WCAG (Web Content Accessibility Guidelines), 446
- accountability, AI and, 91
- activities
 - 404 error, 177
 - alternative designs, 59
 - anthropomorphism, 188
 - appliances, 223
 - assumptions, 72–73
 - bank security, 116–117
 - behavioral changes, 183
 - claims, 72–73
 - cognition, information structure and, 105
 - cognitive walk-through, 564–565
 - computer game investigation case study, 508
 - contextual inquiry, 400
 - conversation breakdown, 142
 - critical incident analysis, 328
 - dashboard, 373–374
 - data analysis, 317–318
 - data gathering for readers, 274–275
 - data recording, 268
 - data visualization, 369
 - design, 31–32
 - design patterns, 486–488
 - design thinking, 456–457
 - double diamond design, 39–41
 - ethnographic data at Royal Highland Show, 512–513
 - evaluation, 497, 498, 499
 - face-to-face conversations, 139
 - Fitbit Chart evaluation, 502
 - Google Analytics, 570–571
 - Google Trends, 352
 - GUIs (graphical user interfaces), 201–202
 - Happyornot terminal, 167
 - heuristic comparison, 560–561
 - icons, 207–208
 - in-depth
 - conceptual models, 95
 - interfaces, 254–255
 - mental models, 132–133
 - timepiece, 65
 - instructing interaction type, 82–83
 - interaction design tools, 492
 - interface metaphors, 78, 441–442
 - iPad usability testing, 530, 531
 - iPods, 14–15
 - Lean UX, 480
 - memory, 112
 - mental models, 123–124
 - menus, 204–205
 - mobile interfaces, 221
 - multimedia interfaces, 211–212
 - navigation app, 443–444
 - observation for data gathering,
 - 287–288
 - in the field, 289–290
 - think-aloud technique, 296–297
 - online likes, 377
 - Painpad field study, 541
 - personas, 406–407
 - Polar A370 evaluation, 502
 - prototype scenarios, 448–449
 - questionnaires, 283, 285–286
 - requirements, 394–395
 - scenarios, 411–412
 - shoplifting detection, 378–379
 - social interaction, 138
 - storyboarding, 428, 450–451
 - student activities, 365–366
 - team makeup, 12
 - telepresence, 148–149
 - think-aloud protocol of analysis, 326–327
 - usability, 63–64
 - user experience aspects, 23–24
 - user identification, 56–57
 - user involvement, 49
 - user research, 476
 - website design, 219
 - activity tracker observation
 - data scraping, 299–300
 - diaries, 298–299
 - interaction logs, 299–300
 - web analytics, 299–300
 - adaptive interfaces, 194
 - advertising, emotional interaction and, 168
 - Affdex emotional AI, 179
 - classification of emotions, 179–180

- affect, 169
- Affectiva emotional AI, 181
- affective computing, 166, 179–182
- affective interfaces, 194
- affinity diagrams, 323
- affordance, 30
- agent-based software, 175
- AgileUX, 473–475
 - aligning work practices, 477–480
 - BDUF (big design up front), 477–478
 - cycle, 475, 478–480
 - documentation, 481–484
 - iteration zero, 475
 - sprint, 475
 - timebox, 475
 - user research, 473–477
- AI (artificial intelligence), 90–91
 - accountability and, 91
 - affective computing, 166
 - control, 93–94
 - emotional AI, 166, 179–182
 - affective expression in speech, 179
 - biosensors, 179
 - body movement sensors, 179
 - cameras, 179
 - facial expressions, 179
 - gesture sensors, 179
 - transparency and, 91
 - visions, 91
- All the Feels, 181
- alternative courses, use cases, 415
- alternative designs, 50
 - cross-fertilization, 58
 - generating, 58–59
 - prototyping, 62
 - selecting, 59–64
 - usability engineering, 62
- alternative hypothesis, 533
- Alzheimer's disease, 117
- Amazon Echo, 7
- ambient interfaces, 194
- analytic frameworks, 329–341
 - content analysis, 329, 332–333
 - conversation analysis, 329, 330–331
 - voice assistants, 330
 - discourse analysis, 329, 331–332
 - grounded theory, 329, 334–338
 - interaction analysis, 329, 333–334
 - systems-based, 329, 338–341
- analytical evaluation, 515
- analytics, 506, 514
 - learning analytics, 568
 - visual analytics, 568
- web analytics, 567–568
 - ClickTale, 572
 - Clicky, 572
 - Crazy Egg, 572
 - Google Analytics, 568–570
 - KISSmetrics, 572
 - Mox Analytics, 572
 - off-site, 568
 - on-site, 568
 - TruSocialMetrics, 572
- analyzing data. *See* data analysis
- anatomy of prototypes, 430
- annotating, external cognition and, 130–131
- annoying interfaces, 174–179
 - 404 error, 177
 - Anna from IKEA, 176–177
 - Clippy, 175–176
- anthropomorphism, 187–190
- anti-aliasing, 206
- The App Generation* (Gardner and Davis), 122
- Apple, VoiceOver, 17
- appliances, 222–223
- applications, 30–31
- approaches to interaction design, 43
- AR (augmented reality), 152
- assumptions, 71–72
- ATM (air traffic management), 58
- atomic requirements shell, 388
- attention
 - design implications, 108
 - multitasking, 105–109
 - phone use while driving, 107–108
- audio, data recording and, 267
- augmented reality (AR) interfaces
 - decision making, 241–242
 - history, 243
 - HUDs (head-up displays), 242
 - mapping, 242–243
 - mixed reality, 241
 - research and design, 244
 - virtual content, 243
- Autodesk, SketchBook, 479
- automatic emotions, 169
- averages
 - mean, 312
 - median, 312
 - mode, 312
 - outliers, 314
 - quantitative analysis, 311–312
- Aware Home, 503
- axial coding, grounded theory, 334, 337–338
- Axure RP, 491

- Balsamiq, 491
- Barney the dinosaur, 187
- BDUF (big design up front), 477–478
- behavioral change, 182–186
 - environmental concerns, 185
 - Fitbit, 184
 - HAPIfork, 183–184
 - social norms, 185
- behavioral level of the brain, 171
- bias, evaluation and, 515
- big data, 89, 350. *See also* data at scale
- biometric data, 181
- brain-computer interfaces (BCI), 250
- brain levels, 171
 - behavioral level, 171
 - reflective level, 171
 - visceral level, 171
- breadcrumb navigation, 217
- Break-Time Barometer, 157
- BSI (British Standards Institute), 500

- CAD, command interfaces, 195–196
- cameras, 117–118
- Caplin Systems, 405–406
- capturing requirements, 386
- Cardiac Tutor*, 209
- cascading menus, 203
- case studies, 302
 - deaf telephony, 455–456
 - evaluation
 - computer game investigation, 507–509
 - ethnographic data at Royal Highland Show, 510–513
 - results, 514–515
 - field study Painpad, 538–540
 - prototyping, paper, 432
 - UX design, 483–484
- categorizing data, 324–327
 - grounded theories, 334
- CHI (Computer-Human Interaction)
 - conference, 262
- chunks of information, 114–115
- The Circle* (Eggers), 411
- citizen science, 354–356, 514
- claims, 71–72
- ClickTale, 572
- Clicky, 572
- Clippy, 175–176
- co-design, 46
- co-presence, social interaction and
 - awareness, 150–151
 - physical coordination, 150
 - shareable interfaces, 152–157

- Cogapp, 13
- cognition
 - attention, 102, 103
 - clear goals, 103
 - design implications, 108
 - information presentation, 103
 - multitasking, 105–109
 - context, 102
 - decision-making, 103
 - distributed cognition, 127–129
 - experiential, 102
 - external cognition
 - annotating, 130–131
 - cognitive tracing, 130–131
 - computational offloading, 130
 - memory load reduction, externalization, 129–130
 - fast thinking, 102
 - information structure, 104
 - learning, 103, 119
 - decision-making, 121–122
 - interactive apps, 120
 - interactive books, 120
 - listening, 120–121
 - natural-language interfaces, 120
 - planning, 121–122
 - problem-solving, 121–122
 - reading, 120–121
 - reasoning, 121–122
 - speaking, 120–121
 - speech recognition systems, 120
 - tactile interfaces, 120
 - listening, 103
 - memory, 103, 111
 - memory load, 113–118
 - personal information management, 112–113
 - perception, 102, 109
 - design implications, 109
 - proprioceptors, 109
 - planning, 103
 - problem-solving, 103
 - processes, 102–103
 - reading, 103
 - reasoning, 103
 - reflective, 102
 - slow thinking, 102
 - speaking, 103
- cognitive frameworks
 - distributed cognition, 127–129
 - embodied interaction, 131–132
 - external cognition
 - annotating, 130–131
 - cognitive tracing, 130–131

- computational offloading, 130
 - memory load reduction, externalization, 129–130
 - gulfs of evaluation, 125–126
 - gulfs of execution, 125–126
 - information processing, 126–127
 - reaction time, 126
 - mental models, 123–125
- cognitive impairment, 17
- cognitive tracing, external cognition and, 130–131
- cognitive walk-throughs, 506, 560–566
 - SSCW (streamlined cognitive walk-through), 566
- collapsible menus, 204
- collecting data. *See* data collection
- command interfaces, 194, 195–196
 - research and design, 196
 - TextSL, 196
 - virtual words, 196
- commercial projects, product owner, 44
- common ground, design space and, 74
- compromises in prototyping, 429–434
- computational offloading, external cognition and, 130
- computer game investigation case study, 507–509
- conceptual design, conceptual models, 434
 - expanding, 444–445
 - interaction types and, 442
 - interface metaphors and, 439–442
 - interface types and, 442–443
- conceptual models, 434
 - analogies, 74
 - expanding, 444–445
 - interaction types and, 432
 - interface metaphors, 439–442
 - interface types and, 432–443
 - mappings, 75
 - metaphors, 74
 - relationships between concepts, 74
- conceptualizing interaction, 71–72
 - design space
 - common ground, 74
 - open-mindedness, 74
 - orientation, 74
- conclusions *versus* data, 264–265
- concrete design, 444–445
 - accessibility, 445
 - inclusiveness, 445
- conscious emotions, 169
- consistency, 29
- console interfaces, 247
- constraints, 28–29
- construction, 421–422, 457
 - physical computing, 458–462
 - SDKs (software development kits), 463–464
- content analysis, 329, 332–333
- contextual design, 323, 435
- contextual inquiry, 400
 - apprenticeship model, 400
 - artifact model, 401
 - context principle, 400
 - cool concepts, 401
 - collaboration model, 401
 - day-in-the-life model, 401
 - relationship model, 401
 - decision point model, 401
 - flow model, 401
 - focus, 401
 - interpretation, 400
 - joy of life concepts
 - accomplish, 401
 - connection, 401
 - identity, 401
 - sensation, 401
 - joy of use concepts
 - direct in action, 401
 - hassle factor, 401
 - learning delta, 401
 - partnership principle, 400
 - physical model, 401
 - sequence model, 401
 - Wall Walk, 401
- contextual menus, 204
- controlled experiments, 515
- convenience sampling, data gathering, 261
- conversation analysis, 329, 330–331
 - voice assistants, 330–331
- conversational mechanisms, 140–141
- conversing interaction type, 81, 83–85
- Cooper, Alan, 2, 222–223
- Cooper, NielsonNorman Group, 13
- cooperative design, 46
- copyrights, 61
- counterbalancing, experiments, 535
- Crazy Egg, 572
- critical incident analysis, 320, 327–328
- cross-cultural design, 446
- cross-fertilization, alternative design and, 58
- CrowdFlower, 359
- crowdsourcing, 46, 353–358, 515
 - citizen science, 514
 - consent form, 263
 - evaluation, 513–514
- cultural differences among users, 16
- cultural probes, 398

- customer journey maps, 451–452
- CuteCircuit, 245
- D3 (Data-Driven Documents), 371
- dark patterns, 25
- data, 259–260
 - versus* conclusions, 264–265
 - versus* information, 264–265
 - numbers, abuse of, 309
 - personal, collecting, 353
 - qualitative, 308
 - quantitative, 308
 - second source, 352
 - socially acceptable use, 378
 - sources, combining, 364–366
- data analysis, 307–308, 309, 351
 - frameworks, 329–341
 - content analysis, 329, 332–333
 - conversation analysis, 329, 330–331
 - discourse analysis, 329, 331–332
 - grounded theory, 329, 334–338
 - interaction analysis, 329, 333–334
 - systems-based, 329, 338–341
 - interviews, 309–311
 - observation, 311
 - outliers, 314
 - presentation, 342–344
 - storytelling, 344–345
 - structured notations, 342–344
 - summarizing findings, 345–346
 - questionnaires, 311
 - question design and, 313–314
 - sentiment analysis, 358–359
 - SNA (social network analysis), 359–363
 - spreadsheets, 314
 - thematic analysis, 322–324
 - tools, 341
 - videos, 321–322
- data at scale, 349–350
- data cleansing, 308
- data collection
 - crowdsourcing, 353–358
 - field studies, 537
 - personal data, 353
 - scraping, 352
 - second source data, 352
 - sensing data, 356–357
 - sentiment analysis, 358–359
 - SNA (social network analysis), 359–363
 - sources, combining, 364–366
 - tracking reactions, 362–363
- data ethics canvas, 379
- data ethics principles, 379
- data gathering, 260
 - case studies, 302
 - focus groups, 268
 - goal setting, 260
 - grounded theory, 334
 - interviews
 - alternative types, 277
 - Cisco WebEx, 277
 - experience enrichment, 277–278
 - focus groups, 268, 271–272
 - question development, 272–276
 - retrospective, 277
 - running, 276
 - semi-structured, 269–271
 - Skype, 277
 - structured, 269
 - unstructured, 268–269
 - Zoom, 277
 - observation, 287
 - activity tracking, 298–300
 - controlled environments, 295–298
 - in the field, 288–295
 - participants
 - behaviors, 273
 - incentives, 264
 - informed consent, 262–264
 - number of, 261–262
 - recording, 267
 - relationships with, 262–264
 - study population, 261
 - variations in types, 270–271
 - pilot studies, 264
 - questionnaires, 278–279
 - administering, 283–285
 - design, 279
 - online *versus* paper and pencil, 286
 - question format, 280–283
 - response format, 280–283
 - structure, 279
 - requirements activity, 396–398
 - sampling, 261
 - convenience sampling, 261
 - saturation sampling, 261
 - snowball sampling, 261
 - technique choice, 300–303
 - technique combining, 300–303
 - triangulation, 264–265
- data interpretation, 307–308
- data presentation, 307–308
- data recording
 - audio, 267
 - notes, 266–267

- photos
 - and audio, 267
 - and notes, 266–267
- video, 267
- data requirements, 392
- data visualization, 366–375
 - D3 (Data-Driven Documents), 371
 - dashboard, 370, 371–372
 - sensing data, 374–375
 - spectrograms, 368–369
 - visual literacy, 366
- deaf telephony, 455–456
- deceptive technology, 186
- decision-making
 - design implications, 122
 - heuristics, 121
 - learning and, 121–122
 - theories, 121
- Dedoose, 341
- deductive approach to qualitative data, 321
- design, 421–422
 - accessibility and, 17–18
 - affordance, 30
 - appliances, 222–223
 - applications, 30–31
 - augmented reality (AR) interfaces, 244
 - brain-computer interfaces (BCI), 250
 - command interfaces, 196
 - consistency, 29
 - constraints, 28–29
 - copying, 61
 - double diamond, 38–39
 - examples
 - remote control, 5–7
 - voice-mail system, 3–4
 - feedback, 27–28
 - gesture-based systems, 231
 - graphical interfaces, 202
 - icons, 208–209
 - inclusiveness and, 17–18
 - menus, 205
 - mobile interfaces, 221–222
 - multimedia interfaces, 211–212
 - multimodal interfaces, 233, 235
 - participatory, 453–454
 - persuasion and, 24–25
 - practices, interaction design and, 10
 - questionnaires, analysis and, 313–314
 - RITW (research in the wild), 54–55
 - touchscreens, 228
 - users and, 9
 - UX design, 471
 - visibility, 26–27
 - VR interfaces, 215–216
 - VUI (voice user interface), 226
 - wearables, 247
 - website design, 218
 - what to design, 7–9
- design concept, 75–76
- Design Council (United Kingdom), 38
- design fiction, 411
- design maps, 451–452
- design patterns, 484
 - anti-patterns, 488
 - dark pattern, 488
 - pattern language, 484–486
 - Swiss Army Knife Navigation, 485
- design probes, 398
- design space
 - common ground, 74
 - open-mindedness, 74
 - orientation, 74
- design thinking, 456–457
- Designer's Model, 93
- desktop, graphical, 76
- dialog boxes, 199
- diaries, requirements activities, 397
- DiCoT (distributed cognition of teamwork), 92, 339–341
- different-participant design, 534–535
- digital forgetting, 117
- digital ink, 226–227
- direct manipulation framework, 85–86
- directional edges, social networks, 360
- discourse analysis, 329, 331–332
- DisplayR, 359
- distributed cognition, 92, 127–129
- documentation, AgileUX, 481–484
- documents, metadata, adding, 113
- double diamond of design, 38–39
 - evolutionary production, 51
- drones, 248–249
- DSDM (Dynamic Systems Development Method), 64
- dynamic signs, 563
- dyslexia, 120
- Echo, 184–185
- ecological validity, 515
- economic principle of prototyping, 430
- edges, social networks, 360
- effectiveness, usability and, 19
- efficiency, usability and, 20
- email, 140
- embodied interaction, 131–132

- emojis, 172–173
- emoticons, 138
- emotional AI, 166, 179–182
 - Affdex, 179
 - Affectiva, 181
 - affective expression in speech, 179
 - biometric data, 181
 - biosensors, 179
 - body movement sensors, 179
 - cameras, 179
 - facial expressions, 179, 180
 - gesture sensors, 179
- emotional design, 166
 - expressive interfaces and, 172–174
- emotional interaction, 165–166
 - advertising, 168
 - affective computing, 166, 179–182
 - annoying interfaces, 174–179
 - anthropomorphism, 187–190
 - behavioral change, 182–186
 - emotional AI, 166, 179–182
 - emotional design, 166, 172–174
 - expressions, 169
 - expressive interfaces, 172–174
 - persuasive technologies, 182–186
- emotional skills, 169
- emotions
 - affect, 169
 - automatic, 169
 - conscious emotions, 169
 - driving behavior, 169
 - stage fright, 169
 - UX and, 166–172
- empirical measurement, HCI and, 49
- engineering *versus* prototyping, 434
- Enguino, 458
- environmental concerns, 185
- environmental requirements, 392
- ERSs (error reporting systems), 47
- ESM (experience sampling method), 299
- essential use cases, 415
- ethical design concerns, 375–383
- ethnographies, 292–294
 - online, 294–295
 - requirements activities, 398
 - Royal Highland Show case study, 510–513
- evaluation, 50, 496
 - analytical evaluation, 515
 - analytics, 506, 514, 549
 - bias, 515, 519
 - case studies
 - computer game investigation, 507–509
 - ethnographic data at Royal Highland Show, 510–513
 - results, 514–515
 - cognitive walk-through, 506
 - controlled experiment, 515
 - controlled settings directly involving users, 500, 501–504
 - usability testing, 501–504
 - crowdsourcing, 513–514, 515
 - ecological validity, 515, 518
 - expert review or crit, 515
 - field study, 515
 - Fitt's law*, 506, 550, 576
 - formative, 499, 515
 - heuristic, 505–506, 515, 549
 - aesthetic and minimalist, 551
 - consistency and standards, 551
 - design guidelines and, 559–561
 - documentation, 551
 - error messages, 551
 - error prevention, 551
 - flexibility and efficiency, 551
 - golden rules, 559–561
 - performing, 554–559
 - principles and, 559–561
 - recognition *versus* recall, 551
 - system/real world match, 550
 - system status visibility, 550
 - user control and freedom, 550
 - websites, 553–554
 - in-the-wild study, 515
 - informed consent, 515–517
 - learning analytics, 506
 - living labs, 503–504, 515
 - location, 498–499
 - methods, 496
 - combining, 506–507
 - influences, 517–519
 - natural settings involving users, 500, 504–505
 - opportunistic, 507
 - participant information, 516–517
 - participants, 515
 - predictive evaluation, 515
 - predictive modeling, 549–550
 - reasons, 496–497
 - reliability, 515, 518
 - scope, 515, 519
 - settings not involving users, 500, 505–506
 - subject/topic, 497–498
 - summative, 500, 515
 - timing, 499–500
 - usability lab, 515

- usability testing, 501–504
- user studies, 515
- users, 515
- validity, 515, 518
- walk-throughs, 549
 - cognitive, 561–566
 - pluralistic, 566–567
- web analytics, 506
- evaluation studies
 - experiments, 523
 - experimental design, 534–535
 - hypotheses testing, 533–534
 - participants needed, 543
 - statistics, 536
 - field studies, 523, 536–537
 - in-the-wild studies, 538–541
 - participants needed, 543
 - usability testing, 523
 - iPad case study, 528–533
 - labs and equipment, 525–528
 - methods, 524–525
 - participants needed, 542
 - tasks, 524–525
 - users, 524–525
- evolutionary production, 51
- evolutionary prototyping, 433
- Excel, 314
- exoskeletons, 232
 - clothing, 245
- expanding menus, 203–204
- expectation management, 44
- experience maps, 453
- experiential cognition, 102
- experimental design
 - counterbalancing, 535
 - different-participant design, 534–535
 - matched-participant, 535
 - order effect, 535
 - pair-wise, 535
 - same-participant design, 535
 - within-subjects design, 535
- experiments, 503, 523
 - experimental design
 - counterbalancing, 535
 - different-participant design, 534–535
 - matched-participant, 535
 - order effect, 535
 - pair-wise, 535
 - same-participant design, 535
 - within-subjects design, 535
 - hypotheses testing
 - alternative hypothesis, 533
 - null hypothesis, 533
 - one-tailed hypothesis, 533
 - two-tailed hypothesis, 533
 - variables, 533
 - statistics, 536
 - t*-tests, 536
- expert review or crit, 515
- exploring interaction type, 81, 86–87
- expressions, emotions, 169
- expressive interfaces, 172–174
 - animated icons, 173
 - emoticons, 172–173
 - Nest thermostat, 174
 - sonifications, 173
 - vibrotactile feedback, 173
- external cognition
 - annotating, 130–131
 - cognitive tracing, 130–131
 - computational offloading, 130
 - memory load reduction, externaliza-
tion, 129–130
- externalization, memory load reduction,
129–130
- eXtreme Programming, 64, 65
- face-to-face conersations, 139–142
 - conversation breakdown, 142
 - conversational mechanisms, 140
 - rules for, 140–141
- Facebook, 145–146
- facial coding, 179
- facial expressions
 - automated facial coding, 179
 - biosensors, 179
 - cameras, 179
- false alarms, 554–555
- fast thinking, 102
- FATE (fairness, accountability, transparency,
explainability), 379, 380–382
- feedback, 27–28
- field observation, 288
 - conducting, 291
 - degree of participation
 - participant observer, 289
 - passive observer, 289
 - ethnographies, 292–295
 - frameworks, 289
 - planning, 291
- field studies, 515, 523, 536–537
 - in-the-wild studies, 538–541
 - Painpad, 538–540
- filtering, memory and, 111
- financial transactions, 113–114
 - MFA (multifactor authentication), 116

- Fitbit, 184
 - fitness trackers, 245
 - flat menus, 203
 - flow, 23
 - focus groups, 268, 271–272
 - formative evaluations, 499, 515
 - frameworks, 88–89, 92–93
 - Designer’s Model, 93
 - System Image, 93
 - User’s Model, 93
 - functional requirements, 390, 392
 - fundamental prototyping, 429
 - futuristic scenarios, 410

 - games, idle games, 336
 - gathering data. *See* data gathering
 - GBIF (Global Biodiversity Information Facility), 355
 - GDPR (general data protection regime), 379, 516
 - gesture-based systems, 229–231
 - GitHub, 489–490
 - Google Analytics, 568–571
 - air quality monitoring, 573–574
 - Google Design Sprints, 51–52, 52–54
 - Google Glass, 245–246
 - Google Lens, 88
 - Google Now card, 79
 - Google Sheets, 314
 - Google Trends, 352
 - Google Ventures, 52–53
 - Gorman, Mary, 390
 - GoToMeeting, 267
 - graphical desktop, 76
 - graphical interfaces. *See* GUIs (graphical user interfaces)
 - grounded theory, 329, 334–338
 - axial coding, 334, 337–338
 - categories, 334
 - comparisons, 335
 - data gathering and, 334
 - development, 334
 - diagrams, 335
 - example, 336–338
 - idle games and, 336
 - memos, 335
 - open coding, 334, 337–338
 - questioning and, 335
 - selective coding, 334, 337–338
 - thematic analysis, 335–336
 - word, phrase, sentence analysis, 335
 - GUIs (graphical user interfaces), 82, 89, 194
 - icons, 197, 205–208
 - research and design, 208–209
 - menus, 197
 - cascading menus, 203
 - collapsible, 204
 - contextual menus, 204
 - expanding menus, 203–204
 - flat menus, 203
 - mega menus, 203
 - research and design, 205
 - pointing devices, 197
 - web forms, 200–201
 - WIMP, 197–198
 - windows, 197, 198
 - activation time, 202
 - dialog boxes, 199
 - highlights, 199
 - management, 202
 - research and design, 202
- gulfs of evaluation, 125–126
- gulfs of execution, 125–126
-
- hands-on simulations, 209–210
- HAPIfork, 183–184
- haptic interfaces
 - exoskeletons, 232
 - MusicJacket, 232
 - ultrahaptics, 231
 - vibrotactile feedback, 231
- HBI (human-building interaction), 252
- HCI (human-computer interaction)
 - empirical measurement, 48, 49
 - focus on users and tasks, 48
 - frameworks, 92–93
 - versus* ID (interaction design), 10–11
 - interfaces videos, 195
 - iterative design, 48, 49
 - lifecycle models and, 51
 - paradigm shift, 89
 - sustainable HCI, 185
 - visions, 89
- hedonic heuristic, 559
- heuristic evaluation, 515
 - briefing sessions, 554
 - debriefing sessions, 554
 - design guidelines and, 559–561
 - evaluation period, 554
 - golden rules, 559–561
 - hedonic, 559
 - principles and, 559–561
 - websites, 553–554
- heuristics, 505–506
 - decision making, 121
- Heuristics* (Nielsen), 559
- high-fidelity prototypes, 428–429, 431

- highlights, graphical interfaces, 199
- HITs (human intelligence tasks), 513–514
- horizontal prototyping, 433
- human mirroring, telepresence robots, 149
- hypotheses testing
 - alternative hypothesis, 533
 - null hypothesis, 533
 - one-tailed hypothesis, 533
 - two-tailed hypothesis, 533
- variables
 - dependent, 533
 - independent, 533
- icons, 205–208
 - research and design, 208–209
- IDE (integrated development environment), 458
- IDEO, 13
 - TechBox, 60
- idle games, 336
- in situ studies, RITW (research in the wild), 54–55
- in-the-wild studies, 515, 538–541
- iNaturalist.org, 355–356
- incentives for data gathering participants, 264
- incidental learning, 119
- inclusiveness in design, 17–18
- index cards, prototyping and, 427
- inductive approach to qualitative data, 321
- information
 - versus* data, 264–265
 - structuring, 104
 - web pages, 110
- information processing, 126–127
- informed consent form, 515, 516–517
- informed consent from data gathering participants, 262–264
- inspections, heuristic evaluation, 550–561
- instructing interaction type, 81, 82–83
- intelligent interfaces, 194
- intentional learning, 119
- interaction analysis, 329, 333–334
- interaction design
 - academic disciplines, 10
 - alternative designs, 50
 - approaches, 43
 - consultancies, 13
 - design practices, 10
 - double diamond, 38–39
 - evaluation, 50
 - versus* HCI, 10–11
 - interdisciplinary fields, 10
 - lenses, 9–10
 - methods, 9–10
 - overview, 1–2
 - philosophies, 9–10
 - prototyping, 50
 - requirements discovery, 50
 - team makeup, 11–12
 - tools, 491–492
 - user experiences, 9
- interaction types, 81–82, 442
 - conversing, 81, 83–85
 - exploring, 81, 86–87
 - instructing, 81, 82–83
 - manipulating, 81, 85–86
 - responding, 81, 87–88
- interactive apps, learning and, 120
- interactive books, learning and, 120
- interactive products, 1–2, 36
 - marble answering machine, 4
 - what do design, 7–9
- interactive tabletops, 235–236
- interdisciplinary fields, interaction design and, 10
- interface metaphors, 78–80, 439–442
- interfaces
 - adaptive, 194
 - affective, 194
 - ambient, 194
 - annoying, 174–179
 - appliances, 222–223
 - augmented reality (AR)
 - decision making, 241–242
 - history, 243
 - HUDs (head-up displays), 242
 - mapping, 242–243
 - mixed reality, 241
 - research and design, 244
 - virtual content, 243
 - brain-computer interfaces (BCI), 250
 - command, 194, 195–196
 - research and design, 196
 - console, 247
 - drones, 248–249
 - gesture-based systems, 229–231
 - graphical (GUIs), 194, 197–198
 - icons, 197, 205–208
 - menus, 197, 203–205
 - pointing devices, 197
 - research and design, 202
 - web forms, 200–201
 - WIMP, 197–198
 - windows, 198–199
- haptic
 - exoskeletons, 232
 - ultrahaptics, 231
 - vibrotactile feedback, 231

- HCI videos, 195
- intelligent, 194
- introduction, 193–194
- invisible, 194
- mobile, 194
- mobile devices, 219–222
 - research and design, 221–222
- multimedia, 194, 209–212
 - development reasons, 210–211
 - hands-on simulations, 209–210
 - research and design, 211–212
- multimodal, 194, 232–233
 - research and design, 233, 235
- natural, 194
- NUI (natural user interface), 252–253
- pen-based devices, 226–228
- robots, 247–248
- selecting, 254
- shareable, 235, 443
 - interactive tabletops, 235–236
 - research and design, 237
 - shared editing tools, 237
 - SmartBoards, 235, 236
- smart, 194, 251–252
- speech, 194
- tangible, 194, 238, 443
 - research and design, 241
 - Resnick, Mitchel, 238
 - RFID tags, 238
 - tangible computing, 238–241
 - Tinkersheets, 238
 - VoxBox, 240
 - Zuckerman, Oren, 238
- touchless, 194
- touchscreens, 228–229
- virtual reality, 443
- VR (virtual reality), 212, 213, 214, 215–216
- VUI (voice user interface)
 - barge in, 225
 - dialogues, 225
 - machine learning algorithms, 224
 - phone apps, 225
 - research and design, 226
 - routing, 224
 - speech-to-text systems, 224
 - voice assistants, 225
- wearables, 245–247
- website design, 216
 - breadcrumb navigation, 217
 - languages, 217
 - mobile devices, 217
 - research and design, 218
- interviews for data gathering
 - alternative types, 277
 - analysis, 311
 - Cisco WebEx, 277
 - experience enrichment, 277–278
 - focus groups, 271–272
 - questions
 - checkboxes, 275–276
 - closed, 272
 - guidelines, 273
 - open, 272
 - requirements activities, 397
 - retrospective interviews, 277
 - running, 276
 - semi-structured, 269–271
 - Skype, 277
 - structured, 269
 - unstructured interviews, 268–269
 - Zoom, 277
- invisible interfaces, 194
- IoT (Internet of Things), 8, 89
- iPad usability testing, 528–533
- iPods, 14–15
- IRBs (institutional review boards), 516
- ISO (International Standards Organization), 500
- iterative design, HCI and, 49

- Jira, 388

- Kanban, 64
- Kinect (Microsoft), 230
 - SDK, 463
- KineticDress, 245
- KISSmetrics, 572
- Knowledge Navigator (Apple), 90

- lab-in-a-box, 526
- lab-in-a-suitcase, 526
- labs, usability testing and, 525–528
- Lean UX, 476–477, 480
- learnability, usability and, 20–21
- learning
 - decision-making, 121–122
 - design implications, 119–120
 - dyslexia, 120
 - incidental learning, 119
 - intentional learning, 119
 - interactive apps, 120
 - interactive books, 120
 - listening, 120–121
 - natural-language interfaces, 120
 - planning, 121–122
 - problem-solving, 121–122
 - reading, 120–121
 - reasoning, 121–122

- speaking, 120–121
- speech recognition systems, 120
- tactile interfaces, 120
- learning analytics, 506
- lifecycle, evaluation, 499–500
- lifecycle model, 51–52
- LilyPad, 458
- listening
 - design implications, 121
 - learning and, 120–121
 - permanence, 120
- living labs, 503–504, 515
- low-fidelity prototypes, 426, 431
 - card-based, 449–457
 - index cards, 427
 - Rainbow system, 426
 - sketching, 427
 - storyboarding, 426, 447–449
 - Wizard of Oz, 428
- Mac icons, 172–173
- machine learning, VUI (voice user interface), 224
- Madgex, 13
- maker movement, 461–462
- Manifesto for Agile Software Development*, 64–65
- manipulating interaction type, 81, 85–86
- manners, voice assistants and, 178–179
- marble answering machine, 4
- matched-participant design, 535
- Mechanical Turk, 513
- medicine, gesture-based systems, 230
- mega menus, 203
- memorability, usability and, 21–22
- memory, 111
 - Alzheimer's disease, 117
 - amnesia, 117–118
 - design implications, 118
 - digital forgetting, 117
 - filtering, 111
 - load reduction, externalization, 129–130
 - long-term, triggering, 118
 - memory load, 113–118
 - PIM (personal information management), 112–113
 - seven chunks of information, 114
 - short-term, 114–115
 - smartphones, 112
- mental models, 123–125
 - erroneous, 124
 - UX (user experience), 124
- menus
 - cascading, 203
 - collapsible, 204
 - contextual menus, 204
 - expanding, 203–204
 - flat, 203
 - mega, 203
 - research and design, 205
- MERboard, 292–293
- metadata, documents, adding, 113
- metalinguistic signs, 563
- metaphors, 78
 - interface metaphors, 78–80
 - popularity, 80
- MFA (multifactor authentication), 116
- micro-interactions, 23
- Microsoft, Kinect, 230
- Miller, George, 114–115
- miscommunication, requirements and, 387
- mobile devices
 - QR readers, 220
 - smartphones, 219–220
 - tablets, 219
- mobile interfaces, 194, 217, 219–222
 - research and design, 221–222
- mobile phone, using while driving, 107–108
- models, 88, 93
 - conceptual models, 92
- MonkeyLearn, 359
- MOOCs (massive open online courses), 506, 568
- Mox Analytics, 572
- multimedia interfaces, 209
 - development reasons, 210–211
 - hands-on simulations, 209–210
 - research and design, 211–212
- multimodal interfaces, 194, 232–233
 - research and design, 233, 235
- multitasking, 105–109
 - phone use while driving, 107–108
- multitouch surfaces, 228
- MusicJacket, 232
- narratives, presentations and, 344–345
- natural interfaces, 194
- natural-language interfaces, learning and, 120
- Nest thermostat, 174
- Nestor Navigator, 324
- .NET Gadgeteer, 459
- NIST (National Institute of Standards and Technology), 500
- NN/g (Nielsen/Norman) Usability Consulting Group, 501
- nodes, social networks, 360
- NodeXL, 360
- nonfunctional requirements, 390
- normal course, use cases, 415

- notes, data recording and, 266–267
- NUI (natural user interface), 252–253
- null hypothesis, 533
- Nvivo, 341
- observation, data gathering
 - activity tracking
 - data scraping, 299–300
 - diaries, 298–299
 - interaction logs, 299–300
 - web analytics, 299–300
 - analysis, 311
 - controlled environments, 295
 - think-aloud technique, 296–297
 - ESM (experience sampling method), 299
 - in the field, 288
 - conducting, 291
 - degree of participation, 290
 - ethnographies, 292–295
 - frameworks, 289
 - planning, 291
 - requirements activities, 396
 - stopping, 291
- OERs (open educational resources), 506, 568
- OFE (online feedback exchange) systems, 46
- older uses, 16
- OmniGraffle, 491
- one-tailed hypothesis, 533
- online crashing analysis, 47
- open coding, grounded theory, 334, 337–338
- Open Data Institute, data ethics canvas, 379
- open-mindedness, design space and, 74
- open source software, 489–490
- opportunistic evaluation, 507
- order effect, experiments, 535
- organizational environment, requirements and, 392
- orientation, design space and, 74
- ownership, users and, 44–45
- pair-wise design, 535
- PalmPilot, 422
- paradigms, 88, 89–90, 93
- PARCTAB system, 237
- participants, 515, 516–517
 - vulnerability, 517
- participants in data gathering
 - incentives, 264
 - informed consent, 262–264
 - number of, 261–262
 - recording, 267
 - study population, 261
- participatory design, 46, 453–454
- patents, 61
- pattern language, 484–486
- patterns, 76
- pen-based devices, 226–228
- percentages, quantitative analysis, 311–312
- perception
 - design implications, 109
 - proprioceptors, 109
- permanent impairment, 17
- personal data, 382–383
 - collecting, 353
- personas, 403–404, 414
 - persona-driven development, 405–406
 - primary, 404
 - styles, 404
- persuasive design, 182
- persuasive technologies, 182–186
- phishing scams, 186
- photos, data recording and, 266–267
- physical computing, 458–462
- physical environment, requirements and, 392
- physical impairment, 17
- Physikit project kit, 357
- pilot studies, data gathering, 264
- PIM (personal information management), 113
- planning
 - design implications, 122
 - learning and, 121–122
- pluralistic walk-throughs, 566
- Pokéman Pikachu device, 182
- poor design, voice-mail system, 3–4
- predictive evaluation, 515
 - Fitt's law*, 576–577
- predictive models, 576–577
- Presence Project, 398
- presentation of findings
 - storytelling, 344–345
 - structured notations, 342–344
 - summarizing findings, 345–346
- probing during interviews, 269
 - requirements, 398–399
- problem-solving
 - design implications, 122
 - learning and, 121–122
- problem space, 41–43
- process model, 51–52
- product owner, 44
- proof of concept, 69
- proprioceptors, 109
- protocols, 311
- prototyping, 50, 62, 421–422
 - anatomy of prototypes, 430
 - compromises, 429–434
 - 3D printing, 423

- economic principle, 430
- versus* engineering, 434
- evolutionary, 433
- filters, 429–430
- fundamental, 429
- horizontal, 433
- manifestations, 429–430
- prototypes
 - card-based, 449–457
 - compromises, 429–434
 - description, 422–424
 - generating, 447–57
 - high-fidelity, 428–429
 - low-fidelity, 426–428, 447–457
- reasons for, 424–425
- throwaway, 433
- user empathy, 437–439
- vertical, 433
- provocative probes, 399
- QA (quantified-self) movement, data
 - collection and, 353
- QR readers, 220
- qualitative analysis, 309
 - categorizing data, 324–327
 - critical incident analysis, 320, 327–328
 - thematic analysis, 322–324
 - affinity diagrams, 323
- qualitative data, 308
 - analysis
 - deductive approach, 321
 - inductive approach, 321
 - numbers, 309
- quantitative analysis, 309, 311–320
- quantitative data, 308
- questionnaires, 278–279
 - administering, 283–285
 - computer game investigation case study, 508
 - data analysis, 311
 - design, 279
 - analysis and, 313–314
 - online *versus* paper and pencil, 286
 - question format, 280–283
 - response format, 280–283
 - check boxes, 280
 - Likert scales, 280–281
 - ranges, 280
 - rating scales, 280, 282
 - semantic differential scales, 281
 - structure, 279
- Raspberry Pi, 459
- reading
 - design implications, 121
 - learning and, 120–121
- reasoning
 - design implications, 122
 - learning and, 121–122
- reflective cognition, 102
- reflective level of the brain, 171
- reliability, evaluation and, 515
- remote control design, 5–7
- remote conversations
 - Facebook, 145–146
 - human mirroring, 149
 - social presence, 146
 - telepresence, 144–145
 - robots, 146–147
 - videoconferencing, 143–144
 - VideoWindow, 143
- requirements, 385–386
 - atomic requirements shell, 388
 - brainstorming, innovation and, 402–403
 - capturing, 386
 - contextual inquiry, 400
 - apprenticeship model, 400
 - context principle, 400
 - cool concepts, 401
 - focus, 401
 - interpretation, 400
 - joy of life concepts, 401
 - joy of use concepts, 401
 - partnership principle, 400
 - data, 392
 - data gathering, 395–398
 - discovering, 50
 - environmental, 392
 - organizational environment, 392
 - physical environment, 392
 - social environment, 392
 - technical environment, 392
 - functional, 390, 392
 - miscommunication and, 387
 - nonfunctional, 390
 - overview, 387–388
 - personas, 403–408, 414
- probes
 - cultural probes, 398
 - design probes, 398
 - provocative probes, 399
 - technology probes, 399
- product dimensions, 390–391
- scenarios, 408–414
- security, 393
- usability engineers, 393
- usability goals, 392–393
- user characteristics, 392

- user experience goals, 392–393
- user stories, 388–390
- Volere framework, 388
- requirements activity
 - data gathering, 396–398
 - observation, 396
 - purpose, 386
- research
 - appliances, 222–223
 - augmented reality (AR) interfaces, 244
 - brain-computer interfaces (BCI), 250
 - command interfaces, 196
 - gesture-based systems, 231
 - graphical interfaces, 202
 - icons, 208–209
 - menus, 205
 - mobile interfaces, 221–222
 - multimedia interfaces, 211–212
 - multimodal interfaces, 233, 235
 - touchscreens, 228
 - VR interfaces, 215–216
 - VUI (voice user interface), 226
 - wearables, 247
 - website design, 218
- research and design, shareable interfaces, 237
- responding interaction type, 81, 87–88
- retrospective interviews, 277
- RITW (research in the wild), 54
- robots, 248
- rules for conversation, 140–141
- safety, usability and, 20
- same-participant design, 535
- sampling, data gathering, 261
 - convenience sampling, 261
 - saturation sampling, 261
 - snowball sampling, 261
- SAS (Statistical Analysis Software), 341
- saturation sampling, data gathering, 261
- scams, phishing, 186
- scenarios, 408–414
- science fiction as inspiration, 91
- scope, evaluation and, 515
- scraping data, 352
- Scrum, 64
- SDKs (software development kits), 463–464
- search engines, 78
- Second Life, 196
- second source data, 352
- security, 116
 - financial transactions, 116
 - MFA (multifactor authentication), 116
 - requirements and, 393
- selective coding, grounded theory, 334, 337–338
- self-tracking, personal data collection and, 353
- semi-structured interviews for data gathering, 269–271
- semiotics, 562–564
- Senseboard, 459
- SenseCam, 117–118
- sensing data, 356–357
 - visualization, 374–375
- sensory impairment, 17
- sentiment analysis, 333, 358
- shareable interfaces, 235
 - interactive tabletops, 235–236
 - research and design, 237
 - shared editing tools, 237
 - SmartBoards, 235, 236
 - social translucence, 155
- shared editing tools, 237
- short-term memory, 114–115
- SigniFYing Message, 563
- Siri, 84
- situational impairment, 17
- Sketch, 491
- sketching, prototyping, 427
- Skype, 260
- slow design, 483
- slow thinking, 102
- smart buildings, 89
- Smart Citizen, 356
- smart interfaces, 194, 251–252
- Smart Living Lab, 504
- smart TVs, interaction, 6
- SmartBoards, 235, 236
- smartphones, 219–220
 - memory, 112
 - video data recording, 267
- SNA (social network analysis), 359–363
- snowball sampling, data gathering, 261
- social computing, 136
 - emoticons, 138
 - Facebook, 145–146
- social engagement, 158–161
- social environment, requirements and, 392
- social interaction, 135–136
 - being social, 136–138
 - co-presence
 - awareness, 150–151
 - physical coordination, 150
 - shareable interfaces, 152–157
 - email, 140
 - face-to-face conversations, 139–142
 - human mirroring, 149
 - remote conversations, 143–149

- social engagement, 158–161
- social media, 136
- social networks, 360
- social norms, 185
- social translucence, 155
- socially acceptable use of data, 378
- speaking
 - design implications, 121
 - learning and, 120–121
- speech interfaces, 194
- speech recognition systems, learning and, 120
- speech-to-text systems, 224
- spreadsheets, 76
 - data analysis, 314
- SPSS (Statistical Package for the Social Sciences), 341
- SSCW (streamlined cognitive walk-through), 566
- stage fright, 169
- stakeholders, 56
- StakeNet, 56
- StakeRare, 56
- Star Trek*, holodeck, 91
- static signs, 563
- statistics, 536
- storyboarding, 426, 447–449
- storytelling, presentations and, 344–345
- streamers, biometric data, 181
- structured interviews for data gathering, 269
- structured notations, 342–344
- STS (socio-technical systems) theory, 338–339
- StudentLife, 364–366
- Subrayaman, Ramanath, 45
- summarizing findings, 345–346
- summative evaluation, 500, 515
- Superflux, 94
- sustainable HCI, 185
- Swiss Army Knife Navigation, 485
- System Image, 93
- systems-based, 329
- systems-based analysis, 338–341
 - DiCoT (distributed cognition of team-work), 339–341
 - STS (socio-technical systems) theory, 338–339
- t*-tests, 536
- tablets, 219
- tactile interfaces, learning and, 120
- tangible interfaces, 194, 238
 - research and design, 241
 - Resnick, Mitchel, 238
 - RFID tags, 238
 - tangible computing
 - Code Juniper, 239
 - littleBits, 238
 - MagicCubes, 238–239
 - MicroBit, 238
 - Torino, 239
 - Tinkersheets, 238
 - VoxBox, 240
 - Zuckerman, Oren, 238
- technical debt, 472–473
- technical environment, requirements and, 392
- technology, RITW (research in the wild), 54–55
- Technology as Experience* (McCarthy and Wright), 15
- technology probes, 399
- Teixeira, Carlos Robert, 331–332
- telepresence robots, 146–147
- telepresence rooms, 144–145
- temporary impairment, 17
- TextSL interface, 196
- Thackara, John, 9
- thematic analysis, 322–324
 - affinity diagrams, 323
 - grounded theory, 335–336
- theories, 88, 92, 93
- theory, RITW (research in the wild), 54–55
- think-aloud protocol, 296–297, 324
 - requirements activities, 397
- Third Age suite, 438–439
- throwaway prototyping, 433
- Tidy Street project, 185
- tinkering, 429
- TiVo remote control design, 5–7
- Tobii Glasses Mobile Eye Tracking glasses, 527
- tools
 - data analysis, 341
 - interaction design, 491–492
- Top Trumps, 398–399
- touchless interfaces, 194
- touchscreens
 - multitouch surfaces, 228
 - research and design, 229
 - single, 228
- transparency, AI and, 91
- triangulation, data gathering and
 - investigator triangulation, 264
 - methodological triangulation, 264
 - triangulation of data, 264
 - triangulation of theories, 264
- TruSocialMetrics, 572
- trustworthiness, 378
- Twitter, digital volunteers, 158
- two-tailed hypothesis, 533
- ubiquitous computing devices, 89
- UI (user interface), design, 9
- ultrahaptics, 231

- Uninvited Guests*, 94
- unstructured interviews for data
 - gathering, 268–269
- usability
 - effectiveness, 19
 - efficiency, 20
 - learnability, 20–21
 - memorability, 21–22
 - safety, 20
 - utility, 20
- usability lab, 515
- usability specification, 501
- usability testing, 501–504, 523
 - case studies, iPad, 528–533
 - labs and equipment, 525–528
 - methods, 524–525
 - tasks, 524–525
 - users, 524–525
- use cases, 415
 - alternative courses, 415
 - essential use cases, 415
 - normal course, 415
 - user story, 415
- user-centered approach, 47–49
- user involvement, 43–44
 - cooperative design, 46
 - crowdsourcing and, 46
 - degrees, 45–46
 - participatory design, 46
 - post-product release, 46–47
- user research, 473–477
- user stories
 - requirements and, 388–390
 - use cases, 415
- user studies, 515
- users. *See also* UX (user experience)
 - adults *versus* children, 16
 - behavior, 48
 - characteristics, 48
 - consulting, 48
 - context of use, 48
 - cultural differences, 16
 - design and, 9
 - empathy, 437–439
 - evaluation and, 515
 - expectation management, 44
 - focus on, 48
 - goals, 48
 - identifying, 55–56
 - interaction design, 9
 - needs, 57
 - older persons, 16
 - ownership, 44–45
 - product owner, 44
 - Screen Checkers, 56
 - stakeholders, 56
 - tasks, 48
 - usability testing, 524–525
 - vulnerability, 517
 - Young Parents, 56
- User's Model, 93
- utility, usability and, 20
- UX design, 471
- UX designers, 471
- UX (user experience), 2. *See also* users
 - aesthetics, 15
 - citizen science and, 355–356
 - consultancies, 13
 - content, 15
 - desirable aspects, 22
 - emotional appeal, 15
 - emotions and, 166–172
 - experiential aspect, 15
 - functionality, 15
 - goals, 22–23
 - look and feel, 15
 - mental models, 124
 - micro-interactions, 23
 - overview, 13–15
 - proof of concept and, 69–70
 - undesirable aspects, 22
 - usability, 15
 - visualizing data, 367
- UXD (user experience design), 14
- V&A museum, 152
- validity, evaluation and, 515
- variables
 - dependent, 533
 - independent, 533
- vertical prototyping, 433
- vibrotactile feedback, 231
- video
 - data recording and, 267
 - interaction analysis, 333–334
- videoconferencing, 143–144
- VideoWindow, 143
- virtual worlds, command interfaces and, 196
- visceral level of the brain, 171
- visibility, 26–27, 93
- Visio, 491
- visions, 88
 - AI (artificial intelligence), 91
 - HCI, 89
 - Knowledge Navigator (Apple), 90
- visual analytics, 568
- visualizing data, 366–375

- D3 (Data-Driven Documents), 371
- dashboard, 370, 371–372
- sensing data, 374–375
- spectrograms, 368–369
- visual literacy, 366
- voice assistants, 225, 330–331
 - manners and, 178–179
- voice-mail system design, 3–4
 - marble answering machine, 4
- VoiceOver (Apple), 17
- Volere requirements framework, 388
- Volkswagen *fun theory*, 183
- VR (virtual reality) interfaces, 212
 - arcades, 213
 - CAVE (Cave Automatic Virtual Environment), 212
 - perspective, 213
 - presence, 212
 - reporting, 213
 - research and design, 215–216
 - travel, 214
 - viewpoints, 213
- VUI (voice user interface)
 - barge in, 225
 - dialogues, 225
 - machine learning algorithms, 224
 - phone apps, 225
 - research and design, 226
 - routing, 224
 - speech-to-text systems, 224
 - voice assistants, 225
- vulnerability of users, 517
- walk-throughs
 - cognitive, 561–566
 - cognitive walk-throughs, SSCW (streamlined cognitive walk-through), 566
 - pluralistic, 566–567
- Wall Walk, 401, 434
- WCAG (Web Content Accessibility Guidelines), 446, 500, 557
- wearables, 245–247
- web analytics, 506
 - ClickTale, 572
 - Clicky, 572
 - Crazy Egg, 572
 - Google Analytics, 568–570
 - KISSmetrics, 572
 - Mox Analytics, 572
 - off-site, 568
 - on-site, 568
 - TruSocialMetrics, 572
 - visitor tracking, 572–573
- web forms, 200–201
- web pages, information structure, 110
- website design, 216–217
 - breadcrumb navigation, 217
 - heuristic evaluation, 553–554
 - languages, 217
 - mobile devices, 217
 - research and design, 218
- Weibo, 332
- white space, 109
- Wi-Fi, mental models, 123
- widgets, windows and, 202
- WIMP (Windows, Icons, Menus, and Pointers), 89, 197–198
- windows (GUIs), 197, 198
 - activation time, 202
 - dialog boxes, 199
 - highlights, 199
 - management, 202
 - widgets and, 202
- within-subjects design, 535
- Wizard of Oz prototyping, 428
- World Wide Web, 76
- written language permanence, 120
- Xerox, 76
- Xerox Star computer interface, 77
- Zoom, 260