



MORGAN & CLAYPOOL PUBLISHERS

The Design of Implicit Interactions

Wendy Ju

*SYNTHESIS LECTURES ON
HUMAN-CENTERED INFORMATICS*

John M. Carroll, *Series Editor*

The Design of Implicit Interactions

Synthesis Lectures on Human-Centered Informatics

Editor

John M. Carroll, *Penn State University*

Human-Centered Informatics (HCI) is the intersection of the cultural, the social, the cognitive, and the aesthetic with computing and information technology. It encompasses a huge range of issues, theories, technologies, designs, tools, environments, and human experiences in knowledge work, recreation and leisure activity, teaching and learning, and the potpourri of everyday life. The series publishes state-of-the-art syntheses, case studies, and tutorials in key areas. It shares the focus of leading international conferences in HCI.

[The Design of Implicit Interactions](#)

Wendy Ju

March 2015

[Core-Task Design: A Practice-Theory Approach to Human Factors](#)

Leena Norros, Paula Savioja, and Hanna Koskinen

March 2015

[An Anthropology of Services: Toward a Practice Approach to Designing Services](#)

Jeanette Blomberg and Chuck Darrah

February 2015

[Proxemic Interactions: From Theory to Practice](#)

Nicolai Marquardt and Saul Greenberg

February 2015

[Contextual Design: Evolved](#)

Karen Holtzblatt and Hugh Beyer

October 2014

[Constructing Knowledge Art: An Experiential Perspective on Crafting Participatory Representations](#)

Al Selvin and Simon Buckingham Shum

October 2014

Spaces of Interaction, Places for Experience

David Benyon

September 2014

Mobile Interactions in Context: A Designerly Way Toward Digital Ecology

Jesper Kjeldskov

July 2014

Working Together Apart: Collaboration over the Internet

Judith S. Olson and Gary M. Olson

November 2013

Surface Computing and Collaborative Analysis Work

Judith Brown, Jeff Wilson, Stevenson Gossage, Chris Hack, and Robert Biddle

August 2013

How We Cope with Digital Technology

Phil Turner

July 2013

Translating Euclid: Designing a Human-Centered Mathematics

Gerry Stahl

April 2013

Adaptive Interaction: A Utility Maximization Approach to Understanding Human Interaction with Technology

Stephen J. Payne and Andrew Howes

March 2013

Making Claims: Knowledge Design, Capture, and Sharing in HCI

D. Scott McCrickard

June 2012

HCI Theory: Classical, Modern, and Contemporary

Yvonne Rogers

May 2012

Activity Theory in HCI: Fundamentals and Reflections

Victor Kaptelinin and Bonnie Nardi

April 2012

Conceptual Models: Core to Good Design

Jeff Johnson and Austin Henderson

November 2011

Geographical Design: Spatial Cognition and Geographical Information Science

Stephen C. Hirtle

March 2011

User-Centered Agile Methods

Hugh Beyer

2010

Experience-Centered Design: Designers, Users, and Communities in Dialogue

Peter Wright and John McCarthy

2010

Experience Design: Technology for All the Right Reasons

Marc Hassenzahl

2010

Designing and Evaluating Usable Technology in Industrial Research: Three Case Studies

Clare-Marie Karat and John Karat

2010

Interacting with Information

Ann Blandford and Simon Attfield

2010

Designing for User Engagement: Aesthetic and Attractive User Interfaces

Alistair Sutcliffe

2009

Context-Aware Mobile Computing: Affordances of Space, Social Awareness, and Social Influence

Geri Gay

2009

Studies of Work and the Workplace in HCI: Concepts and Techniques

Graham Button and Wes Sharrock

2009

Semiotic Engineering Methods for Scientific Research in HCI

Clarisse Sieckenius de Souza and Carla Faria Leitão

2009

Common Ground in Electronically Mediated Conversation

Andrew Monk

2008

Copyright © 2015 by Morgan and Claypool

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

The Design of Implicit Interactions

Wendy Ju

www.morganclaypool.com

ISBN: 9781627052672 print

ISBN: 9781627052689 ebook

DOI 10.2200/S00619ED1V01Y201412HCI028

A Publication in the Morgan and Claypool Publishers series

SYNTHESIS LECTURES ON HUMAN-CENTERED INFORMATICS #28

Series Editor: John M. Carroll, Penn State University

Series ISSN 1946-7680 Print 1946-7699 Electronic

The Design of Implicit Interactions

Wendy Ju

Stanford University and
California College of the Arts

SYNTHESIS LECTURES ON HUMAN-CENTERED INFORMATICS #28



MORGAN & CLAYPOOL PUBLISHERS

ABSTRACT

People rely on implicit interaction in their everyday interactions with one another to exchange queries, offers, responses, and feedback without explicit communication. A look with the eyes, a wave of the hand, the lift of the door handle—small moves can do a lot to enable joint action with elegance and economy. This work puts forward a theory that these implicit patterns of interaction with one another drive our expectations of how we should interact with devices. I introduce the Implicit Interaction Framework as a tool to map out interaction trajectories, and we use these trajectories to better understand the interactions transpiring around us. By analyzing everyday implicit interactions for patterns and tactics, designers of interactive devices can better understand how to design interactions that work or to remedy interactions that fail.

This book looks at the “smart,” “automatic,” and “interactive” devices that increasingly permeate our everyday lives—doors, switches, whiteboards—and provides a close reading of how we interact with them. These vignettes add to the growing body of research targeted at teasing out the factors at play in our interactions. I take a look at current research, which indicates that our reactions to interactions are social, even if the entities we are interacting with are not human. These research insights are applied to allow us to refine and improve interactive devices so that they work better in the context of our day-to-day lives. Finally this book looks to the future, and outlines considerations that need to be taken into account in prototyping and validating devices that employ implicit interaction.

KEYWORDS

machine, communication, technology, computers, interface, automation

Contents

Acknowledgments	xiii
1 Introduction	1
1.1 Background	3
1.2 The Theory of Implicit Interactions	3
1.2.1 Point of View	4
1.3 Book Overview	4
2 The Theory and Framework for Implicit Interaction	7
2.1 Actors in the Framework	7
2.2 Implicit Interaction Framework	9
2.2.1 Attentional Demand	9
2.2.2 Initiative	10
2.3 Interaction Paradigms	10
2.3.1 Foreground Reactive	11
2.3.2 Foreground Proactive	11
2.4 Interaction Trajectories	12
2.5 Interaction Analogues	18
2.6 Interaction Pitfalls	19
2.7 Conclusion	19
3 Opening the Door to Interaction	21
3.1 Identifying Implicit Interactions	22
3.2 Analyzing Implicit Interactions	22
3.3 Designing Implicit Interactions	25
3.4 Implicit Interaction Patterns	26
3.5 Verifying/Validating Patterns	27
3.6 Wizard-of-Oz Techniques	28
3.7 Video Prototyping Techniques	29
3.8 Findings	30
3.9 Conclusion	30
4 Light and Dark: Patterns in Interaction	31
4.1 The Difference between Pattern Diagrams vs. State Diagrams	31

4.2	Command-Based Interaction	35
4.2.1	Breakdowns Invoke the Interactive Paradigm	37
4.2.2	Common Errors: Misperception, Misjudgment, Misexecution	38
4.2.3	Repairs	39
4.2.4	Basic Patterns	40
4.3	Interactions with “Automatic Switches”	41
4.3.1	Automatic Systems Hide Problems	41
4.3.2	Using Feedforward to Show Future Actions	42
4.3.3	Testing Background Proactive Systems	42
4.4	Interactive Switches	43
4.4.1	Interaction Errors	45
4.4.2	Repairs and Overrides	46
4.4.3	Guiding Lights	47
4.5	Dark Patterns	48
4.6	Invisibility Is a Consequence and Not a Cause of Good Design	49
4.7	Conclusion	49
5	Action and Reaction: The Interaction Design Factory	51
5.1	Proxemics, or, Dancing with the Materials of Design	51
5.1.1	The Role of Proxemics in Interaction	52
5.1.2	Proxemics in Interaction Design	53
5.2	Intention	55
5.2.1	The Dark Pattern of Secret Robotic Plotting	56
5.2.2	The Role of Intention in Interaction	58
5.2.3	Intention in Interaction Design	58
5.3	Consistency	59
5.3.1	The Head-in-a-Head Problem	59
5.3.2	The Role of Consistency in Interaction	61
5.3.3	Consistency in Interaction Design	61
5.4	Conclusion	61
6	Driving into the Future, Together	63
6.1	The Rationale for Automation	63
6.2	Interaction Design On the Critical Path	64
6.3	The Right Analogue Is Key	64
6.4	Transfer of Control	65
6.5	Watch Out for People	66
6.6	Know When to Ignore the Rules	67

Bibliography	69
Author Biography	77

Acknowledgments

This book is the culmination of much research undertaken over many years by a large number of collaborators beyond myself. I would like to thank Leila Takayama, David Sirkin, Lawrence Neeley, Doug Dooley, Brian Lee, and Arna Ionescu for their contributions to the stories, ideas, and theories here. I also owe a huge debt of gratitude to my many mentors: Terry Winograd, who helped me to hone the intellectual aspects of this body of work over many meetings, over many years; Bill Verplank, whose habit of sketching his thinking I have adopted wholesale; Larry Leifer, who has provided continuous amused support and lateral thinking; also, Scott Klemmer and Don Norman, who have helped me to clarify my ideas through their patient and persistent questioning.

I would also like to acknowledge the many people who assisted me with the writing of this monograph. In particular, Corina Yen, Nikolas Martelaro, and Sarah Lewis all deserve credit for helping to transform my Ph.D. dissertation into material that is more generally palatable to read.

I wish to thank my family, particularly my parents, my husband Justin Olson, and sons Case Olson and Simon Olson for their support throughout this project.

CHAPTER 1

Introduction

Imagine, for a second, a doorman that behaves like an automatic door does. He does not acknowledge you when you pass by or approach. He gives no hint which door can or will open—until you wander within six feet of the door, whereupon he flings the door wide. If you arrive after hours, you might stand in front of the doors for a while before you understand that the doors are locked because the doorman's blank stare gives no clue.

If you met such a doorman, you might suspect psychosis. And yet this behavior is typical of our day-to-day interactions not only with automatic doors but any number of interactive devices. Our cell phones continue to ring even when we are obviously trying to stifle them. Our smartphone alarm clocks go off in the morning even if we are already awake, staring right at the screen and poking at the device. Our computers only know to ring every so many minutes before an important meeting, jerking us in and out of distraction rather than easing us out of our task and toward our next activity. The infiltration of computer technologies into everyday life has brought these interaction crises to a head, to the point that many feel it is necessary to take a digital Sabbath.

Everyday Interactives, Pervasive and Mundane

While the lion's share of attention around interaction design is given to new devices and platforms that introduce new interaction paradigms—computers, webpages, smart phones, tablets—more mundane interactive devices surround us as well in our everyday lives: vending machines, ticket turnstiles at a train station, information kiosks at a museum, GPS-enabled navigation systems, auto-flush toilets, motion-triggered room-lights, and automatic automobile door locks. More and more passive mechanical devices are making a significant transition to become electronic and computationally enhanced digital machines. Our cars and coffee machines are acquiring screens and buttons where they once had dials and knobs, embedded chips and sensors where they once had springs and gears. In this transition, there are novel possibilities for old genres of devices, but there are also new snags for their users.

While we might be willing to dedicate a lot of time to the newest and whizziest gadgets we have acquired, the demand for our attention will not scale if explicit interaction is required by every new computer-enabled device. As we design interactions for multitudes of devices distributed throughout our world, we need an approach for interaction that is different than the one used for traditional computer user interfaces. Many in the research world believe that it is possible for us to make smart machines that are context-aware and autonomous, so that they won't need the input

2 1. INTRODUCTION

of the user at all. Implicit interactions represent another view: they suggest that people are critical to interactions, but that the engagements can be designed to be far less jarring or disruptive than what we experience now.

Interactions Require Social Savvy

It is for these everyday devices and in-the-world contexts that an understanding of the rules that govern human interaction is really critical. Many have suggested that more sophisticated artificial intelligence or elaborate networks of sensors may be the solution to the problem of obnoxious and overbearing machines. However, this book considers another complementary perspective: that sensing and computation need to be augmented with an understanding of the unstated expectations we people have from our interactive counterparts. We have protocols that govern when and how we engage one another, when and how we transact, when and how we take leave, how much attention is demanded, and how much attention is given. Our lack of sophistication in designing interactive devices that do not constantly demand input or attention can be remedied through increased awareness and understanding of the implicit protocols that govern our everyday interactions with one another.

This foray into implicit interactions represents a major departure from traditional human-interaction design. Traditional human-computer interactions and human-machine interactions have focused on the realm of explicit interactions, where the use of computers and interactive products rely on the overt input and output common in command-based or graphical user interface-based interfaces. In these interfaces, the agency lies wholly with the user or machine: You use a mouse to select the file you'd like to open. You order a robot to fetch you a drink. The alarm clock tells you to wake. The GPS tells you where to go. In implicit interactions, agency and initiative are less clearly assigned. Implicit interactions are based on inputs and outputs that are negotiated and jointly performed—perhaps the computer desktop makes some file icons slightly larger and easier to click on based on what else you're working on; perhaps a robot offers you a drink because it thinks you look thirsty. No party is clearly in charge; actions do not equate to agency if they take place in the context of implied demands or offers. These types of implicit interactions are an inevitable part of what some call “smart” products, products whose actions contain some degree of agency or of activity that occurs without the explicit behest or awareness of the user.

These interactions are becoming more prevalent as interactive technologies begin to be applied to situations where people are not able to exercise full agency and control—when users are novices, or in arenas such as the automobile, the meeting room, or the home, where the person is physically, socially, or cognitively engaged.

Developing Intuitions for New Interactions

In some ways, designers of interactive devices are in new territory, and cannot know what to expect from users. In other ways, though, the answers lie within us. We humans have an abundance of experience employing implicit interactions in our day-to-day interactions with one another. Over the course of their lives, people learn behavioral patterns and codes that enable them to negotiate their interactions with one another. We often employ them without conscious thought: we modulate our speaking volume based on ambient noise level, use smaller words when explaining things to children, and hold the door open for others when we see that their arms are full. These accommodations do much to smooth our day-to-day interactions with one another and yet are made without explicit command. The success of these interactions relies less on extraordinary intelligence and more on sophisticated negotiation of changing contexts and subsequent behaviors.

1.1 BACKGROUND

In my years working at the intersection of research and design on interactive devices ranging from kitchen counters that guide people to cook to chairs that rearrange themselves when guests come over, I've often been struck by how the difference between the right interaction and the wrong one was often subtle and difficult to talk about. The more complex the devices, the more difficult the problem. Today, devices are designed by interdisciplinary teams who often do not speak the same professional language or follow a common practice. Debating, sketching, and discussing designs has grown more challenging as interaction design grows as a field because interaction designers do not have established norms or tools for conversing about interaction. It does not help that the most successful interactions are often somewhat invisible, as we cannot attribute functionality to mechanisms we do not discern.

1.2 THE THEORY OF IMPLICIT INTERACTIONS

The Implicit Interaction Framework supports *The Theory of Implicit Interactions*, which posits that people rely on *conventions of interaction* to communicate queries, offers, responses, and feedback to one another. These everyday implicit interaction patterns with one another drive our expectations of interactive device behavior. Thus, designers can learn to make better interactive devices by honing their ability to recognize and dissect day-to-day interactions to understand how they work. By understanding how people cue attention, indicate potential for action, signal agreement or refusal, interaction designers can develop an intuition for successful interaction patterns. This will help designers to create more effective and comprehensible products and services.

1.2.1 POINT OF VIEW

This book is written from a juxtaposition of design and research perspectives. It has two goals. The first is to help researchers to analyze interactions, discern patterns, and characterize user responses. At the same time, this book is intended to empower designers working on designing implicit interactions, specifically to give them a stronger intuition to draw upon as they design products and services that have an interactive or behavior component to them.

1.3 BOOK OVERVIEW

In this book, I outline the Implicit Interaction Framework as a tool, which maps the trajectories of interactions by looking at the interplay of initiative and the dynamics of attentional demand over the course of simple transactions. By mapping these activities and patterns on this framework, we designers are better able to analyze the pragmatics of interactions, to understand specific interactive practices at play, to communicate and share their intuitions with clients and collaborators, and to develop successful interaction analogs. [Chapter 2](#) explores the Implicit Interaction Framework in detail, looking at how the framework interaction essentializes interactions to make it easier to apply the method of designing through interaction analogues. In [Chapter 3](#), we closely examine interactions with doormen and automatic doors, to understand in detail how interactions are initiated, the elaborate dance and negotiation that precedes every engagement. [Chapter 4](#) looks at the transition from automation to interaction through the example of light-switching. In [Chapter 5](#), we look at an example of how implicit interactions perform in contexts where people and interactive devices are working together to get somewhere by looking at interactions with vehicles. Finally, in [Chapter 6](#), we revisit the question of how to design and evaluate interactions with the implicit interaction perspective in mind. This design-based approach has two main objectives: to be generative—that is, to guide designers in a constructive fashion in designing implicit interactions—and to be generalizable—that is, to suggest techniques and methods that are applicable to interaction designers working on a wide array of ubiquitous computing scenarios.

This book is written for designers and researchers of interactive systems. It does not assume a lot of technical knowledge about how such systems are operationalized on a technological level. Instead, I focus on aspects of timing, behavior, attention, and initiative that everyday people should understand. That said, because it is written for designers and researchers, the perspective taken throughout this work is that of the designer. When I speak of “users” or “people,” I do not mean ourselves, but rather the other people who will be interacting with the devices we create.

By shining a light on implicit interactions, my intent is to help designers notice these implicit communicative activities taking place in their everyday life, to illuminate key interaction techniques and patterns using the Implicit Interaction Framework, and to illustrate how these techniques and patterns can be used by designers to develop novel ways of engaging people with interactive devices.

By using a framework, the design process can be more systematic and less ad-hoc; designers still have important work to do, they will just be able to mine known territory instead of having to dig around in the dark. It is my hope that, when these social mechanisms that help us engage with one another in everyday life are made clearer, we can come to expect elegance and economy in our day-to-day interactions with machines.

CHAPTER 2

The Theory and Framework for Implicit Interaction

To help us better understand and design implicit interactions, I introduce the *Implicit Interaction Framework*. This framework helps us to understand the dynamics of interactions as they unfold, and to identify common patterns in interactions that can assist us when we are analyzing interactive activity or designing novel interactions. It is meant to be used in the translation between observations of human-human interactions and the design of human-device interactions.

This framework is premised on a theory that implicit interactions function through regular patterns of communication; hence, it hypothesizes that interactions designed in accordance with these patterns will be recognizable and effective. Lucy Suchman posited, in fact, that human-computer communication was a special case of human communication in which the resources available to the participants is limited (Suchman, 1987). Following Brenda Laurel's metaphor of "the computer as theatre" (Laurel, 1993), we speak of the person and the interactive systems as being actors on the stage of our framework. The interactions are communications, but they are also performances that we are watching, directing, scripting (Field, 1994), and evaluating.

2.1 ACTORS IN THE FRAMEWORK

This framework models interactions as the exchange between two entities. One of the key premises in this work is that people engage interactive objects or media *socially*, as if they were other people, even if there is no mistaking the object for a person (Reeves and Nass, 1996). In the diagrams in this chapter, we assume one entity is a person and the other is an interactive object. The interactive object is commonly a computer, robot, or interactive device, but it could also be something larger or more systemic (like the voice interface for the computer on the Starship Enterprise in *Star Trek*.) The framework provides a useful basis for understanding elementary interactions between two interactants. In everyday practice, interactions can be far more complex, involving more people or objects that work in concert or in an uncoordinated fashion.




Figure 2.1: We are interested in the interplay between what people see, think, and do in an interaction, and what objects perceive, deduce, and perform.

Bill Verplank says, “Interaction Designers answer three questions: How do you do? How do you feel? How do you know?” (Verplank, 2003) As designers of interactive systems, we also need to think about what the interactive object can perceive, deduce, and what actions it can perform. More importantly, we need to understand how what one entity sees, does, and deduces affects the other.




Figure 2.2: The Implicit Interaction Framework shows the dimensions of interactive system behavior.

2.2 IMPLICIT INTERACTION FRAMEWORK

The Implicit Interaction Framework (Figure 2.2) divides the space of possible interactions between our two actors along the axes of attentional demand and initiative. Attentional demand is the attention demanded of the user by the computer system. It is generally described by the degree of cognitive or perceptual *focus*, *concentration*, and *consciousness* required of the user. Interactions that demand the user's attention are *foreground interactions*, and interactions that evade the user's attention, and in fact, may elude notice, are *background interactions*.

Initiative encompasses who initiates an interaction and to what degree. The framework presumes the perspective of interactive system designers, so interactions initiated by the user are *reactive interactions*, and interactions initiated by the system are *proactive interactions*. By characterizing interactions in this way, we are able to generalize about the capabilities and features of whole classes of interactions in a domain-independent fashion.

Early on in human-computer interactions, researchers associated attentional demand and initiative with one another. Bill Buxton, for instance, defined attentional ground this way: “What we mean by Foreground are activities which are in the fore of human consciousness—intentional activities. Speaking on the telephone, or typing into a computer are just two examples. By Background, we mean tasks that take place in the periphery—‘behind’ those in the foreground. Examples would include being aware of someone in the next office typing, or the light in your kitchen going on automatically when you enter it, as opposed to you manually flicking the switch” (Buxton, 1995).

What the Implicit Interaction Framework helps to illustrate is that attention and initiative do not always go hand-in-hand. Buxton’s examples contrast deliberate user-initiated interactions—typing into a keyboard or switching on a light—to device-initiated interactions that users scarcely notice. However, device-initiated interactions can demand attention—a cell phone ringing, for example, gets a person’s attention even though the device initiates the interaction. Also, user-initiated interactions can occur in the attentional background—for example, if the lights go out when I leave a room.

2.2.1 ATTENTIONAL DEMAND

Attentional demand can be manipulated by adjusting the perceptual prominence of objects. This may be done through visual organization techniques, such as contrast, hierarchy, and weight (Lupton, 2004), as well as more dynamic means, such as pointing or placing (Clark, 2003). This type of visual display formatting has been most notably used to improve situation awareness in airplane cockpits (Andre et al., 1991). Interaction design research on the use of such techniques to present ambient information to users engaged in some other task has been pursued at the MIT Media Lab (Wisneski et al., 1998) and Berkeley’s Group for User Interface Research (Matthews et al., 2004)

10 2. THE THEORY AND FRAMEWORK FOR IMPLICIT INTERACTION

among others. Saskia Bakker has pointed out that tangible and handheld objects in particular can enable peripheral interactions that do not require major attentional effort (Bakker, 2010).

Another way to affect the degree of attention demanded is through *abstraction*. By combining elements into a larger whole, the user is presented with less detail. *Chunking* is an example of an abstraction technique wherein experts are able to comprehend complex situations (such as the state of a chessboard) with greater ease because they are able to parse the scene into familiar subcomponents (Chase and Simon, 1973). Gestalt laws suggest that chunking leads to an “integrating of awareness” where people are able to identify a whole (say, a particular person’s face) without being able to identify the details that make up the whole (Polanyi, 1967).

2.2.2 INITIATIVE

Initiative is salient in situations where actors work together to accomplish a task. “Every day we engage in activities in which we have to coordinate with others to succeed,” says Clark. “Face to face, we have systematic, economical and robust techniques of arranging for joint activities” (Clark, 1996). The distinction of who initiates an interaction is a critical one in human-human communication.

Device-initiated interactions (where the device does more than merely react to a user’s immediate command) have existed for quite some time—alarm clocks, for example, or tutorial systems. However, the range of device-initiated interaction has grown a great deal with the advent of inexpensive computation, sensing, and networking technology. Thus, the factor of initiative is becoming increasingly important as more and more interactive devices are able to perform autonomous action. Autonomous action enables devices to take action without engaging a user’s explicit attention.

Proactive objects operate in a realm of greater presumption, and so it is common that they need ways of seeing, discerning, and reasoning about the world (Tennenhouse, 2000). This explains why most forays into proactivity, such as the research performed at Microsoft Research (Czerwinski et al., 2000), University of Karlsruhe (Schmidt, 2000), and Georgia Tech (Salber et al., 1999), have been oriented to the technological issues of sensing, aggregating data, developing user and task models, and performing inference (Pantic and Patras, 2006). The importance of both attention and initiative in human-machine communication have been recurring themes in the research of Eric Horvitz (Horvitz, 1999; Horvitz et al., 2003), but interestingly his work does not look at these factors together as a way to map interaction.

2.3 INTERACTION PARADIGMS

Figure 2.3 shows the Implicit Interaction Framework, with descriptions of interactions typified by each quadrant.




Figure 2.3: Characteristic interactions for each quadrant of the Implicit Interaction Framework.

2.3.1 FOREGROUND REACTIVE

Interaction paradigms tend to belong primarily in one quadrant or another of the Implicit Interaction Framework. For example, “direct manipulation,” which is what we call the interaction paradigm associated with windows, icons, mouse and pointers, typically takes place in a person’s attentional foreground, and the computer reacts and responds to control maneuvers made by the person on the computer desktop. The previously dominant “command-line” paradigm was similarly foreground-reactive.

2.3.2 FOREGROUND PROACTIVE

Interactions that alert people about things they don’t know about, or which guide people through an interaction are foreground-proactive. That is to say, these are interactions where the device takes initiative in the interaction. This is common with clocks, timers, and alarms, and also is used in direction-based applications, such as automotive navigation or online tutorial interactions.

The same device might incorporate modes that lie in all four quadrants. Take, for example, digital video recorders, such as the genre-creating TIVO. DVRs can act much as a videocassette re-

12 2. THE THEORY AND FRAMEWORK FOR IMPLICIT INTERACTION

corder does, recording shows when you command (foreground reactive), or recording shows you've pre-programmed (background reactive). In addition, the device can suggest shows it thinks you might want to watch (foreground proactive) or even pre-record shows it is pretty sure you want to watch (background proactive).




Figure 2.4: The DVR's different record functions on the Implicit Interaction Framework.

2.4 INTERACTION TRAJECTORIES

Although it is useful to characterize basic interactions paradigms based on attention demanded and initiative, the most useful aspect of the Implicit Interaction Framework is its ability to map the dynamic changes in attentional demand and initiative that take place throughout an interaction. For example, let us consider how a snooze alarm works.

The alarm is initially set through some foreground-reactive task where the user picks a time in the future that the alarm will go off.




Figure 2.5: User setting alarm clock is a foreground reactive task.

For situations where the alarm is recurring, it is a good idea for the alarm clock to show some indicator that an alarm is set as feedback, both to avoid the user forgetting that an alarm is set, and to more clearly indicate when an alarm is not set. Presumably because users know when they've just set the alarm, the alarm set displays are always pretty subtle.




Figure 2.6: Next step: the clock may indicate whether the alarm is set or not.

Alarms, like all alerts, are proactive—if the intended audience for the alarm were aware enough to initiate an alarm, then the alarm would not be required. When the set time arrives, the alarm goes from proactive-background to proactive-foreground to get the user’s attention and wake them up!




Figure 2.7: Next step: alarm ringing moves the trajectory of the interaction into the foreground proactive space.

The snooze function is a foreground-reactive interaction, but it only works as a snooze if snoozing the alarm takes less attention than setting an alarm from scratch.




Figure 2.8: Next step: the user may hit the snooze button.

Readers who are not early birds are probably familiar with the snooze-alarm cycle. Particularly hardened readers might even be familiar with a common failure mode of many alarms, which are only able to perform a finite number of snooze cycles before the alarm just doesn't go off anymore.




Figure 2.9: Next step: snooze alarm cycle—replete with failure mode—diagrammed on the Implicit Interaction Framework.

As you can see, the Implicit Interaction Framework helps make clearer the whole of the interactions people have with the alarm clock—not just the moments of explicit interaction, but also the latent implicit interactions that are really critical to the function of the snooze. It can also help us to identify better alternatives to the current program.




Figure 2.10: Next step: instead of using a PROACTIVE-BACKGROUND display for the end of snooze, the alarm should use a PROACTIVE-FOREGROUND display.

2.5 INTERACTION ANALOGUES

The Implicit Interaction Framework was created to map insights gained from studying human-human interaction to the improved design of interactive devices. Thus, as important as what is in the Implicit Interaction Framework is what is *not* diagrammed. For instance, the framework doesn't indicate how the alarm is implemented. Is it a mechanical clock? Is it a digital clock? Is the clock a metaphorical stand-in for children that come in and wake their parents in the morning? These details do not necessarily change the trajectory of the interaction. Similarly, the Implicit Interaction Framework doesn't include anything about the user's conceptual model or motivations. The emphasis is on the back and forth between the interactants, and the amount of attention allocated to the interaction at each phase. A lot of details are left out.

By focusing solely on the attentional demand and initiative, it becomes easier to consider just the interactive behaviors, and thus to form *interaction analogues* between human-human interactions and human-device interactions. For example, this perspective makes it easier for a designer to observe how a skilled human navigator offers directions to help drivers find their way and map

those insights onto a how a good GPS app should perform a similar task. On some level, these behaviors reflect social protocols that ease communication, and the framework helps to highlight the functional aspects of those protocols. With the framework, it doesn't matter who is interacting, or what they are interacting about, only how.




Figure 2.11: In the Implicit Interaction Framework, no one cares if you're a dog.

2.6 INTERACTION PITFALLS

Conversely, the interaction design framework can be used to identify common pitfalls or dark patterns that interactions can fall into. Often, as designers, we can see that an interaction is failing but don't always understand why. We might not understand why people are annoyed at an interactive avatar that intermittently pops up on our screen until someone says, "Oh, it's like someone hanging up on you because that avatar never says goodbye!" We have found over and over in our research that the social protocols that govern interpersonal interaction, such as welcome and leave-taking, also apply to device interactions as well.

2.7 CONCLUSION

I have presented a framework for implicit interaction that characterizes interactions based on attentional demand and initiative—factors that are pertinent to any interaction, regardless of domain. This framework can be used by designers as a lens on their interaction design problems, and can help them leverage existing linguistic, sociological, or ethnographic techniques to the end of designing better human-computer interactions. The framework supports the *Theory of Implicit Interaction*.

tions, making it easier to see the *conventions of interaction* to communicate queries, offers, responses, and feedback to one another. This basis in convention enables people to communicate efficiently with other people even in new, unfamiliar situations because they can use the common language of implicit interaction to negotiate the interactions. Furthermore, it allows implicit interactions to be applied to the design of interactive devices to improve people's ability to "communicate" with interactive devices.

Because implicit interactions have convergent features due to the constraints imposed by the human in the loop, knowledge about the interactions can be generated and generalized—key components to any area of academic research. This portability of solutions from one domain to another also enables design solutions to be passed from one design researcher to another, enabling designers of interactive objects to develop generalized interaction patterns for different classes of interactions.

CHAPTER 3

Opening the Door to Interaction

The simple act of walking up to a new door can raise many questions:

Where is the door?

Does it open?

What way?

How quickly or slowly?

Incorporating other people into the context multiplies our uncertainty:

Who will approach the door first?

Will they hold the door for me?

Or am I to hold the door for them?

What is a socially appropriate speed with which to approach the door?

Novel interactions can be awkward, but they also make us conscious of how much coordination occurs beneath the surface in our day-to-day familiar interactions. Lucy Suchman aptly pointed out, “The practical problem with which the designer of an interactive machine must contend is how to ensure that the machine responds appropriately to the user’s actions... Every action assumes not only the intent of the actor, but the interpretive work of the other in determining its significance” (Suchman, 1987). Each volley of the interaction is built around a series of contingencies based on response and interpretation.

Interactions around doors are a great example of the issues designers face in designing interactions that are supposed to be seamless and require little attention. Because implicit interactions occur outside of the user’s notice or initiative, they can be challenging to design; they must be lightweight and low attention. It is difficult to go up and ask people about the interactions, because they often don’t notice them consciously, or even have the words to describe what they feel happened or what they think should happen. The seamlessness and invisibility of successful interactions can be difficult for the designers themselves to see all of the mechanisms at work within the interaction. We can’t improve implicit interactions through the brute force solutions of longer explanations or more training, because these solutions serve to increase the weight and attention on the interaction. Thus, it is important that the designers of implicit interactions pay greater attention to the interplays between interactants.

One important function, then, of the implicit interaction framework is that it helps us to map out how existing interactions unfold. These trajectories make explicit that which is invisible in day-to-day life. The method we promote to design implicit interactions is to identify these trajectories and then to use them as patterns for new interactions. In this chapter, we illustrate the implicit interaction framework being used to analyze interactions around doors. We will then use these trajectories as interaction patterns, first for the behavior of interactive doors, and then to the design of other devices that need to encourage people to approach and engage in order for interaction to happen. This chapter functions as a case study introducing how the implicit interaction framework can be used for analysis and design of interactions.

This use of the implicit interaction framework for patterning depends on the theory of implicit interactions; it assumes that we can look to human-human interaction patterns that will carry over to human-device interaction. Embedded in this theory is the idea that it is the conventions of communication—what message is conveyed, how and by whom—that is operative in the interaction. The framework is not concerned about the details of who is interacting.

3.1 IDENTIFYING IMPLICIT INTERACTIONS

One of the first and hardest steps in designing and analyzing implicit interactions is discerning where there is an implicit interaction in the first place. From the perspective of the implicit interaction framework, we are always looking for a dyad, a pair of interactants, who are performing some sort of joint action, some activity that requires communication or coordination to pull off. In the following example, we look at how a doorman interacts with a passerby. This is a nice example, because few people think of themselves as having interacted with a doorman unless they engage in a conversation about the weather, or the local sports team—but if they even see him, there can be an interaction, even if they do not walk into his building.

3.2 ANALYZING IMPLICIT INTERACTIONS

Here is a scene from the play of everyday urban life that looks at how a doorman interacts with passersby.

SCENE 1a:

SETTING: On a sidewalk at the entrance to a building in the middle of the block.

ROLES: Doorman, Passerby

SEQUENCE:

Doorman: (stands in front of the door, wearing a red uniform)

Passerby: (walks down street, on a path that will pass the door)

Doorman: (spots person walking down street)

Passerby: (notices doorman with red finery in front of the door, but keeps on walking)

Doorman: (puts gloved hand on door handle)

Passerby: (slows down a little, and looks into the doorway)

Doorman: (opens door slightly)

Passerby: (keeps walking past door; turns to look down street)

Doorman: (lets door shut, and takes hand away from the door handle)

In this scripted example, the doorman employs proactive, low-attention techniques to signal his capability for opening doors. He did this through overt preparation, when he put his gloved hand on the door handle, and through an enactment technique, by pulling the door open a little as a suggestion. In some ways, it can be difficult to appreciate the elegance of these interactions unless we consider the following alternatives:

SCENE 1b:

SETTING: On a sidewalk at the entrance to a building in the middle of the block.

ROLES: Doorman, Passerby

SEQUENCE:

Doorman: (stands in front of the door, wearing a red uniform)

Passerby: (walks down street, on a path that will pass the door)

Doorman: (spots person walking down street) Hey!

Passerby: (surprised, notices doorman with red finery in front of the door. He keeps on walking, but more slowly, looking around to see if the doorman is talking to him.)

Doorman: Hey, do you want to come in?

Passerby: (slows down a little, and looks into the doorway) Um, maybe?

Doorman: I'm willing to let you in! The door opens this way!

Passerby: Uh, no, I'm just going to keep going down the street here.

24 3. OPENING THE DOOR TO INTERACTION

Doorman: Okay, then I won't open the door for you.

SCENE 1c:

SETTING: On a sidewalk at the entrance to a building in the middle of the block.

ROLES: Doorman, Passerby

SEQUENCE:

Doorman: (stands in front of the door, wearing a red uniform)

Passerby: (walks down street, on a path that will pass the door)

Doorman: (keeps staring straight ahead, may or may not spot person walking down street)

Passerby: (either doesn't notice the door, or doesn't think she can go in, keeps on walking)

Doorman: (stands silently)

Passerby: (walks on by, not slowing down at all)

In Scene 1b, the doorman is far more overt. His actions are focal, and his communication style is exclusive; a person couldn't talk on a phone or interact with friends and engage with the doorman at the same time, the way she could in Scene 1a. Also, the act of refusal is more awkward in Scene 1b, as if the additional expenditure of effort and attention increases the sense of obligation to interact.

In Scene 1c, the doorman is far less overt than in 1a. If he is aware of and willing to engage with the passerby, it is not clear. The result is not actually neutral—it is unfriendly and even a bit hostile. If the passerby *did* want to enter the building, the lack of engagement or initiative on the part of the doorman might make her think twice.

Here are elements from these scenes mapped on the implicit interaction framework.




Figure 3.1: Basic interaction pattern with doorman.

3.3 DESIGNING IMPLICIT INTERACTIONS

An interaction designer whose job was to design the user experience of people entering the doorman's building might use these mapping techniques to think about what the back and forth between the doorman and the person should be. They might also think about alternatives for each step that might accomplish the same ends: what if the doorman's hands are full? What if the passerby is distracted?

Analogously, an interaction designer designing an automatic door can use the doorman pattern to motivate questions such as how the door draws attention to itself, how it communicates its role as a portal, and how it introduces its affordance. Such steps sometimes can be accomplished implicitly: the door's mere physical form serves to draw attention and communicate its "door-ness."

The designer also can look for clever ways to achieve the effects of each step: by opening a little when a person walks by, for example, the automatic door can simultaneously draw attention, define its role as a door, and introduce its ability to open automatically by softly humming in overt preparation or making a show of jiggling its handle as enactment.

3.4 IMPLICIT INTERACTION PATTERNS

The patterning of interactions helps designers to determine the roles, setting, and sequence of the interaction to be designed. This helps to situate the designed interaction so that it takes into account context and environment much in the way that we do naturally for human interaction. The interaction analogues allow the designer to imagine functionally equivalent actions, mapping the capabilities of the automatic door against the techniques employed by the doorman, without slavishly and literally replicating his actions.

This second interaction sequence is patterned after the interaction from Scene 1a, with an automatic door that mimics the doorman's implicit behaviors analogously.

SCENE 2:

SETTING: On a sidewalk at the entrance to a building in the middle of the block.

ROLES: Door, Passerby

SEQUENCE:

Door: (with sign that says "Automatic Door")

Passerby: (walks down street, on a path that will pass the door)

Door: (sensors notice motion down the street)

Passerby: (notices door frame, and keeps on walking)

Door: (makes a soft motor hum noise, as if preparing to open)

Passerby: (slows down a little and looks into the doorway)

Door: (opens a little, jiggling its handle)

Passerby: (keeps walking past door; turns to look down street)

Door: (lets door shut)

Let us compare the implicit interaction framework trajectories of Scenes 1a and Scene 2 by comparing Figures 3.1 and 3.2.




Figure 3.2: Interaction pattern with automatic door, analogous to the doorman in Figure 3.1.

Clearly drawing analogous framework patterns can be useful as a thought experiment or collaboration tool. It can help to anchor a thorough analysis of the originating interaction, or discussion about how existing interactions unfold, or how new ones would work. However, it is worth asking: Do these analogous interaction patterns work in real life?

3.5 VERIFYING/VALIDATING PATTERNS

Frameworks are simplified representations that help people understand a topic, so we can't really speak about whether a framework is true or untrue. However, when it comes to the patterns, we can check to see if interactions designed with analogous patterns generated from the implicit interaction design framework are received the way we would expect them to be.

In my research with colleagues at Stanford University, we looked at this question in part by performing in-person field experiments and online video prototypes looking at how people interpreted different automatic door gestures. By using a lever bar, a hidden person was able to move the glass building door so that it appeared to open and close automatically (see Figure 3.3b). A monitor sat in the halls, appearing to wait for someone. When people passed by, the confederate surreptitiously used a walkie talkie call button to signal to the door operator to open a door (Ju and Takayama, 2009).




Figure 3.3: Snapshots showing the setup of the Wizard-of-Oz door study from A) outside and B) inside. Note the secret monitor on the left, and the Wizard-of-Oz door opener on the right.

Based on the implicit interaction framework and theory, we expected people to interpret different door gestures to have different meanings, meanings that would reflect the perceived welcome or willingness of the door to interact. In these studies, people who were walking toward a door that “overtly prepared” to open by swinging open a little and then pausing before opening the rest found the door to be more welcoming and urging than doors that just opened, or that opened and then closed on them before they arrived. In follow-up studies which showed different videos of a person walking up to an automatic door that opened a different way in each clip, study participants were more likely to agree that the door that prepared-then-opened and the door that opened and then closed seemed to be thinking.

3.6 WIZARD-OF-OZ TECHNIQUES

Now, we might protest, “It’s not *really* an automatic door! This is really a test of how people interact with men who hide behind doors with levers.” One of the things implied by the Implicit Interaction Theory is that such a substitution is permissible; from the interactant’s perspective, the way in which the interaction is generated is not important. It is of no consequence if a person or a machine generates the motion. Only the effectively communicated behavior or action matters to the person walking by. This design technique, where we substitute a person performing a function that is intended in the longer term to be performed by a machine, is called “Wizard of Oz,” after the faux-magical character in Frank L. Baum’s books (Baum, 1990). Wizard-of-Oz techniques have been used by human-computer interaction researchers to design speech and natural language interfaces (Thomas, 1976) (Dahlbeck et al., 1993), multimodal input (Salber and Coutaz, 1993), contextual input (Hudson et al., 2003), gestural input (Höysniemi et al., 2004), as well as to prototype appropriate system responses (Maulsby, 1993; Horvitz et al., 1998; Klemmer et al., 2000).

In our experiment, situating the Wizard-of-Oz study in the field gave us the benefit of seeing how people would respond to gesturing doors in a natural setting, particularly how they would respond the first time they encountered such a door. However, with this experimental setup, it was difficult to ask people to evaluate their reactions toward different door gestures in the context of other possible gestures; once they saw how the door was actually operated, it would be harder to interpret the movements as coming from the door itself. In addition, we were concerned about the effect that the natural variations in door gestures might have on people's interpretations. Finally, we found that the people encountering the door were usually on the way from one place to another, and were generally too impatient to write more than a couple of words in the written responses. Thus, it was necessary to find other ways of validating these interaction designs that allow people to compare gestures in a context where they have a reasonable representation of the interactive experience but are able to provide more detailed feedback.

3.7 VIDEO PROTOTYPING TECHNIQUES

For a secondary experiment, we decided to use video prototypes of gesturing doors, so that participants would be able to compare the different door gestures and scenarios, would all be looking at the same door gestures, and would have more time to explain what they felt different door gestures meant and how they responded.

Video prototyping is a promising approach both for designing novel interactions as well as developing convincing empirical evaluations with large participants. In the words of Syrdal, et al.: "Using video rather than live interactions for evaluating prototype behaviors allows for a focused investigation into specific issues involved in the development of a system, without requiring a fully functional platform" ([Syrdal et al., 2008](#)). In this approach, designers and researchers generate videos of possible human-robot interaction scenarios and have them evaluated in a systematic manner. Because video prototypes are easy to produce, they allow designers to test out how a system should be designed prior to building the system in whole. In addition, video prototypes are useful for interaction research, because it is easier to show a large number of people the same interaction, and because it is easier to disseminate video and solicit responses than to recruit as many people to witness the same interaction in person.

The results from our video prototype study were largely consistent with the responses gained from the field; the effect of the interaction was clearly more vivid in the in-person experiments, but those participants weren't able to compare different door opening techniques. In addition, the video prototype study elicited far more descriptive responses to the open-ended questions, because the study participants did not have to be accosted while they were on their way somewhere.

3.8 FINDINGS

These experiments are obviously only early forays into understanding how people interpret the meanings of various autonomous gestures on the part of automatic doors. However, they suggest that the “offer” patterns generated analogously from the implicit interaction framework were interpreted as such by passersby. Analogous gestures were interpreted to signal intentional communication compared to the simple door-opening motion we used as a control. Despite the novelty of gesturing doors, untrained interactants “intuitively” read the gestures in systematic ways. This suggests that people have a common understanding of door interaction and interpretation of the meaning of door gestures, comparable to interpretations of human gestures (McNeill, 2005).

My colleagues and I have used these Wizard-of-Oz methods to evaluate interactions with kiosks, telepresence robots, and interactive furniture, and we have consistently found that video prototypes and Wizard-of-Oz experiments provide a strong indicator of how people will respond to the actual devices later. The people in our studies and experiments are not “fooled” into thinking that the device is smart. They understand the functional intent of the actions and know how to respond without being deeply concerned about the authenticity of the intelligence or reasoning underneath these communications.

3.9 CONCLUSION

In this chapter, we showed how the implicit interaction framework could be used to generate interaction patterns to help make interactions more coherent. One interesting aspect about the implicit interaction framework is that, because it doesn’t specify the *mechanism* by which initiative is taken or attention is commanded, it both allows and demands ingenuity on the part of the designer in creating the interacting entity. For example, a more literal analog of the doorman might be a doorman robot, one that would look and gesture and manipulate the door in exactly the same way a person would. The framework has nothing to say about whether a doorman robot or an automatic door is *preferable*. It suggests, though, that the critical functional component of the interaction around the door has to do with signaling information about the availability of the door to interact with. How those functions are achieved is up to the designer.

In addition, we used a number of special techniques—Wizard-of-Oz prototyping, video prototyping, as well as field experimentation and online experiments—to help us evaluate the design of our interactions from a broad swath of people before having to delve into the details of how to engineer the sensing and response. The incorporation of behaviors, responses, and evaluations from a wide number of people helps to make sure that our patterns are based on accepted conventions and cultural understandings rather than individual whim and eccentricity. At the same time, it is important to note that since the patterns are contextually and environmentally situated, it is important to know that there will be local patterns and norms that vary from place to place.

CHAPTER 4

Light and Dark: Patterns in Interaction

Now we've seen how to use the implicit interaction framework to map interactions, and then the implicit interaction theory allows us to use the framework to create analogous interactions. Let us move on to use the implicit interaction framing to consider some very common interaction patterns. We will use the question of how to turn on lights as our everyday example.

Engineers and designers are often reluctant to treat simple mechanisms like light switches as if they were agents with beliefs and desires. After all, they simply enact the will of the user mechanically. However, as Shoham describes, "It is perfectly coherent to treat a light switch as a (very cooperative) agent with the capability of transmitting current at will, who invariably transmits current when it believes that we want it transmitted and not otherwise; flicking the switch is simply our way of communicating our desires..." (Shoham, 1993). This is particularly true as the introduction of automation starts to blur the line between simple mechanisms and systems with more autonomy. "Whereas once this ascription of the light switch's function was unnecessary, increasingly, as we turn lights on and off with programming, with the help of sensors, remotely, via wifi or Bluetooth, the intentional stance becomes increasingly apt for describing the light" (Dennett, 1989).

4.1 THE DIFFERENCE BETWEEN PATTERN DIAGRAMS VS. STATE DIAGRAMS

Consider the classic light switch: it can get flipped up to turn the lights on, and down to turn the lights off. Easy.

The classic way that engineers are taught to diagram something like a light switch is using a state machine framework. The basic state machine for a light switch looks like this:




Figure 4.1: State diagram of a light switch.

The state diagram captures all of the states a lighting system can be in, and how the system can transition from one state to another. The great thing about state machine diagrams is that they enumerate the different states a machine can be in, and how those states are related to one another; this makes the diagrams a good way to model intended behavior. These diagrams are good to use when the system being described is a closed-system. What this type of representation *doesn't* capture is the interaction between the user and the lighting system, even in this very simple scenario. While a good state diagram will represent how the system will respond gracefully to unexpected input/events, it does not help people find these events.

One problem with using a state transition diagram is that it focuses all of the attention on the device's states; what is perceived or determined by the user is inferred, as is the range of the user's own actions. One example of this is the diagram of the Smart Light Switch by Cooperstock, et al.: You can see the designer automatically start thinking of whether there is motion or no motion based on whether the motion occupancy sensor perceives motion or not. This might have nothing to do with how the user sees the same environment, but this diagram doesn't provide any accommodation for the user's perspective at all.




Figure 4.2: State diagram of a Smart Light Switch described by Cooperstock et al. (1997).

While some (like Thimbleby (2010)) argue that state diagrams can go beyond describing devices and machines to describe what people do, such models assume people are operating with a static and perfect model of the interaction and cooperating as if they are part of the mechanism. However, it is not a good idea to use state diagrams to model human-machine interactions. In a state diagram, we need to articulate all the modes that are possible in the system we are designing. As designers, it is not possible for us to say what state the user is in, and it is a mistake to assume we can. At the same time, the state diagrams have a positivity bias; the diagrams are not meant to show all that can happen, only what is desired to happen. The system should be designed to take in a variety of different inputs and actions from the user, without making too many assumptions about the user's own state or mode.




Figure 4.3: People's conceptual models do not map cleanly to state diagrams.

Interaction patterns mapped on the interaction design framework, by contrast, are not system descriptions. The model for interaction is communicative and transactional, rather than mechanistic. This type of depiction focuses on patterns of activity, desired and undesired, and in fact helps us understand the functional meaning or purpose of the different activity patterns toward communicating the intention of the interactants engaging one another. The focus is on the transaction, on interplays of initiative, and on attention. When we design something like the interaction with a light switch, we should first analyze the space of possible interactions, and then design the state diagram of the light switch afterward.




Figure 4.4: The Rasmussen SRK model applied to the light.

It is also worth contrasting the implicit interaction framework with more cognitive models, such as Rasmussen's SRK model. Rasmussen pointed out in his seminal morphology of mental models for man-machine systems (Rasmussen, 1979) that there are actually different hierarchies of description simultaneously at work. In the S-R-K model (skills, rules, knowledge), the actual task of flipping a light switch is a skill. The associated assumption that flipping a switch will turn on a light is an example of rules. The deciding of whether the light should be on or off at any moment is knowledge. This model creates a hierarchy of the kinds of information, actions, and errors people may make when interacting with the world. However, it focuses on the user's fixed cognitive model rather than the unfolding and negotiated understanding that occurs during an interaction.

Our implicit interaction pattern approach toward interaction design is inspired by, but distinct from, HCI Design Patterns (Borchers, 2001). HCI Design Patterns, themselves derived from Charles Alexander's pattern language for towns, buildings, and construction (Alexander, 1977), capture the essence of a successful solution to a recurring usability problem in interactive systems; they focus designers on known solutions to previously encountered problems so that similar solutions can be applied to another design. This type of pattern focuses on the solution or the interactive device, and its visible features—a model of function in terms of objects and properties. The interaction patterns we use in this book, in contrast, are models of function, focusing on patterns of communication between interactants. These patterns put aside the issue of machine intelligence about the context or situation at hand, focusing instead on attentional demand and initiative as the critical components.

4.2 COMMAND-BASED INTERACTION

The classic light switch can be viewed as a command-based interface. Command-based interfaces are those that are foreground-reactive, and such actions take place in the upper-left-hand quadrant of the implicit interaction framework. (Throughout this and later chapters, we will represent common interaction tactics by representing them with brackets and capital letters, such as {COMMAND} so that it is easy to identify them.)




Figure 4.5: Simple framework diagram of a {COMMAND}, user turning on the light as a foreground-reactive task.

Actually, though, there's a bit missing here. Sometimes we flip the switch, and the light doesn't turn on. Maybe the bulb is broken. Maybe someone bought a cheap fluorescent light with an unforgivable delay. Maybe the light we're controlling is actually in the next room. All we know is, something failed in the interaction. What do we do then?

It is helpful to be able to reflect this in the diagram. Hence, to distinguish the difference between successful and unsuccessful interactions, we can mark the action with an exclamation point for successful conclusion, and an X for an error.




Figure 4.6: An exclamation point or X helps to indicate whether an interactive maneuver was successful or not.

Commonly light switches have a mechanical “snap” to them that makes it so that the light switch tends to flip either to on or off, but doesn’t easily sit in between. This snap feeling gives people positive feedback that they successfully changed the state of the switch, regardless what else is going on down the line in the circuit. In some ways, we could map this as the materials of the switch “conversing” with the user, with the snap being placed in the proactive foreground space. However, since the conversation is really more of an echo to help the user affirm that their command was sent, and people usually feel mechanistically about such feedback, I prefer to diagram feedback this way, with quote marks around the node to show the feedback.




Figure 4.7: Quote or echo marks around the point help to show the positive {FEEDBACK} the user gets in response to flipping the switch.

4.2.1 BREAKDOWNS INVOKE THE INTERACTIVE PARADIGM

We normally have a mental model of the light switch as a mechanical device or tool. In contrast, often when the light fails to come on after we flip a switch, we become more aware that the one-to-one relationship we have constructed between the flip of the switch and the illumination of the room is actually the result of a more complex chain, one that is usually robust but does break down now and again. When the light switch fails to produce light, we both become more aware of it, and it also seems stranger to us. This phenomenon was described by the philosopher Heidegger as *breakdown* (Heidegger, 1962). From an interaction perspective, because the light switch suddenly seems like it has its own agenda, possibly separate from our own, we should transition from modeling the light switch in the foreground reactive space alone, and start using the foreground proactive quadrant.




Figure 4.8: The failure of the light to go on in response to the switch leads to {BREAKDOWN} when user is suddenly more aware of the system.

4.2.2 COMMON ERRORS: MISPERCEPTION, MISJUDGMENT, MISEXECUTION

One good thing about separating out the user's actions from the light's reactions is that it helps us more finely parse the space of potential interaction problems. For example, issues of misperception (a problem with the sensing), misjudgment (a problem with the decision making), or misexecution (a problem with the attempted action) might have similar effects, but different remedies.

Similarly, it is important to distinguish *whose* error it is. Sometimes the problem is on the user's side. Issues of misperception or misjudgment map to what Don Norman calls the "gulf of evaluation," whereas misexecution is due to the "gulf of execution" (Norman, 2002). However, when the system is no longer merely mechanical, it is possible for the system, too, to have problems with perception, judgment, and execution. Indeed, Bellotti et al. have pointed out that with sensing systems such as the simple one we're describing, it may make more sense to describe the communicative rather than the cognitive questions and challenges in an interaction (Bellotti et al., 2002). Distinguishing these is important, because often we need to design systems to make it easy

to diagnose and recover from mistakes of different flavors. We can distinguish these on the framework diagrams by drawing the “X” that indicates error in front of, on, or after the interaction point.

	Diagram	Example of User Error	Example of System Error
Misperception	X•	User doesn't find the switch	Light doesn't get signal from the switch
Misjudgment	X	User is using the wrong switch for the light	Light is programmed incorrectly
Misexecution	•X	User doesn't push hard enough to flip the switch	Light is mechanically unable to light

Figure 4.9: Diagrams and examples for different types of user and system error.

For users, it is difficult to reconcile errors if they cannot identify the cause. Even for the classic light switch, there are both user errors and system errors to consider. It is important for designers to try both to prevent and disambiguate these errors. Light switches that have an LED on the switch, for example, help people find the switch, and know if they've flipped it enough to change the switch state.

4.2.3 REPAIRS

What do people normally do when lights don't come on when they expect them to? In my observations, they try to turn them on again, usually a bit more emphatically. *Reiteration* can be a successful repair technique if the problem is system misperception, and perhaps re-executing the signal, louder or more emphatically, is an indication of the user's mental model of the error. However, if in fact the lack of light is due to a problem with system execution—say, a little bit of latency in the response of the lighting due to CFL bulbs, for instance—the user's actions just exacerbate the problem. Mapping these interactions on the Implicit Interaction Framework helps to frame the issue as one of miscommunication. Miscommunications are common in our interactions, so it is important to contemplate how to design repair mechanisms when interactions go awry. This is different than in system design, where run-time errors are far less common and usually the indication of a design flaw rather than part of everyday function.




Figure 4.10: Implicit Interaction Framework showing user's attempt to {REPAIR} lighting {BREAKDOWN} through {REITRATION} by “switching louder.”

The other common methods of repairing a light not turning on speak to the communicative mental model people employ even when interacting with a simple machine like a light with a light switch. They sometimes *re-address*, trying other switches in case they misjudged which switch controls the light. Sometimes they *re-assess*—to see if there is a crucial missing piece of information that might clarify the situation: another additional switch on the light itself, perhaps.

4.2.4 BASIC PATTERNS

The diagrams we've employed in exploring our interactions with the basic light switch are actually common enough that we call them *patterns*. In the previous examples, we labeled the diagrams with the particulars of the interaction in question, but without those details we are left with the basic patterns for interaction: command, feedback, execution.

Even though we were talking about a light switch when we first generated these diagrams, these patterns would be the same even if we were using push-button switches, or pull-chain lights, or voice commands. The type of feedback, or the occurrence of particular errors might be different (for example, a voice or sound-based switch, or a toggle, does not have the persistent reflection of state that a flip switch has).

We have discussed some very basic patterns that show interactions that don’t work out well, or fail. Basic patterns include misperception, misjudgment, misexecution, and can occur on either the person or the system side. It is important to consider these when designing interactions because the fact that there are two independent entities means that communication failures are common and repair mechanisms need to be put into place.

4.3 INTERACTIONS WITH “AUTOMATIC SWITCHES”

Often, we want lights to go on and off without having to have someone flipping a switch. For example, people are terrible about turning lights off when they don’t need them on anymore, and this is a waste of electricity. A really basic solution is to put the lights on a timer, so that they are on for a fixed duration of time. This is a background-reactive interaction. If you have left the room by the time the lights go out, the whole issue of controlling the lights is out of sight and out of mind. One thing about things that happen in the background, though, is that you can’t fix what you don’t see.

4.3.1 AUTOMATIC SYSTEMS HIDE PROBLEMS

If the lights don’t go out after you’ve left the room, how will you know?




Figure 4.11: Automatic Light Switch mode error leads to {ALIASING} so that the person cannot easily tell when the timer is going to turn the light off or not.

It is important for the lighting system to provide some indication about what is going to happen in the future, so that you can perform a repair if it is necessary. Without this indication, the two future states are aliased; a person has no way of knowing what state the lights will be in an hour after she leaves.

4.3.2 USING FEEDFORWARD TO SHOW FUTURE ACTIONS

An example of a mechanism that does indicate the future state is a hot-tub timer. Usually a person has to turn the dial from off to a certain amount of time—say 15 minutes—and then they can see when the time is almost up, and reset it, or they can be sure that if they leave a bit early that the tub will stop bubbling after a few minutes.




Figure 4.12: Diagram of a hot-tub timer. The {FEEDFORWARD} signal enables the user to {OVERRIDE} the settings, correcting when the lights will go out.

4.3.3 TESTING BACKGROUND PROACTIVE SYSTEMS

Another example of an automatic switch is a dawn-dusk light sensor that outdoor flood lamps are often equipped with. One thing that you might do to make sure that your flood lamp is actually on when you test it during the day is to simulate darkness by covering the light sensor.




Figure 4.13: {SIMULATION} procedures allow user to test the future function of a system.

The issue with background tasks, whether they are reactive or proactive, is that they occur beneath a person's notice. If there are problems, they will take place outside of the user's notice. Hence, it is important to use forwarding techniques to help the user identify possible issues when their attention is available. In the case of the hot-tub timer, the feedforward is built into the system interface. In the case of the dawn-dusk sensor, the user initiated the forward simulation. It is important that interactive systems provide means for performing these tests. Otherwise, a failed smoke alarm would be indistinguishable from a working smoke alarm until the moment an obvious fire occurred! In both cases, the generalized solution is that we raise actions that happen below the attentional horizon line into the realm of explicit notice to identify or prevent problems.

4.4 INTERACTIVE SWITCHES

Up until this point, we talked about pretty basic switches, switches where a person either controlled the switch or set up an automatic response to time, daylight, etc. There has not been the sense up until this point that there is another entity controlling the lights unless the user runs into a problem.

However, now let's consider "smart switches" which use motion occupancy sensors. From an engineering perspective, switches using occupancy sensors are no different than those using light sensors or timers. However, because they are responding to our actions (or lack thereof) these

44 4. LIGHT AND DARK: PATTERNS IN INTERACTION

switches feel different. When the lights go out, we feel wronged, because we have been misjudged, and harshly. It is as if there is another person controlling the lights that we need to reason with.

Many of us are familiar with the solution; we have to wave our arms to turn the lights back on. This is form of repair called an override. Whereas a correction is a fix of one's own mistake, an override is way changing someone else's course of action.




Figure 4.14: A smart switch {MISJUDGEMENT} of room occupancy, followed by {OVERRIDE}.

What if we had a light switch that dimmed the lights rather than turning them out when it sensed a lack of motion? The method of repair could be much the same—the flapping of the arms to indicate “I’m still here”—but somehow, the level of presumption is less offensive. It is not only important in the interaction to have a path to override, but to structure the interaction so that the system seems open to being overridden.




Figure 4.15: A re-designed smart switch: {MISJUDGEMENT} of room occupancy is followed by dimming rather than extinguishing of lights, followed by {OVERRIDE}.

4.4.1 INTERACTION ERRORS

What are errors that can be made in the interaction with the smart switch? Like previous switches, the switch might misperceive the situation, it might misjudge what to do, or it might mis-execute its intended action. The specifics of what the switch is trying to perceive, judge, and execute might be more complex, but the class of errors is the same. For example, it is not always the case that people want the light on when they are in the room—they might want to take a nap, or test a lighting device, or take a low-exposure photo. In these cases, the misjudgment is easily repaired as long as there is a way to override the default action.

One of the terrible failure cases with smart light switches is if the light itself burns out—people sit there flapping for a long time before figuring out what is going on. It is actually a good idea for there to be a light next to the occupancy sensor that is on when it senses and blinks when it detects motion, so that people can diagnose whether the sensor is perceiving and what it is judging. These signals help people to discern the difference between misperception, misjudgment and misexecution, which is important because the repair mechanism for each is different.

However, now there is a host of interaction errors that also needs to be considered. Every action is also potentially a message, and that message can be too loud, or not loud enough. Each series of actions is considered a behavior, and that behavior can be happening too slowly, or not quickly enough, too rarely, or not often enough. Shall we just ask the user what is desired, then? Well, every address to the user must command the right amount of attention, and not be too presumptive, or over- or under-solicit information.




Figure 4.16: Showing classic error patterns of {PRESUMING}, {BELABORING}, {OVERSOLICITATION}, and {UNDERSOLICITATION}.

4.4.2 REPAIRS AND OVERRIDES

Unlike the mechanistic errors, the list of things that can go wrong with an interaction seems to be endless. Often we might hear that smart switches are more trouble than they are worth, even though automated light switches and thermostats can literally yield billions of dollars in energy savings (Lu et al., 2010; Williams, 2012). Clearly, the measure of success for an automated system is asymmetrical; a few errors can mar hundreds of hours of successful function.

There are two causes for hope. The first is that we are people designing for people, and we have some ability to plumb the depths of our own interaction experience for intuitions about how best to design future interactions if we can frame the analogies right. The second cause for hope is this: people, for all their vaunted intelligence, make inference mistakes all the time, and are usually forgiven.

When people go out on a limb, taking initiative in the face of uncertainty, they engage in compensating measures, hedging their actions with techniques such as overt subtlety (where actors make a show of how unobtrusive they are trying to be) or pre-emptive apology (where actors may bow their heads, scrunch up their faces, or raise their shoulders as they execute an action to indicate an apology if their initiative is unwelcome). Research on interruptions at Carnegie Mellon ([Fogarty et al., 2005](#)) and Microsoft Research ([Czerwinsky et al., 2000](#)) that have focused primarily on *when* to interrupt could be complemented by research on *how* to interrupt. There are conventional ways to act proactively, even in the face of uncertainty, and these are a matter of sociable design rather than technological intelligence. This “facework,” sometimes called social graces, *savoir-faire*, diplomacy, or social skill ([Goffman, 1967](#)) is critical to smoothing over the inevitable mistakes in interaction. In our examples here, we have seen that people often try to recover or repair errors through reiteration and override. They also try to predict and prevent failures by employing feedforward and simulation. When all else fails, it can be useful to decrease the cost of misjudgments by decreasing the presumption of the automated systems actions. The smart switch that dims the lights slowly rather than turning them out all at once is like the doorman that opens the door a little rather than flinging it open at every prospective passerby.

4.4.3 GUIDING LIGHTS

Another way that we could address the issue of people forgetting to turn off the lights is to have the light switch cue the person to act. For instance, in my home I like to just use natural light during the day. The problem is that if the room is already bright with sunlight you often can't tell that a light has been left on. I often think lights should respire softly if they are on in the day, gently enough so that it is not bothersome if you meant to have them on for some reason, but clearly enough so that you would remind a person who had forgotten. Like in the previous example of motion sensors, where the person needed to have a model of what the light was perceiving, judging, and attempting to execute, so in this case the system needs to have a model of what the person can see, think, and intends to do. The implicit interaction framework helps us realize that we need to have a cue that the light is on, but it doesn't have to be the visual cue of the light itself. Depending on the context, it's possible that the right clue might be a sound, or a nudge, or a message.




Figure 4.17: Sometimes {REMOTE ALERT OF STATE} is required to repair errors.

4.5 DARK PATTERNS

The previously discussed errors in classic, automatic, and interactive switches assume good-faith engagement on the part of both the user and the system; there are also “dark patterns” we could consider. A *dark pattern* is a special kind of pattern, defined as: “a type of user interface that appears to have been carefully crafted to trick users into doing things (where these user interfaces) are carefully crafted with a solid understanding of human psychology, and they do not have the user’s interests in mind” (Brignull, 2011). In their paper on “Dark Patterns in Proxemics” Greenberg et al. draw a distinction between dark patterns and *anti-patterns*. “Anti-patterns are designs that *unintentionally* result in a negative experience or even harm” (Greenberg et al., 2014). These include the Captive Audience or the Attention Grabber (when the system takes advantage of an otherwise unwilling audience to initiate interaction), and the Bait-and-Switch (where the content of the interaction offered at the outset of engagement is different than the one provided in the end).

One common instance of the bait-and-switch uses misdirection; for instance, in a room with one wall switch and one light, a user may believe that the switch should control the light. A devious designer could bait-and-switch the user and use the switch to, say, turn on a hidden web camera. These dark patterns are often the result of misaligned interests between the interactants. Since a lot of the design thinking and research on interactions assumes good-faith cooperation, future investigation and profiling of dark patterns would definitely be worthwhile.

4.6 INVISIBILITY IS A CONSEQUENCE AND NOT A CAUSE OF GOOD DESIGN

In general, I have avoided using traditional user interface-type solutions such as dialogue boxes or speech interfaces. This is not because implicit interactions need inherently to be non-verbal. However, much of what makes implicit interactions work is that they function at people's attentional periphery. Dialogue-based interactions tend to be focal; it is difficult to carry on more than one dialogue at a time. The implicit interactions we have been exploring, however, make only occasional bids for attention. This makes them more appropriate for placement in people's everyday environments.

The classic light switch only acts in response to the person, and most of us are so familiar with it that we use it without thinking. The smart light switch *can* be an improvement upon the classic light switch, but only if its plans of action are sensible, and if it makes its plans clear enough so that they can be overridden or corrected if need be. Both types of switches can seem invisible in use; neither attains its invisibility through subtlety.

This is a point important enough that it is a mantra of mine: *Invisibility is a consequence and not a cause of good design*. Invisibility in use doesn't come from subtlety in communication. It comes from clarity of communication, even if that communication doesn't come through writing or speech. It is nigh impossible to attain clarity of communication by obscuring the signal.

4.7 CONCLUSION

In this chapter, I used the case study of how to switch on lights as a way of examining several fundamental patterns in interaction as they present themselves in the implicit interaction framework. I considered the role both of positive patterns, which we try to emulate, and error patterns, which we seek to avoid, mitigate, or reconcile. I also briefly considered dark patterns, where the intentions of the interactants might not be in line with one another. By mapping out potential problems and solutions using the implicit interaction framework, we are able to consider a wider space of interaction possibilities than are suggested by the more commonly used framework of the state diagram. The framework also helps us to consider new possibilities of ways to interact that weren't present in previous systems. In this way, the implicit interaction framework can make interactions more robust and inventive.

CHAPTER 5

Action and Reaction: The Interaction Design Factory

Up until this point, we have focused on using the Implicit Interaction Framework for analyzing and making simple the use of analogues and patterns. We have been looking at how the pacing, dynamics, signaling, and expectations of interactions can be mapped into the space of initiative and attentional demand so that our interactions with machines can be informed by our experiences interacting with people. However, interaction designers are rarely tasked with these simple types of problems; more routinely, they are asked to design intuitive and easy-to-use interactions for novel technologies, often ones without a lot of easy analogues to fall back on. As our designs move further into uncharted territories, it becomes increasingly important to have a playbook of successful interaction strategies to draw from.

Reeves and Nass pointed out in the Media Equation that people's interactions with interactive systems are fundamentally social, so that any factors that govern interpersonal relationships should be expected to apply to human-machine interactions (Reeves and Nass, 1996). But how do we convert these understandings of social factors into actual designs for novel interactions? By identifying key social factors and mapping how they affect the initiative and attentional demand throughout the interaction, the Implicit Interaction Framework makes it easier for designers to map how to control the "levers" they have available to guide the interactive experience. In this way, designers can better grasp how to draw upon their social intuition and experiences to create interactions that have the desired interactive effect. This chapter illustrates the process of turning social science insights about the factors that affect human interaction to novel interaction designs that are effective. We focus on three factors related to implicit interactions—proxemics, intention, and consistency—and look at how they were applied to shape the design of interactive whiteboards, personal robots, and telepresence systems, respectively.

5.1 PROXEMICS, OR, DANCING WITH THE MATERIALS OF DESIGN

When people are engaged in discussions, they tend to move closer and farther from one another depending on what they are speaking about, and what relationship they have to one another. Some of the reasons for this behavior are pragmatic. You stand farther from strangers and closer to those

who are familiar. You speak more of public topics and share information with people who are a few arm's lengths away, and draw closer to discuss more private and intimate matters (Hall, 1969). And yet, these proxemic patterns are so ingrained in our communication patterns that they endure when the pragmatic reasons for them don't apply. In the online virtual community Second Life, for example, players communicate via typed statements, and they generally use a "god's eye" view, so they don't need to be closer or farther than one another to gain more or less of a view of a fellow avatar. Nevertheless, researcher Nick Yee found that the behavioral patterns of interpersonal distance that have been documented in the physical world still hold in online virtual environments. Male avatars stand farther from other male avatars, while female avatars stand closer together. When forced to stand closer together than is comfortable, avatars would avert mutual gaze. Just as in the physical world, male-male dyads preferred less intimacy, and were less likely to look at each other indoors where large interpersonal distances were not allowable.

This "hidden dimension" of interpersonal distance is one that people often use without explicit notice. Hall noticed that people had a definite sense of personal territory, which he classified into four zones: intimate, personal, social, and public (Hall, 1969). As their titles suggest, different types of interactions take place in different zones, and people move around to maintain distances appropriate for different types of interactions they are having. We feel uncomfortable if we are too close or too far from others for the interactions we are engaged in.

5.1.1 THE ROLE OF PROXEMICS IN INTERACTION

We engage in similar cyclic changes in physical distance with objects and materials we are working closely with. Think, for instance, of how you interact with a whiteboard when you are sketching out ideas or a drawing. You move in close to scribble your ideas on the board, and, after a stretch, you take a step back to reflect on whether the newly added nugget is spelled right, if it needs amending, how it relates to other bits on the board. Then you step back in and do more writing. These rhythms are present in all of our creative activities (McKim, 1972; Cross, 2001), but at a large display or canvas, the evidence of this cyclic switching is manifest physically (Tang and Leifer, 1989; Gill and Borchers, 2008).

When people are drawing on a shared whiteboard as part of discussions with others, the physical back-and-forth becomes even more pronounced. In the observational research around traditional whiteboard usage I performed with my colleagues, for instance, the changes in physical distance from the board correlated with changes in design activity (Ju et al., 2006). Design teams engaged in informal meetings would cycle between phases of drawing and analysis; these changes corresponded with changes in their physical proximity to the whiteboard. Users would stand close to the board when they were writing, farther back when discussing written artifacts in detail, or farther back still when engaging in meta-discussion. The visual artifacts of these engagements are exceedingly brief—in general, people draw the minimum needed to support the conversation and

very few of the drawn artifacts of whiteboard meetings are coherent on their own. People stand close to the board and write a few things, step back and talk to one another for several minutes, go back and write a bit more. These patterns serve to underline the fact that the process of creating the drawings is as important as the drawings themselves (Bly, 1988). In fact, in the case of collaborative design meetings, it seems clear that engaging in the iterative process of design, having a collaborative conversation about the many possible solutions, and eliminating false leads is more important than the drawn artifact in influencing the final design (Tang and Leifer, 1989). Hence, the physical location of the people interacting is a better indicator of the phase and nature of the design activity than the written artifacts themselves.




Figure 5.1: Proxemic factors correlate with changes in activity in a wide variety of interactions.

5.1.2 PROXEMICS IN INTERACTION DESIGN

The key insight for interaction designers, then, is that the bodily position, orientation, and eye gaze correlate in meaningful ways with shifts in contextual activity or meaning. This is critical because bodily positions, orientations, and eye gaze are relatively easy to detect compared to changes in conversational topic or tone, or shifts in interactive activities. By using physical shifts as a proxy signaling shifts in task or topic, an interactive system can respond “intelligently” without needing actual intelligence, interrupting at appropriate times or making suggestions when they are most useful. The same action, taken at a slightly different time, might have a completely different outcome.

How might awareness of the dynamics of proxemics influence whiteboard design using the frameworks and methods described previously? For instance, sketching activity tends to be bursty, and people generally perform things like role shifts or pen color changes between bouts of actions during periods of new thinking as well as more intermittent moments of reflection about the activity. With an interactive whiteboard, it is possible to have the whiteboard’s behavior respond in kind to the changing phases and episodes of activity.




Figure 5.2: Observations of people interacting at a whiteboard suggest that location in front of the board correlates with changes in task activity and cognitive function.

The whiteboard, then, can attend to the background reactive cue of a person's physical proximity as a signal associated with activity the person is engaged in. For instance, the digital tool tray that lets me switch ink colors can be hidden and slide out only when I take a step back when I'm pausing to reflect on my sketches a bit. My physical location might tell more about what aspect of the design activity I am really engaged in than the sketches on the board itself. When the whiteboard responds to the background reactive cue, it transitions from being a natural and unconscious behavior to something that is reflected and converted into a communicative signal by the whiteboard. It is important that this overt and attention-grabbing cue occur at a phase in the drawing and designing activity when a person is receptive to input, suggestions, and other communications, rather than during the active writing phase when they really need to focus on just writing. In general, timing is really crucial when transitioning from Background-Reactive to Foreground-Proactive, because it demands major shifts in locus of attention and cognition on the part of the people involved.




Figure 5.3: Observations of people interacting at a whiteboard suggest that location in front of the board correlates with changes in task activity and cognitive function.

Proxemics, then, can be used as a key factor for designing interactions not only with white boards, but large public displays (Vogel and Balakrishnan, 2004; Marquardt and Greenberg, 2012), as well as any number of ubiquitous computing systems (Weiser, 1991) in the world. Technologies that help to sense the factor of proximity, such as infrared distance sensors, Bluetooth low energy modules, and gaze trackers, can help to instrument a wide variety of applications where implicit interactions are desirable because proximity is so crucial a factor to human interaction.

5.2 INTENTION

Intention, the indication of what a person plans to do next, helps us with expectation management in our day-to-day interactions. In our original doorman study, for example, the doorman's overt preparations to open the door help to express intention to people who otherwise might not know if they would be permitted to enter. Intention precedes action, and, to a large extent, we are not aware of how we express our intentions before we act, nor how we observe it and factor it into our own expectations; intention is expressed and absorbed naturally and obliviously. We tend to notice the importance of things like intention more when we are unpleasantly surprised by something that someone else does.

5.2.1 THE DARK PATTERN OF SECRET ROBOTIC PLOTTING

We have largely avoided discussing robots, because the need for communicative analogues does not stem from or rely on human likeness. There are naturally some applications for robots where there is a need for the robot to look and act like people. Realistically, though, robots are often designed to take over tasks that people aren't good at, like lifting heavy things, or performing repetitive tasks. They often need to be non-humanoid to fulfill their function. In any case, even robots that look like people don't really function like people; they have more in common as far as sensors or cognitive matter with a smart phone than a human. Because robots are not like us, we people can have difficulty understanding what they do. This gap in understanding can be frustrating—even lethal. Even with their very limited deployment to industrial environments, robots have killed and maimed dozens of people in the last 30 years (Markoff and Miller, 2014). Most of the injuries and deaths have happened when humans who are maintaining the robots make an error or violate the safety barriers, and because the robots behave differently than people expect them to. As robots start to be deployed to work alongside people, it becomes increasingly important that the robots' intentions are made clearer to the people that they work with.

Although I have never encountered a robotic doorman, I have worked alongside a robot that can open doors: the Willow Garage PR2. The PR2 is a personal robotics platform that is roughly humanoid: it is human height and has arms, a head, and a moving base (Bohren et al., 2011). The PR2 was meant to operate in offices and homes, alongside people, and much development work was put into having the robot navigate the rooms, halls, doors that people use everyday. Opening a door is hard for the PR2; it has to identify that some feature in the wall is actually a door and door frame, it has to locate the door handle, it has to figure out how to grab and turn the door handle, it has to push the door open, and only then can the PR2 pass through the door.

Many a time, the roboticists and computer scientists working at Willow Garage would see a PR2 stopped in front of a door with its head bowed. It was hard to tell if it was trying to sense the handle, or if it was calculating its arm trajectory, or if it had accidentally parked itself in an inopportune place. Very often, people—even the robot researchers that work with the PR2 everyday—would step in between the robot and the door, or wave their hands in front of the robot's "eyes" to see if it was looking. This actually had the effect of prolonging the time it took the PR2 to open the door, because it would need to start rescanning the door anew. It could even potentially be dangerous; if a person tried to pass around the PR2 to get through the door just when it was about to reach out and grab the handle or push the door, a person could be injured.




Figure 5.4: Overt preparation helps to signal intent and allows people to perform avoidance or override in response. Covert plotting can lead to misunderstanding, injury, or death.

For the PR2, the problem was that the people around it didn't understand what it was trying to do. In a people-free factory, a robot doesn't have to worry about communicating when it is sensing or thinking in order to sense or think without interruption. Then again, a people-free factory doesn't usually have doors! Because the PR2 is interacting in human territory, it needs not only to contend with doors, but also people's expectations for what it will do.

Research I performed with colleagues at Willow Garage and Pixar Animation showing animations of the PR2 performing various tasks indicates that performing forethought greatly improves people's understandings of the robot's intentions (Takayama et al., 2011). People who see videos of a robot performing "thinking" movements before executing an action—the saccadic shifts in gaze, the tilting of the head, the scratching of the brow—were significantly surer of their descriptions of what the robot was doing than people who didn't see the robot perform forethought. These robots that had showed forethought were rated to be more appealing and approachable. In addition, robots that react visibly to their success or failure at a task were rated to seem smarter than robots that did not react. In fact, the robot that seems disappointed in its own failure seems smarter to people than the robot that does not respond to its successful completion of a task.

5.2.2 THE ROLE OF INTENTION IN INTERACTION

From a design perspective, having the robot perform forethought was an important example of feedforward; actual forethought is all that is needed to open the door, but the performance of forethought is how others around the robot understand that the robot means to open the door. Similarly, the reaction of the robot provides a key piece of feedback. When the robot reacts happily or in a disappointed matter around its success or failure, it helps people understand what the robot's intentions were. Hence, a robot that can recognize that it failed to open a door is still perceived as smarter than a robot that can successfully open a door but doesn't seem to register the outcome of the action.

Intention helps us understand what is likely to happen next in interaction. The overt preparation that was mentioned in [Chapter 3](#) used a high attentional demand signal to convey a willingness or intention to open the door. By communicating intention, we can synchronize joint action: we can let others know to get out of the way, or to join in on our count. It enables people to perform override. It aids people in monitoring the situation. It prevents alarm or surprise.

5.2.3 INTENTION IN INTERACTION DESIGN

Sometimes intention is conveyed by relatively static cues—facial expression, head orientation, or bodily position. However, intent is conveyed dynamically. Maya Cakmak's work in robot handovers, for example, showed that spatial and temporal contrast was key in making people understand when a robot intended to hand them a drink; if the robot just held out a drink or moved too slowly, people were not sure that the robot wasn't holding the drink out for some other reason ([Cakmak et al., 2011](#)). Very often the language of intent is very contextual, and requires interactants to know something about the structure of the task, knowledge of cause and effect, or reasoning about the interaction partner's goals or beliefs.

We mentioned earlier that intention can be thought of as a flavor of feedforward (what will happen next?) or feedback (what just happened?). However, to express feedforward and feedback in ways that are easily mapped. Just as the doorman could make a feint toward opening the door as a way of showing an intention to open the door without saying a word, a robot could hold its hand out or lean toward the door as ways of expressing intent without resorting to words, beeps, or icons that need to be interpreted to be understood. Perhaps the most crucial lesson for interaction designers is that, in interactive devices, ways of having the system express what it is about to do need to be built into the system from the outset. The expression of intent for action can be as important as the action itself in interactive systems.

5.3 CONSISTENCY

Consistency is an important factor in interpersonal relations, because when different aspects of a person's demeanor suggest different things, their interactive partners can be confused and grow distrustful. In face-to-face interactions, inconsistencies in verbal and non-verbal cues are often interpreted as a sign of deceit (Kraut, 1978) and cause increased cognitive load (Fiske and Taylor, 1991) and mistrust (Ekman and Friesen, 1974). This principle of consistency carries over to computer agents; Isbister and Nass, for example, found that people prefer computer agents whose verbal and non-verbal cues were consistently introverted or extroverted to those whose cues were mixed (Isbister and Nass, 2000). This consistency-attraction effect is so strong that it trumps similarity-attraction, where people prefer personality types like their own (Nass and Lee, 2001). People respond positively toward agents that display consistency in verbal and non-verbal cues compared to those with mismatched cues (Isbister and Nass, 2000). They prefer that agents use gaze and gesture to provide contextual grounding for the agent and user's shared experience to those that do not (Cassell and Thorisson, 2000).

When we are interacting with others, we may make efforts to unify our body language, our facial expressions, and our vocal inflections so that they support the message that we are voicing, or we might subconsciously undermine them through inconsistent cues if we don't really believe what we're saying (Argyle, 1988). However, in the design of machines, inconsistency is often the norm, because inconsistency is just what happens unless significant efforts are being made to avoid it.

5.3.1 THE HEAD-IN-A-HEAD PROBLEM

Nowadays, it is increasingly common for people to communicate with co-workers through telepresence robotics systems. These systems feature the hybrid of a live video representation of the remote worker and a physical platform, often with human-body-like proportions, that gives that worker some amount of physical presence in the central work site. This setup enables remote workers to communicate naturally with their distant peers, have a similar presence as on-site workers (in terms of body size and location), and often allows them to move or be moved around the central workplace to engage in day-to-day informal office interactions (Kraut et al., 1990).

This configuration of a video display of the remote user on an articulating base creates a confusing juxtaposition, an issue we'll call the head-in-a-head problem. The remote worker actually has two representations in the central meeting space: that of the video feed, and that of the physical platform. These two channels are likely to portray inconsistent non-verbal facial or gestural cues. This lack of consistency causes people to feel unsure about what the remote person is communicating, and also to make that person seem more dominant, less friendly, and less involved (Sirkin and Ju, 2012).

It might seem that having a base that doesn't move provides a neutral stage on which to perform on-screen actions. However, the research suggests otherwise. Groom, et al. explain this in their own studies on behavioral realism: "People expect bodies not only to serve as decorations suggesting identity, but also as functional units intended to interact with the environment and to communicate information" (Groom, 2009).




Figure 5.5: Even interactions that demand more attention can be preferable to those where the various aspects of the system seem inconsistent.

What we found in these experiments was that people were a lot more sure of how to interpret the different behaviors they observed when the on-screen and in-space motions were consistent. In a follow-up study, we showed a scenario of a person using a similar telepresence setup to interact with two co-located collaborators. Different participants saw different versions of the same scenario, with the remote participant using just on-screen motion, or on-screen plus in-space motion. When the remote collaborator's telepresence system moved in concert with her expressions and gestures, she was perceived to be friendlier, less dominant, and more involved. More interestingly, in the same conditions, the on-site design collaborators were rated as being more equal. Perhaps participating in a meeting through static video chat creates a perception of inequality between participants. This inequality may cause remote teammates to appear higher, or perhaps lower, in stature. For example,

the static platform may have caused the remote participant to appear dictatorial, making requests and issuing instructions from behind a flat panel. After seeing him lead the discussion, one participant commented: “I thought [the remote participant] sounded a little rude; it seemed like he tried to hurry [the local participant] and the other person at one point.” Re-introducing physical motion and gesture by the remote teammate may ameliorate that perception, and put the participants back on more even terms relative to one another.

5.3.2 THE ROLE OF CONSISTENCY IN INTERACTION

On the whole, our study points to the importance of considering all of the channels of communication present in embodied or automated systems. Even if one mode or channel of information is dominant or focal, the other channels color people’s interpretations of what is communicated. Consistency and reinforcement helps people to interpret signals correctly; inconsistency creates confusion and doubt. The presence of multiple communication channels has the power to add or detract, and the effect can enhance or undermine trust and credibility.

From the design perspective, being able to use consistency or inconsistency raises the dynamic range of interaction: a system can convey a message with one channel, but convey confidence or uncertainty about the message through consistent or inconsistent parallel channels.

5.3.3 CONSISTENCY IN INTERACTION DESIGN

For designers, the key lesson is to consider all the different ways that a system is conveying information; often aspects of the design that are assumed to be “neutral” are anything but. Communications researchers know the axiom that you cannot *not* communicate (Watzlawick et al., 1967). A corollary should be that every aspect of the system communicates. Merely lowering the attentional demand of some aspect of the system—for instance, by not moving—does not remove it as an influence in the interaction. So much is interpreted to be part of an interactive systems’ overall demeanor: if a system has a voice, the tone and inflection of the voice can affect the interpreted message. If a system interface is primarily visual the aesthetic layout and graphical decisions can bias a person’s reading. In embodied systems, small accommodations made to reinforce actions—orienting the head or eyes toward the person being addressed, slowing or pausing actions to emphasize the primary function—are important. Normally, designers aim to keep all aspects of a design consistent, in order to promote trust and confidence, but it may be that incorporation of inconsistent signals can help, too. They can help to convey a need for heightened awareness or redoubled engagement.

5.4 CONCLUSION

Jamer Hunt says, “The ubiquity of moments of interaction—of sites of speaking and listening—means they will increasingly become contested areas: who will have the power to initiate, respond,

and ultimately control the conversation? The designers who plan, choreograph, and implement these interactions will have to subtly calibrate their presence in our lives, or else risk creating a din that drives us all into noise-cancelled hermitic retreat" ([Antonelli, 2011](#)).

In this chapter, we've thought more about how to map factors into the design of interactions with novel devices. Proxemics, intentionality, and consistency are far from the only factors that influence social interactions amongst people, but an understanding of each is hugely generative for interaction designers trying to design interactions that are understandable without being noisy. These factors are like instruments that play in concert with the rhythm of the interaction patterns introduced in earlier chapters. Used well, they can make for rich and harmonious interactions. Used poorly, they can be disruptive and disturbing. What these factors have in common is that they are social factors. There is a vast literature describing the role these and related elements play in human interactions, and these can be sourced by interaction designers looking for novel territory in human-machine interactions.

CHAPTER 6

Driving into the Future, Together

This book has laid out an approach for designing human-machine interactions that mines our insights about everyday social interactions for creating novel interactions that nevertheless work in a familiar way. Many of the technologies and techniques for implementing these interactions can be highly technical. It can be difficult to make sensors that perceive a situation as a person would, and so translational effort is required to establish common understanding between actors. It can be difficult to make a machine move in ways that are expressive or safe. It can be difficult to build complex models of what sorts of behaviors are appropriate in the myriad contexts, with myriad different users, that people encounter all the time. Still, for all the sensing, modeling, or actuating difficulties that underlie human-machine interactions, the standard of proof is always how the interaction works for everyday people.

This book has focused on fairly ordinary and everyday interactions and technologies to make the point that both the need for and the secrets to good interaction design are all around us. Now, as we drive toward the end of this journey, I would like to conclude by showing how the very same ideas and tools that help to make better automatic doors, office whiteboards, and telepresence robots can be applied to the once fantastical and still mind-bending prospect of autonomous vehicles.

6.1 THE RATIONALE FOR AUTOMATION

Nearly any introduction of automation brings strong feelings of resistance (Poppel, 1981). We have gotten along this far in the path of human progress; is more automation really necessary or desirable? Of course some people resist change at every turn, but at the same time, there are valid questions about the nature of the interaction that is introduced, the impact it will have on our actions, the culture that it will impose. It is important for these technologies to take into consideration the context and situations they are being introduced to, and for these technologies to adapt.

In the case of cars, people are quick to point out both the pleasures of driving that would be lost, and of the inevitable accidents and loss of life that will ensue.

The truth is that people are terrible drivers. There are more than 30,000 traffic fatalities annually in the U.S. alone, and \$277 billion in economic losses that occur in the U.S. annually from traffic accidents (NHTSA, 2014). That number doesn't include the billions of dollars of economic losses and environmental damage due to traffic congestion. In many ways people are poorly adapted to the task of driving; in the U.S., at least, it involves long periods of near inactivity and boredom interspersed with rare moments requiring rapid situational assessment, judgment, and physical exe-

cution. Automated vehicles should be better equipped to handle the aspects of driving that require constant vigilance as well as the rarer moments that require split-second quick sensing, judgment, and actuation. In addition, platooning—the grouping together of cars in energy-efficient configurations, can smooth traffic, increase fuel efficiency, and increase operational speeds during periods of heavy traffic (Robinson 2010).

6.2 INTERACTION DESIGN ON THE CRITICAL PATH

The caveat, of course, is that before autonomous vehicles can reduce deaths or traffic, they need to work for and with people. Of the many obstacles—technical, financial, ethical, legal—that developers of autonomous vehicles face, the interaction-centered issues are the most difficult to resolve. How can autonomous cars operate safely on the road with other human-driven cars? How can they anticipate the actions of those who share the streets with automobiles? More importantly, how can the vehicles interact with the drivers—to communicate capabilities and limitations, to forecast actions and explain changes—in ways that inspire trust and confidence? How can the interfaces be designed so that when people need to take control of operating the vehicle that transfer of control goes smoothly?

At the heart of these challenges is the question of how to design interactions. In [Chapter 2](#), I introduced the Implicit Interaction Theory and Framework, which argues that 1) interactions are fundamentally about communication, albeit often non-focal and non-verbal communication and 2) because people already have norms and patterns of communication designed to manage attentional demand and initiative, people will be more accepting of novel interactions patterned off of existing interactions. When we design interactive products that will live in the midst of everyday people, we need to be able to read the messages that these people—drivers, pedestrians, cyclists, the drivers of other cars, construction crews—are giving off with their bodies and their behaviors. We need to communicate in ways that are equally sensible to others, so that actions are coordinated and accidents are avoided. With all of the actors that the car needs to interact with simultaneously, mastering the explicit messaging that has typified human-computer interaction in the realm of personal computers is not sufficient. Implicit interactions will be crucial as autonomous cars make their way in the world at large.

6.3 THE RIGHT ANALOGUE IS KEY

In [Chapter 3](#), we looked at using analogues to adapt interaction solutions from one domain to another. One of the dominant metaphors in the autonomous vehicle space is what is called the H-metaphor ([Flemisch et al., 2003](#); [Norman, 2007](#)). “H” stands for horse, and the idea behind the H-metaphor is that people riding horses collaborate and communicate with the horse to travel through the world. Riders sometimes ride with loose reins, indicating only high-level directions to

the horse, who is then permitted to make decisions about speed and obstacles as is necessary. Other times, riders ride with tight reins, and the horse and rider are more tightly coupled; the rider is able to convey more nuance and decision making in directing the horse, and the horse is able to give more feedback about the interaction through the reins.

I would argue that the H-metaphor is not a good one for autonomous cars. The prospect of drivers being mostly uninvolved in driving but then being expected to take over direct control just at the moment of greatest danger and uncertainty seems like a recipe for disaster. Instead, based on my own experience, I offer another H-metaphor: the husband metaphor. My husband is an autonomous driver. I trust him to drive by himself every day; I never worry once about his capabilities. However, when I sit next to him in the passenger seat, I also participate in driving. I help make decisions about where to go, and suggest alternative routes to take. I warn about potential issues and point out latent hazards that I think my husband might not see. I expect that the most important challenges in interacting with the autonomous car of the future are not the wresting of control and the tight monitoring of activity in periods of excitement or duress. Rather, I think the ordinary tasks of negotiating activities, communicating and reconciling disparate perceptions of the environment, anticipating actions and managing attention—these are more likely to be the tough challenges that we need to tackle so that the car and driver's talents combine synergistically rather than destructively.

The point here is this: when we base our interaction designs on human interactions, the analogue dictates the interaction. The choice of analogue not only drives the interaction patterns that can be used as a template for solutions, it also shapes the very questions and challenges the designers take on.

6.4 TRANSFER OF CONTROL

One thing that is not often emphasized when we talk about autonomous vehicles is that the automation is going to be introduced piecemeal into the high-end vehicles put on the market each year. These automated features then start to be included in mid- to economy-range vehicles over time. Already this is happening. There are new automated driving features introduced in cars every year: anti-lock braking, electronic stability control, adaptive cruise control, lane keeping assist, traffic jam assist, crash avoidance systems—even automatic self-parking. What drivers will experience when they take the road is that there will be more and more places and contexts where the car can step in and take the wheel.

The challenge for the near future, then, will be in transferring control to and from the driver. In [Chapter 4](#) we looked at the evolution of interactions when we transition from traditional mechanical light switches to automated and then “smart” switches. By looking at the patterns

of interaction with each kind of light switch, we unearthed many patterns, problems, as well as possible remedies.

As cars add new automated features, the car will need both to make it clear when an autonomous feature can be used (for instance, when the car is driving above a minimum speed threshold) and when its use needs to end (for example, when the freeway ends). Ideally, it would also monitor the driver to make sure they are available and attending to the car's messages. In these transitions we can imagine a wide array of possible patterns, errors, and remedies. It is also important for designers to think about dark patterns, and whether there is any way to avoid them. For instance, the FBI has added autonomous cars to its watch list, because the agency is concerned about autonomous cars being used in possible terrorist attacks (Harris, 2014).

6.5 WATCH OUT FOR PEOPLE

My current research foray into autonomous cars owes a great debt to the late Professor Clifford Nass, who was an advisor and mentor of mine. One of my favorite Cliff stories was about a trip he took in an experimental autonomous vehicle down California's Highway 280 in the Bay Area. At one point, seemingly out of nowhere, the car suddenly braked. Cliff was incensed. There was no obstacle in front of them, no dog running across the road—the car obviously made a mistake! After his trip, he reported the incident to one of the engineers responsible for programming the car. The engineer looked over the car logs, and said, "Oh, no, the car did what it was supposed to. It was staying out of the blind-spot of a car in the next lane." So, on one hand, the car was very clever about driving; statistically it is dangerous to drive in another car's blind spot, and so it avoided that potentially dangerous situation. On the other hand, it didn't communicate why it had taken this action to Cliff. If Cliff had been driving, he would likely have overridden the car because it wasn't driving in a way that made sense to Cliff.

Although I have emphasized many of the more social and communicative aspects of designing interactions, it is still very difficult for machines to make out much of what we people are doing, and why. One of the challenges in communicating and developing common ground with machines is that we have profoundly different world views. It is both hard for people to understand what machines are looking at, and thinking about, and vice versa. In the realm of research frontiers for interaction design is understanding the things that matter to people—like proxemics, consistency, and intention—and factoring those into deciding how to act and react. Adapting to these factors often involves additional sensing and actuation technology that is dedicated toward interaction. These technologies often seem superfluous to traditional engineers, but are actually critical when we are working with people.

6.6 KNOW WHEN TO IGNORE THE RULES

In designing autonomous vehicles, we often have to consider life and death situations that are less likely in designing more house- or office-bound objects. Because of this, a lot of the design principles that have been introduced to make usable products or enjoyable software need to be questioned. For example, in designing interactive characters in games or for interfaces, we usually want to make characters likeable and trustworthy. The more likeable or trustworthy, the better. In cars, however, the designer needs to worry: if a car seems too likeable will it keep the driver from taking over at a critical moment, lest it be rude? If it seems too trustworthy, will drivers second-guess their own instincts and decision making at critical junctures? Maximizing consumer preference is not enough. Maybe it's important here to estimate the car's actual confidence and trustworthiness, and calibrate its actions accordingly.

Sometimes it can be tempting for design theorists to provide a list of rules or principles to design by. I will refrain. For the better part of this book, we have encouraged the reader to follow examples and patterns, to suss out social rules and follow them. But in the case of design, I say: break the rules. A writer's imagination is limited; reality, far less so. Sometimes designers need to make machines that are scary or threatening. Sometimes it is important to express uncertainty, or undermine trust. Obediently internalizing design truisms also blinds a designer to important possibilities. Instead I will exhort readers: explore. In this book I have emphasized how interactions follow social patterns, but it is also important to know that social patterns are being renegotiated and reinvented every day. Societies develop rituals and patterns to suit our ever-changing contexts, and as designers, we should lead the way in interrogating what is possible on the way to discovering what is desirable.

Bibliography

Alexander, C. A. (1977). *Pattern Language: Towns, Buildings, Construction*. Oxford University Press. [34](#)

Andre, A. D., Wickens, C. D., and Moorman, L. (1991). Display Formatting Techniques for Improving Situation Awareness in the Aircraft Cockpit. *The International Journal of Aviation Psychology*, 1(3): pp. 205–218. [DOI: 10.1207/s15327108ijap0103_2](#). [9](#)

Antonelli, P. (Ed.) (2011). *Talk to Me: Design and the Communication between People and Objects*. The Museum of Modern Art. [62](#)

Argyle, M. (1988). *Bodily Communication*. London: Methuen. [59](#)

Bakker, S., van den Hoven, E., and Eggen, B. (2010). Design for the Periphery. *EuroHaptics 2010*, 71. [10](#)

Baum, L. F. (1900). *The Wonderful Wizard of Oz*. George M. Hill Co. [28](#)

Bellotti, V., Back, M., Edwards, W. K., Grinter, R. E., Henderson, A., and Lopes, C. (2002). Making sense of sensing systems: five questions for designers and researchers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (pp. 415–422). ACM. [DOI: 10.1145/503447.503450](#). [38](#)

Bellotti, V. and Edwards, K. (2001). Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Human–Computer Interaction*, 16(2–4): pp. 193–212. [DOI: 10.1207/S15327051HCI16234_05](#).

Bly, S. A. (1988). A use of drawing surfaces in different collaborative settings. In *Proceedings of the 1988 ACM conference on Computer-supported cooperative work*. ACM. pp. 250–256. [DOI: 10.1145/62266.62286](#). [53](#)

Bohren, J., Rusu, R. B., Jones, E. G., Marder-Eppstein, E., Pantofaru, C., Wise, M., ... and Holzer, S. (2011, May). Toward Autonomous Robotic Butlers: Lessons Learned with the PR2. In *Proceedings of 2011 IEEE International Conference on Robotics and Automation (ICRA)*,. IEEE. pp. 5568–5575. [DOI: 10.1109/ICRA.2011.5980058](#). [56](#)

Borchers, J. (2001). *A Pattern Approach to Interaction Design*. New York: John Wiley and Sons, Inc. [34](#)

Brignull, H. (2011). Dark Patterns: Deception vs. Honesty in UI Design. *Interaction Design, Usability*, Issue 338. [48](#)

Buxton, W. (1995). Integrating the Periphery and Context: A New Model of Telematics. *Proceedings of Graphics Interface '95*, pp. 239–246. [9](#)

Cakmak, M., Srinivasa, S. S., Lee, M. K., Forlizzi, J., and Kiesler, S. (2011). Human Preferences for Robot-Human Hand-Over Configurations. In *Proceedings of 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. pp. 1986–1993. [DOI: 10.1109/IROS.2011.6094735](#). [58](#)

Cassell, J. and Thorisson, K. (2000). The Power of a Nod and a Glass: Envelope vs. Emotional Feedback in Animated Conversational Agents. *Jnl. Applied A.I.* 13(4): pp. 519–538. [59](#)

Chase, W. G., and Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1): pp. 55–81. [DOI: 10.1016/0010-0285\(73\)90004-2](#). [10](#)

Cherubini, M., Venolia, G., DeLine, R., and Ko, A. J. (2007). Let's Go to the Whiteboard: How and Why Software Developers use Drawings. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM. pp. 557–566. [DOI: 10.1145/1240624.1240714](#).

Clark, H. H. (1996). *Using Language*. Cambridge: Cambridge University Press, p. 92. [DOI: 10.1017/CBO9780511620539](#). [10](#)

Clark, H. H. (2003). *Pointing and Placing. Pointing: Where Language, Culture, and Cognition Meet*, Hillsdale NJ: Erlbaum, pp. 243–268. [9](#)

Cooperstock, J. R., Fels, S. S., Buxton, W., and Smith, K. C. (1997). Reactive Environments. *Communications of the ACM*, 40(9): pp. 65–73. [DOI: 10.1145/260750.260774](#). [32](#)

Cross, Nigel (2001). Design Cognition: Results from Protocol and Other Empirical Studies of Design Activity. In: Eastman, C.; Newstatter, W. and McCracken, M. eds. *Design Knowing and Learning: Cognition in Design Education*. Oxford, UK: Elsevier, pp. 79–103. [DOI: 10.1016/B978-008043868-9/50005-X](#). [52](#)

Czerwinski, M., Cutrell, E., and Horvitz, E. (2000). Instant Messaging: Effects of Relevance and Timing. In *People and computers XIV: Proceedings of HCI*. Vol. 2, pp. 71–76. British Computer Society. [10, 47](#)

Dahlback, N., Jonsson, A., and Ahrenberg, L. (1993). Wizard of Oz Studies—Why and How. *Knowledge Based Systems*, 6(4): pp. 258–266. [DOI: 10.1016/0950-7051\(93\)90017-N](#). [28](#)

Dennett, D. C. (1989). *The Intentional Stance*. MIT Press. [31](#)

Ekman, P., and Friesen, W. V. (1974). Detecting Deception from the Body or Face. *Journal of Personality and Social Psychology*, 29(3): 288. [DOI: 10.1037/h0036006](#). [59](#)

Field, S. (1994). *Screenplay*. New York: Bantam Doubleday Del. [7](#)

Fiske, S. and Taylor S. (1991). *Social Cognition*. New York: McGraw-Hill. [59](#)

Flemisch, F. O., Adams, C. A., Conway, S. R., Goodrich, K. H., Palmer, M. T., and Schutte, P. C. (2003). The H-Metaphor as a Guideline for Vehicle Automation and Interaction. NASA Technical Report 20040031835. [64](#)

Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., ... and Yang, J. (2005). Predicting Human Interruptibility with Sensors. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 12(1): pp. 119–146. [DOI: 10.1145/1057237.1057243](#). [47](#)

Gill, S. P., and Borchers, J. (2008). Knowledge in Co-Action: Social Intelligence in Collaborative Design Activity. In *Cognition, Communication and Interaction*, pp. 38–55. Springer London. [DOI: 10.1007/978-1-84628-927-9_3](#). [52](#)

Goffman, E. (1967). On Face-Work. *Interaction Ritual*, pp. 5–45. [47](#)

Greenberg, S., Boring, S., Vermeulen, J., and Dostal, J. (2014). Dark Patterns in Proxemic Interactions: A Critical Perspective. *Designing Interactive Systems* (DIS 2014). [DOI: 10.1145/2598510.2598541](#). [48](#)

Groom, V., Nass, C., Chen, T., Nielsen, A., Scarborough J. K., Robles, E. (2009). Evaluating the Effects of Behavioral Realism in Embodied Agents. *International Journal of Human-Computer Studies*, 67(10): pp. 842–849. [DOI: 10.1016/j.ijhcs.2009.07.001](#). [60](#)

Hall, E. T. (1969). *The Hidden Dimension* (Vol. 1990). New York: Anchor Books. [52](#)

Harris, M. (2014). FBI Warns Driverless Cars Could Be Used as ‘Lethal Weapons.’ The Guardian. com, 16 July 2014. [66](#)

Heidegger, M. (1962). *Being and Time* (J. Macquarrie and E. Robinson, Trans.). New York: Harper and Row. [37](#)

Horvitz, E., Breese, J., Heckerman, D., Hovel, D., and Rommelse, K. (1998). The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. In *Proceedings of the Fourteenth conference on Uncertainty in Artificial intelligence*, pp. 256–265. Morgan Kaufmann Publishers Inc. [28](#)

Horvitz, E. (1999). Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM. pp. 159–166. [DOI: 10.1145/302979.303030](#). [10](#)

Horvitz, E., Kadie, C., Paek, T., and Hovel, D. (2003). Models of Attention in Computing and Communication: From Principles to Applications. *Communications of the ACM*, 46(3): pp. 52–59. [DOI: 10.1145/636772.636798](#). [10](#)

Höysniemi, J., Hämäläinen, P., and Turkki, L. (2004). Wizard of Oz Prototyping of Computer Vision Based Action Games for Children. In *Proceedings of the 2004 Conference*

on Interaction Design and Children: Building a Community. ACM. pp. 27–34. DOI: [10.1145/1017833.1017837](https://doi.org/10.1145/1017833.1017837). 28

Hudson, S., Fogarty, J., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. and Yang, J. (2003). Predicting Human Interruptibility with Sensors: A Wizard of Oz Feasibility Study. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM. pp. 257–264. DOI: [10.1145/642611.642657](https://doi.org/10.1145/642611.642657). 28

Isbister, K., and Nass, C. (2000). Consistency of Personality in Interactive Characters: Verbal Cues, Non-Verbal Cues, and User Characteristics. *International Journal of Human-Computer Studies*, 53(2): pp. 251–267. DOI: [10.1006/ijhc.2000.0368](https://doi.org/10.1006/ijhc.2000.0368). 59

Ju, W., Ionescu, A., Neeley, L., and Winograd, T. (2004). Where the Wild Things Work: Capturing Shared Physical Design Workspaces. In *Proceedings of the 2004 ACM Conference on Computer-Supported Cooperative Work.* ACM. pp. 533–541. DOI: [10.1145/1031607.1031696](https://doi.org/10.1145/1031607.1031696).

Ju, W., Neeley, L., Winograd, T., Leifer, L. (2006). “Thinking with Erasable Ink: Ad-hoc Whiteboard Use in Collaborative Design.” CDR Technical Report #20060928. 52

Ju, W., Lee, B. A., and Klemmer, S. R. (2008). Range: Exploring Implicit Interaction through Electronic Whiteboard Design. In *Proceedings of the 2008 ACM Conference on Computer-Supported Cooperative Work.* ACM. pp. 17–26. DOI: [10.1145/1460563.1460569](https://doi.org/10.1145/1460563.1460569).

Ju, W., and Takayama, L. (2009). Approachability: How People Interpret Automatic Door Movement as Gesture. *International Journal of Design*, 3(2): pp. 1–10. 27

Klemmer, S. R., Sinha, A. K., Chen, J., Landay, J. A., Aboobaker, N., and Wang, A. (2000). Suede: A Wizard of Oz Prototyping Tool for Speech User Interfaces. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology.* ACM. pp. 1–10. DOI: [10.1145/354401.354406](https://doi.org/10.1145/354401.354406). 28

Kraut, R. (1978). Verbal and Nonverbal Cues in the Perception of Lying. *Jnl. Pers. and Soc. Psych.* 36, pp. 380–391. DOI: [10.1037/0022-3514.36.4.380](https://doi.org/10.1037/0022-3514.36.4.380). 59

Kraut, R., Fish, R., Root, R. and Chalfonte, B. (1990). Informal Communication in Organizations: Form, Function, and Technology. In R. Baecker (Ed.) *Readings in Groupware and CSCW: Assisting Human-Human Collaboration*, pp. 130–144. San Francisco: Morgan Kaufman. 59

Laurel, B. (1993). *Computers as Theatre*. 1991. New York: Addison-Wesley. 7

Laurel, B. (Ed.). (2003). *Design Research: Methods and Perspectives*. MIT Press.

Lu, J., Sookoor, T., Srinivasan, V., Gao, G., Holben, B., Stankovic, J., ... and Whitehouse, K. (2010). The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *Pro-*

ceedings of the 8th ACM Conference on Embedded Networked Sensor Systems. ACM. pp. 211–224. DOI: [10.1145/1869983.1870005](https://doi.org/10.1145/1869983.1870005). 46

Lupton, E. (2004). *Thinking with Type: A Critical Guide for Designers*. Princeton Architectural Press. 9

Marquardt, N., and Greenberg, S. (2012). Informing the Design of Proxemic Interactions. *IEEE Pervasive Computing*, 11(2): pp. 14–23. DOI: [10.1109/MPRV.2012.15](https://doi.org/10.1109/MPRV.2012.15). 55

Matthews, T., Dey, A. K., Mankoff, J., Carter, S., and Rattenbury, T. (2004). A Toolkit for Managing User Attention in Peripheral Displays. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*. ACM. pp. 247–256. DOI: [10.1145/1029632.1029676](https://doi.org/10.1145/1029632.1029676). 9

Maulsby, D., Greenberg, S., and Mander, R. (1993). Prototyping an intelligent agent through Wizard of Oz. In *Proceedings of the INTERACT'93 and CHI'93 Conference on Human Factors in Computing Systems*, (pp. 277–284). ACM. DOI: [10.1145/169059.169215](https://doi.org/10.1145/169059.169215). 28

Markoff, J. and Miller, C.C. (2014). Danger: Robots Working. June 17, 2014. *New York Times*, Page D1. 56

McKim, R. H. (1972). *Experiences in Visual Thinking*. Boston:Prindle, Weber and Schmidt. Duxbury Press. 52

McNeill, D. (2005). *Gesture and Thought*. University of Chicago Press. DOI: [10.7208/chicago/9780226514642.001.0001](https://doi.org/10.7208/chicago/9780226514642.001.0001). 30

Nass, C., and Lee, K. M. (2001). Does Computer-Synthesized Speech Manifest Personality? Experimental Tests of Recognition, Similarity-Attraction, and Consistency-Attraction. *Journal of Experimental Psychology: Applied*, 7(3): 171. DOI: [10.1037/1076-898X.7.3.171](https://doi.org/10.1037/1076-898X.7.3.171). 59

National Highway Traffic Safety Administration. (2010). The Economic and Societal Impact of Motor Vehicle Crashses, DOT HS 812 013, May 2014, pp. 5–6. 63

Norman, D. A. (2002). *The Design of Everyday Things*. Basic Books. 38

Norman, D. A. (2007). *The Design of Future Things*. Basic Books. 64

Pantic, M., and Patras, I. (2006). Dynamics of Facial Expression: Recognition of Facial Actions and their Temporal Segments from Face Profile Image Sequences. In *Proceedings of 2006 IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 36(2): pp. 433–449. DOI: [10.1109/TSMCB.2005.859075](https://doi.org/10.1109/TSMCB.2005.859075). 10

Polanyi, M. (1967). *The Tacit Dimension*. Doubleday and Co. 10

Poppel, H. L. (1981). Who Needs the Office of the Future? *Harvard Business Review*, 60(6): pp. 146–155. 63

Rauch, J. (2011). The Origin of Slow Media: Early Diffusion of a Cultural Innovation through Popular and Press Discourse, 2002–2010. *Transformations*, 20.

Rasmussen, J. (1979). *On the Structure of Knowledge: A Morphology of Mental Models*. (Riso-M-2192) Riso National Laboratory, Roskilde Denmark: Electronics Department. 34

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. 7, 51

Robinson, T., Chan, E., and Coelingh, E. (2010). Operating platoons on public motorways: An introduction to the SARTRE platooning programme. In *17th World Congress on Intelligent Transport Systems*, (pp. 1–11). 64

Salber, D. and Coutaz, J. (1993). Applying the Wizard of Oz Technique to the Study of Multi-modal Systems. In *Human-Computer Interaction*, pp. 219–230. Springer Berlin Heidelberg. DOI: 10.1007/3-540-57433-6_51. 28

Salber, D., Dey, A. K., and Abowd, G. D. (1999). The Context Toolkit: Aiding the Development of Context-Enabled Applications. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM. pp. 434–441. DOI: 10.1145/302979.303126. 10

Schmidt, A. (2000). Implicit Human Computer Interaction through Context. *Personal Technologies*, 4(2–3): pp. 191–199. DOI: 10.1007/BF01324126. 10

Shoham, Y. (1993). Agent-Oriented Programming. *Artificial Intelligence*, 60(1): pp. 51–92. DOI: 10.1016/0004-3702(93)90034-9. 31

Sirkin, D. and Ju, W. (2012). Consistency in Physical and On-screen Action Improves Perceptions of Telepresence Robots. *Proceedings of the 7th International Conference on Human-Robot Interaction (HRI 2012)*, Boston, MA. pp. 57–64. DOI: 10.1145/2157689.2157699. 59

Suchman, L. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. New York: Cambridge University Press. 7, 21

Syrdal, D. S., Otero, N., and Dautenhahn, K. (2008). Video Prototyping in Human-Robot Interaction: Results from a Qualitative Study. In *Proceedings of the 15th European Conference on Cognitive Ergonomics: The Ergonomics of Cool Interaction*. ACM. Article 29. DOI: 10.1145/1473018.1473055. 29

Tang, J., and Leifer, L. (1989). Observations from an Empirical Study of the Workspace Activity of Design Teams. In *Proceedings of the First International ASME Conference on Design Theory and Methodology*. ASME. 52, 53

Takayama, L., Dooley, D., and Ju, W. (2011). Expressing Thought: Improving Robot Readability with Animation Principles. In *Proceedings of the 6th International Conference on Human-Robot Interaction*. ACM. pp. 69–76. DOI: [10.1145/1957656.1957674](https://doi.org/10.1145/1957656.1957674). 57

Tennenhouse, D. (2000). Proactive Computing. *Communications of the ACM*, 43(5): pp. 43–50. DOI: [10.1145/332833.332837](https://doi.org/10.1145/332833.332837). 10

Thimbleby, H. (2010). *Press On: Principles of Interaction Programming*. MIT Press. 33

Thomas, F., Johnston, O., and Rawls, W. (1981). *Disney Animation: The Illusion of Life* (Vol. 4). New York: Abbeville Press.

Thomas, J. C. (1976). *A Method for Studying Natural Language Dialogue* (No. RC-5882). IBM Thomas J Watson Research Center, Yorktown Heights, NY. 28

Verplank, W. (2003). *Interaction Design Sketch Book*. Unpublished paper. Available on-line at www.billverplank.com/IxDSketchBook.pdf. 8

Vogel, D. and Balakrishnan, R. (2004). Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. *Proc. ACM UIST*, pp. 137–146. DOI: [10.1145/1029632.1029656](https://doi.org/10.1145/1029632.1029656). 55

Watzlawick, P., Bavelas, J. B., and Jackson, D. D. (1967). *Pragmatics of Human Communication: A Study of Interactional Patterns, Pathologies and Paradoxes*. W.W. Norton and Company. 61

Weiser, M. (1991). The Computer for the 21st Century. *Scientific American*, 265(3): pp. 94–104. DOI: [10.1038/scientificamerican0991-94](https://doi.org/10.1038/scientificamerican0991-94). 55

Williams, A. (2012). *A Meta-Analysis of Energy Savings from Lighting Controls in Commercial Buildings*. Leukos. 46

Wisneski, C., Ishii, H., Dahley, A., Gorbet, M., Brave, S., Ullmer, B., and Yarin, P. (1998). Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information. In *Cooperative Buildings: Integrating Information, Organization, and Architecture*. Springer Berlin Heidelberg. pp. 22–32. DOI: [10.1007/3-540-69706-3_4](https://doi.org/10.1007/3-540-69706-3_4). 9

Wooldridge, M., and Jennings, N. R. (1995). Intelligent Agents: Theory and Practice. *The Knowledge Engineering Review*, 10(2): pp. 115–152. DOI: [10.1017/S0269888900008122](https://doi.org/10.1017/S0269888900008122).

Author Biography



Wendy Ju is Executive Director for Interaction Design Research at Stanford University's Center for Design Research, and an Associate Professor in the Graduate Design Program at the California College of the Arts in San Francisco. Her research is primarily focused on the design of interactive devices, particularly human-robot interaction and autonomous car interfaces.

Wendy has played a broad and long-standing role in the Bay Area design research community. She was the founder and editor-in-chief of *Ambidextrous Magazine*, Stanford's journal of design. She frequently collaborates and consults with local technology companies on interaction technologies, human interaction experiments, and enabling research. Wendy has been instrumental in creating hands-on interaction design curriculum for K-12 students, undergraduate, and graduate programs at Nueva School, Berkeley's College of Environmental Design, Stanford's Electrical Engineering Department, and Stanford's Music Department. Wendy has a Ph.D. from Stanford University in Mechanical Engineering, and a Master's Degree from the MIT Media Lab.