

Designing Personalization in Technology-Based Services

by

Min Kyung Lee

September, 2013

CMU-HCII-13-102

Human-Computer Interaction Institute

School of Computer Science

Carnegie Mellon University

Pittsburgh, Pennsylvania 15213

Thesis Committee:

Jodi Forlizzi (Co-chair), Carnegie Mellon University

Sara Kiesler (Co-chair), Carnegie Mellon University

John Zimmerman, Carnegie Mellon University

Leila Takayama, Google

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

Copyright © 2013 Min Kyung Lee. All rights reserved.

This work was supported by the National Science Foundation (IIS-0121426, CRI-0709077),
Microsoft Research, and the author's KwanJeong Lee Jong Hwan Scholarship and Siebel
Scholarship. Any opinions, findings, conclusions, or recommendations expressed in this material
are those of authors and do not necessarily reflect those of the funding agencies.

Keywords

Personalization, Technology-Based Services, Service Design, Human-Robot Interaction, Human-Computer Interaction, Computer-Supported Cooperative Work, Long-Term Interaction, Cooperation, Rapport, Engagement, Script, Schema, Service Breakdowns, Recovery Strategies, Breakdown Mitigation Strategy, Personal Service Providers, Adaptive Service Design, Co-Creation, Human-Centered Design, Interaction Design

Acknowledgements

It takes ten years to have an idea really one's own — about which one can talk.

The Wind at Djemila by Albert Camus

My passage to crafting my own idea and voice has not been a solitary individual endeavor, but rather one accompanied by many mentors, friends, and family members. My Ph.D. advisors, Sara Kiesler and Jodi Forlizzi, provided essential guidance, support, and encouragement throughout my intellectual journey. Sara Kiesler has exemplified how vigorous curiosity and open-mindedness can push boundaries and transform diverse perspectives into beautiful, novel ideas. Jodi Forlizzi has been an example of strength and determination, bringing design principles into human-computer interaction and transforming design as an academic field. Thank you for being my role models.

John Zimmerman, one of my thesis committee members, has prompted a critical reflection on my work, honing my contribution and making it much more interesting. He once offered a research opportunity to a budding student who had just arrived at CMU. It was my first experience with research and played a critical role in my decision to pursue the Ph.D. program. Thank you.

Leila Takayama, another thesis committee member, has encouraged me to strive to make an impact on both theory and practice. She created a place for me at Willow Garage; since then, she has been incredibly helpful and has shown me what it means to be a powerful yet supportive leader. Thank you.

Several scholars, many of them not cited in my dissertation, deeply influenced my intellectual growth and brought joy to my research career. George Loewenstein, Richard Buchanan, Donald Schon, Wanda Orlikowski, Robert Sutton, Jennifer Aaker, Daniel Kahneman, Amos Tversky, and Herbert Simon: Thank you.

All of my projects were possible thanks to great collaborators such as Siddhartha Srinavasa, Paul Rybski, Reid Simmons, Pamela Hinds, Maxim Makatchev, Kyle Strabalasa, and Anca Dragan. Being in the HCI Institute, I also had some wonderful opportunities for interacting with and receive advice from many bright individuals. Jason Hong, Anind Dey, Matthew Kam, Aniket Kittur, Robert Kraut, Jennifer Mankoff and Haakon Faste: Thank you.

I got through the ups and downs of the Ph.D. program only thanks to my marvelous friends and colleagues at CMU. Tawanna Dillahunt and I shared many laughs and anguishes. Brian Lim took pictures of our cohort, creating permanent reminders of many cherished memories. Bryan Pendleton always made sure our cohort did not starve like graduate students, ate healthily, and slept well. Thank you.

Scott Davidoff introduced me to American dual-income families' lives, New York, American Sci-Fi novels, Citizen Cope, and the charms of Pittsburgh. Turadg Aleahmad, Gahgene Gweon, Gary Hsieh, Soojin Jun, SeungJun Kim, Sunyoung Kim, Matthew Lee, Johnny Lee, Joonhwan Lee, Ian Li, Bilge Mutlu, Peter Scupelli, and Karen Tang have always been quick to offer a hand, and were all good friends and mentors. Thank you.

Jennifer Marlow, Ruogu Kang, Chloe Fan, Rebecca Gulotta, Anthony Chen, Gabriela Marcu, and Sauvik Das are energetic individuals who made my daily life in the HCI Institute colorful. Julia Schwartz, Suyoun Kim, and Stephen Oney helped me keep going strong with writing routines. Queenie Kravitz created an environment where all students, including myself, could find home at work. David Casillas, Kevin Topoloski, and Mark Penney made sure everything went smoothly. Thank you.

While our interactions were brief, I still remember kind words and advice from these genuine individuals: Cristen Torrey, Justin Weisz, Andrew Ko, Moira Burke,

Amy Hurst, Aruna Balakrishnan, Jeff Wong, and Elijah Mayfield. Their words were timely, and it kept me going at critical moments. Thank you.

Miso Kim and I shared many long Pittsburgh days and nights working together and talking about a wide variety of topics. Esther Ahn and I helped each other stay awake for many sleepless nights as master's students. Wooyoung Lee, Jueun Lee, Sooho Park, Takyoon Lee and Suyoun Kim helped me maintain a healthy body and mind. Thank you.

I had incredible opportunities to work with and learn from many research assistants: Stephanie Brown, Pong Sarun Savetsila, John Antanitis, Kimberley Nederlof, Andy Echenique, Sean Kim, Jane Park, Yiwen Jia, Leonard Turnier, Katherine Cuti, Mingwei Hsu, Mitchell Luban, Alex McCluskey, Ahmad Shamsuddin, Junior Baboolall, Victoria Yew, and the Snackbot development team. Thank you.

My final thanks go to my family – my parents, my brother, and Junsung Kim.

Abstract

Personalization technology has the potential to optimize service for each person's unique needs and characteristics. One way to optimize service is to allow people to customize the service themselves; another is to proactively tailor services based on information provided by people or inferred from their past behaviors. These approaches function best when people know what they want and need, and when their behaviors and preferences remain consistent over context and time.

However, people do not always know what they want or need, and their preferences often change. In addition, people cannot always articulate their preferences with the level of detail required for customization. The customized service that they want may be suboptimal for their needs. Finally, personalized services may become obsolete as people's preferences or contexts change, unless systems can detect these changes.

This thesis recasts personalization technology to accommodate uncertainties and changes in people's preferences and goals. I study personal service providers as a model for adaptive personalization that helps people customize their services and that adjusts service according to changes in people's preferences and goals. I derive design strategies for adaptive personalization, two of which I empirically evaluate.

The first strategy adapts service interaction styles to support long-term service usage. The first two studies investigate ways to detect people's preferred interaction styles with a robotic service – whether people treat the system as a relational being or a utilitarian tool – and the efficacy of personalizing service interaction based on this interaction preference. The next study explores how the relational interactions of technology service should be personalized over time in the context of a robotic snack delivery service in a workplace. Two types of adaptive relational interaction are investigated in a longitudinal field experiment

– a social interaction strategy that adapts its conversation topic to knowledge common to an organization, and a personalized interaction strategy that learns about people over time and adapts its interactions accordingly. The results suggest that social and personalized strategies collectively improve people’s cooperation, rapport, and engagement with the service over time; the strategies also influenced social dynamics in the workplace, facilitating the adoption of a robot into an organization.

The second strategy helps people gain insight into their needs and goals when they personalize service offerings. This strategy promotes reflection, helping people think through and articulate their needs and goals. I investigate different design variables for implementing a reflective strategy for technology service. I empirically evaluate its efficacy in the context of Fitbit, a physical activity monitoring service.

This thesis makes contributions to HCI, HRI, and interaction and service design. It broadens the concept of personalization discussed in HCI and HRI; designs and evaluates adaptive personalization strategies that accommodate uncertainties and changes in people’s preferences; draws attention to the dynamic nature of people’s orientations to interactive technologies; and captures the human-centered design process of creating and implementing a robotic service.

Table of Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 14 |
| | 1.1 Thesis Problem and Approach | 16 |
| | 1.2 Thesis Contributions | 19 |
| | 1.3 Thesis Outline | 21 |
| 2 | Service Approach to Personalization | 24 |
| | 2.1 Personalization Research | 24 |
| | 2.1.1 What Is Personalization | 24 |
| | 2.1.2 Personalization Technology: User-Driven and System-Driven Personalization | 29 |
| | 2.2 Service Research | 33 |
| | 2.2.1 What Is Service | 33 |
| | 2.2.2 Technology-Based Service and Service Design Research | 34 |
| | 2.2.3 How a Service Approach Can Help | 37 |
| 3 | Adaptive Service Design | 40 |
| | 3.1 Understanding Service Experiences as Dynamic | 40 |
| | 3.1.1 Service Orientation | 40 |
| | 3.1.2 User Experience Over Time | 43 |
| | 3.2 Designing Adaptive Technology-Based Service | 46 |
| | 3.2.1 Line of Adaptivity | 46 |
| | 3.2.2 A Blueprint for A Robotic Snack Delivery Service | 47 |
| 4 | Detecting Service Orientation with Technology-Based Service | 53 |
| | 4.1 Service Provided by Autonomous Agents | 54 |
| | 4.2 Service Orientation and Interaction Scripts | 55 |

| | | |
|----------|--|-----------|
| 4.2.1 | Roboceptionist Scripts | 56 |
| 4.2.2 | Greeting as an Indicator of the Script | 58 |
| 4.2.3 | Hypotheses | 59 |
| 4.3 | Method | 59 |
| 4.3.1 | Roboceptionist | 60 |
| 4.3.2 | Data Collection and Coding | 61 |
| 4.4 | Results | 63 |
| 4.4.1 | Grounding Behavior | 64 |
| 4.4.2 | Relational Behavior | 66 |
| 4.4.3 | Conversation Topics | 69 |
| 4.4.4 | Sentence Structure | 71 |
| 4.5 | Discussion and Limitations | 71 |
| 4.6 | Design Implications | 73 |
| 4.7 | Summary | 75 |
| 5 | Matching Interaction Style to Service Orientation | 76 |
| 5.1 | Research Context: Service Breakdowns and Recovery | 76 |
| 5.2 | Mitigation Strategies | 78 |
| 5.2.1 | Expectancy-Setting Strategies | 79 |
| 5.2.2 | Recovery Strategies | 79 |
| 5.2.3 | Service Orientation | 81 |
| 5.3 | Study Design | 81 |
| 5.3.1 | Participants | 82 |
| 5.3.2 | Materials | 82 |
| 5.3.3 | Scenarios | 83 |
| 5.3.4 | Procedure | 85 |
| 5.3.5 | Measures | 85 |
| 5.4 | Results | 87 |

| | | |
|----------|---|-----------|
| 5.4.1 | Evaluation of the Robots | 87 |
| 5.4.2 | Impact of Service Breakdown | 88 |
| 5.4.3 | Impact of Expectancy-Setting (Forewarning) Strategy | 89 |
| 5.4.4 | Impact of Recovery Strategies | 90 |
| 5.4.5 | Service Orientation and Recovery | 91 |
| 5.5 | Discussion | 93 |
| 5.6 | Implications | 94 |
| 5.7 | Summary | 95 |
| 6 | Snackbot: Design of Robotic Platform and Service for Long-Term Interaction | 96 |
| 6.1 | Design Approaches to Robotic Systems | 96 |
| 6.2 | Context of Use | 98 |
| 6.3 | Design Goals | 99 |
| 6.4 | Snackbot Design Team | 100 |
| 6.5 | System Overview | 101 |
| 6.5.1 | Hardware | 101 |
| 6.5.2 | Software | 102 |
| 6.5.3 | Form | 102 |
| 6.5.4 | Interaction | 103 |
| 6.6 | Design Process | 103 |
| 6.6.1 | Needs Analysis and Service Concept Generation | 105 |
| 6.6.2 | Observation of Hot Dog Vendor's Interactions with Customers | 105 |
| 6.6.3 | Form Giving and Interaction Design | 111 |
| 6.6.4 | Second Prototype | 119 |
| 6.7 | Lessons Learned | 123 |

7 Longitudinal Field Experiment of Long-Term Personalization Strategies 125

7.1 Moving Social Agents from Labs into the Real World 125

7.2 Personalized Robotic Service Design 127

7.2.1 Personalized Service Design Goals 128

7.2.2 Service Components 129

7.2.3 Interaction Design 132

7.3 Method 135

7.3.1 Field Site 136

7.3.2 Experimental Design 136

7.3.3 Participants 137

7.3.4 Procedure 137

7.3.5 Data Sources 138

7.3.6 Measures 139

7.4 Results: Effects of Personalization 143

7.4.1 Overall Service Usage 143

7.4.2 Manipulation Check 144

7.4.3 Rapport 144

7.4.4 Cooperation 146

7.4.5 Engagement 147

7.4.6 Service Satisfaction 148

7.4.7 Discussion 148

7.5 Results: Changes in Service Orientation 150

7.5.1 Service Orientation with Snackbot 151

7.5.2 Ripple Effect: Influence of Social Context on Orientation 155

7.5.3 Discussion 160

7.6 Implications for Design 161

7.6.1 Using Personalization Strategies 162

7.6.2 Leveraging the Ripple Effect 163

7.7 Limitations 165

7.8 Summary 166

8 Understanding Personalization Process of Experienced Personal Service Providers 167

8.1 The Practices of Human Service Providers as a Model for Technology-Based Services 167

8.2 Method 169

8.2.1 Sample 169

8.2.2 Interview Protocol 171

8.2.3 Analysis 171

8.3 Findings 172

8.3.1 Collecting Client Information 172

8.3.2 Co-creating Personalized Services 175

8.3.3 Evaluating and Adapting Personalized Services 180

8.4 Design Opportunities 184

8.4.1 Defining the Goals and Scope of Personalization 184

8.4.2 Adapting Personalization over Time and Context 188

8.5 Design Space for Co-Creation of Personalized Technology-Based Services 190

8.5.1 Current Personalization Methods for Technology-Based Services 190

8.5.2 Empowering Users When Collecting Information for Personalization 191

8.5.3 Empowering Users When Creating Personalized Services 193

8.5.4 Psychological Ways of Increasing Perceived Control 194

8.6 Limitation and Future Work 194

| | | |
|----------|---|------------|
| 9 | Conclusion | 196 |
| | 9.1 Review of Thesis Contributions | 196 |
| | 9.1.1 Contributions to Personalization, HCI and HRI Research | 196 |
| | 9.1.2 Contributions to Interaction Design Research and Service Research | 198 |
| | 9.2 Ethical Discussion | 198 |
| | 9.2.1 Social Technology | 198 |
| | 9.2.2 Personalization Technology | 200 |
| | 9.2.3 Autonomous Technology and Human Labor | 200 |
| | References | 202 |

1

Introduction

Services around us are increasingly becoming technology-based, delivered by computing technology platforms. *Technology-based services*, such as airport reservation kiosks, online health record access, medical telepresence robots, and massive open online education courses, allow for efficient delivery of services that were previously provided only by human employees. With the development of intelligent ubiquitous sensing technology, new types of technology-based services have emerged as well – health-tracking and management services such as Fitbit and Healthragenous, assistive robots such as HERB, and digital wristbands in Disneyland that monitor users’ journeys through the park and influence their experience.

Personalization has taken on new importance in these technology-based services, as they offer unprecedented opportunities to optimize services for individuals (Mayor-Schonberger & Cukier, 2013; Pariser, 2011; Riecken, 2000). Unlike human service providers, who have limited attention, time, and resources, computing technology can easily keep records of user behaviors and personalize its services for each individual with affordable costs on a mass scale. For example, personal training software can create a personalized regime of feedback through a mixture of different rewards and punishments based on individual usage patterns, personalities and motivators. New initiatives in industry and academia acknowledge the importance of personalization for multiple services, such as online education (Weld, Adar, Chilton, Hoffmann, & Horvitz, 2012) and health information services (Grasso & Paris, 2011).

Previous research on personalization suggests many reasons why personalized services would be more effective than mass-produced services. For example,

previous studies found that compared to uniform service solutions, personalizing service offerings and interaction styles resulted in higher-quality outcomes: greater persuasion; increased satisfaction, liking and loyalty with service; reduced cognitive and physical loads; and more effective service outcomes such as faster task completion time (Adomavicius & Tuzhilin, 2005; Blom, 2000; Fan & Poole, 2006; Tuzhilin, 2009).

Much of this research on personalization technology has been investigated in the domains of media entertainment, information tools, and e-commerce (Shubert, Uwe, & Risch, 2006). These service domains can be characterized as preference-based. That is, people's choices about and uses of services are mainly guided by what they like and are interested in. As these are part of their everyday choices, people tend to have well-defined preferences for these services. For example, people generally know what kinds of movies that they like.

Recently, however, technology-based services are moving beyond these preferences-based domains. For example, in health and education, the efficacy of services is not just determined by whether people like the service or not, but also by whether service solutions satisfy user needs. This knowledge is often derived from domain expertise. In addition, people may not have formed preferences on these services, as they lack domain knowledge and do not deal with these choices every day.

The temporal dimension of service usage is also important in these domains. People's previous choices and service use impact their future choices. Progress for goals or ideal outcomes usually is achieved through repeated service usages, yet continued engagement might be challenging in tasks that, unlike entertainment services, are not be inherently engaging or motivating.

This shift in technology-based services gives rise to the need to investigate personalization for various service outcomes, and to design principles for

supporting uncertainties in user preferences and supporting changes in experience over time.

1.1 Thesis Problem and Approach

This thesis focuses on the problem of designing personalization for technology-based services, addressing the limitations of existing approaches for technology personalization.

Most current personalized systems use a user-driven or system-driven approach to optimize service. User-driven approaches allow people to customize the service themselves (Adomavicius & Tuzhilin, 2005; Blom, 2000; Fan & Poole, 2006).

These approaches function best when people have well-defined preferences and goals. However, people do not always know what they want; they may lack clear insights about their needs and preferences (Riquelme, 2001; Simonson, 2005). For instance, a person may look for advice on dieting, but his hidden goal may be to look more attractive. Customized service based on people's surface needs often fails to satisfy their deeper needs. People may not be able to articulate their preferences with the detail and clarity required for customization (Nisbett & Wilson, 1977; Simonson, 2005). Further, the customized service that people want may be less than optimal for their needs. They may lack domain knowledge, such as sufficient medical expertise, to weigh the value of the advice they seek, or they may have decision biases such as overconfidence (Kahneman, 2003). For example, a person who wants to use a virtual tutoring system may want to learn from an agreeable, friendly tutor; however, a strict tutor may be more suited to the student's learning, despite her preference.

System-driven approaches proactively tailor service based on information provided by people or inferred from their past behaviors. For example, a virtual tutor can adjust its own personality based on an initial survey that measures students' personalities. System-driven personalization, or implicit/intelligent

personalization, is beneficial in that it provides automatically personalized services without requiring effort on the part of the users. It has the potential to offer services that users could not have found themselves, or “personalized environments” that users find helpful, but could not have created by themselves.

System-driven approaches function best when people’s preferences and the context for people’s behaviors remain consistent. However, people’s preferences and context change over time. As people use a service repeatedly, their expectations and relationship with the service evolve. People can use the same service across different contexts. A personalized service may become obsolete unless the system detects these changes. For example, a person who uses exercise tracking software may become bored and lose interest over time as her progress in losing weight slows down, so the system may need to employ different motivation techniques and create fun experiences to sustain engagement.

As system-driven personalization becomes more prevalent, it is important to remember that this type of personalization also has downsides. It recommends what the system thinks would work best for specific individuals or contexts, but it does not do a good job of capturing changes or multi-faceted identities in user experiences. The system may therefore personalize its services incorrectly (Kramer, Noronha, & Vergo, 2000). It may also create a “filter bubble” (Pariser, 2011) or “echo chamber” (Sunstein, 2007), which exposes users only to the content they want to be exposed to, reinforcing the views that they already have and omitting opportunities to encounter new points of views, make discoveries, or experience serendipitous situations. It is unclear how users can control these system-driven learning systems: the mechanisms, methods or delivery of system-driven personalization can compromise user autonomy. In addition, this type of personalization needs to gain users’ trust without causing them to overtrust the system and stop making their own decisions.

This dissertation recasts technology personalization as a service that empowers individuals to find the solutions best suited to them, accommodating uncertainties in people's preferences and needs and dynamically supporting changes in users' needs over time. Services should empower people to discover what they want and how to address their needs, not just once but repeatedly over the course of their experience with the systems, so that the system can adapt to their progress and changes in service use over time.

In order to achieve this, we take a service-oriented approach. Service perspectives allow us to investigate personalization technology that supports uncertainties and changes because they emphasize the collaborative, continuous process of creating values for both users and providers, as opposed to focusing on products. In this thesis, we derive design principles and explore design dimensions based on the observation of personal service providers, as well as a review of service and personalization literature. We explore three principles that highlight important sub-issues in designing personalized technology-based services:

- Understanding service orientation and matching service interaction to that orientation
- Designing for long-term interaction and adapting/personalizing service in response to people's changing experiences to encourage rapport and engagement
- Empowering users to personalize services with technology providers through co-creating strategies that help them discover their preferences and gain insights into their own needs (with right level of assistance)

These principles are empirically evaluated in the context of a robotic delivery service, a physical activity tracking and management service and an anthropomorphic information kiosk. Various methods have been used including

log analysis, design experiments using scenarios, laboratory studies, and longitudinal field experiments.

1.2 Thesis Contributions

This thesis offers new design principles, knowledge about user behaviors, and research methods for designing personalized technology-based services. Taken together, these make contributions to human-computer interaction (HCI), human-robot interaction (HRI), interaction design, and service research.

I. *Design principles for personalized technology-based services.* This thesis offers design principles for creating personalized technology systems that can accommodate uncertainties and changes in user preferences and needs, expanding the definition of personalization.

a. Personalizing for service orientation: This thesis highlights the importance of users' service orientation towards technology-based services, a construct that has rarely been considered in personalization research. We investigate users' service orientation with two types of robotic services, and show that personalizing service interaction to match users' orientation can lead to more effective service outcomes.

b. Personalizing for long-term interaction: This thesis suggests that personalizing service interaction styles to adapt to user behaviors over time, rather than personalizing once in the beginning and keeping the same strategy throughout repeated interactions, can lead to improved service outcomes such as stronger rapport and engagement. This thesis designs social rapport-building strategies that adapt to individual and organizational context over time and evaluates them through a longitudinal field study.

c. Personalizing through co-creation: This thesis investigates a design space in which users and technology can co-create personalized technology-based services. Most approaches to personalization are either user-driven or system-driven; this thesis explores design dimensions for making users and systems equal participants in the personalization process, and empirically evaluates the impact of this approach on users' service experiences and motivations.

II. *Knowledge on sensemaking and adoption of intelligent technology services.*

Empirical evaluation of the design principles described above also provides information on how people make sense of agentic technologies. Existing research on anthropomorphism and the “Computers are Social Actors” paradigm emphasizes the role of individual traits in users' sensemaking of agentic objects, and treats the resulting mental models as stable qualities. In contrast, this thesis suggests that sensemaking is influenced not only by individual characteristics, but also social, organizational, and temporal context, all of which may change over time.

III. *Research methods and knowledge for designing intelligent technology services.*

- a. Conceptual framework for an adaptive service blueprint: This thesis offers a conceptual framework and blueprint for creating and scripting adaptive behaviors of technology-based services.
- b. Integration of knowledge from service and psychology research, and service design methods: Very little service design research utilizes psychology and marketing research on service and incorporates this knowledge into the design of their methods and tools. This thesis attempts to bridge the gap by creating a blueprint that incorporates service orientation and technology adoption research.

- c. Design process and examples of robotic services: This thesis is one of the first studies to apply service- and human-centered design methods to the design of a social robot and the service it offers. Most HRI research centers around the robot's physical form and interaction; less attention has been given to the design of the service it provides. This thesis illustrates a way to employ a service design approach when designing both a product and its service, taking into account multiple stakeholders in the system.
- d. Knowledge for technology-based services: Most studies on technology-based services focus on their adoption process. This thesis attempts to provide design methods for technology-based services, and examines the role that important variables in human-provided services play in technology-based services.

1.3 Thesis Outline

CHAPTER 2 begins by covering related work in personalization and service research.

CHAPTER 3 introduces design guidelines for adaptive service design. Important dimensions include two factors that services need to be sensitive to: Service orientation (CHAPTERS 4 and 5), changes in experience over time (CHAPTERS 6 and 7). The remainder of the thesis is organized into two corresponding areas, followed by CHAPTER 8 where more nuanced model of adaptive service design was investigated by studying the practices of human personal service providers. This work suggests that co-creation is another important aspect of adaptive service.

CHAPTERS 4 and 5 explore ways to detect which type of interaction people prefer, and investigate whether personalization improves service. CHAPTER 4 presents a study analyzing log data from a receptionist robot. The findings suggest that

whether people start interacting with the robot with or without a greeting will likely indicate their orientation towards the robot. CHAPTER 5 evaluates the efficacy of matching interaction style to user preference through an online experiment that explores recovery strategies for robotic service breakdowns. The findings show that tailoring a recovery strategy to one's orientation makes the strategy more effective.

CHAPTERS 6 and 7 explore the first research question in the context of a robotic snack delivery service with the goal of sustaining long-term service engagement. CHAPTER 6 introduces the design and development of a technology-based service for long-term interaction. We used a human-centered design process to design a robot platform called *Snackbot*. Informed by the way service providers personalize their interaction with clients over time, I designed an adaptive robotic service. I compared two conversation strategies: social strategies that adapt the robot's conversational topics to knowledge common to an organization and personalized strategies that help the robot learn about individuals and adapt topics to individuals' contexts.

Using this service platform, CHAPTER 7 introduces a longitudinal field experiment in a workplace tracking 21 users over two months. The overall results suggest social and personalized strategies collectively improve users' cooperation, rapport, and engagement with the service when compared to social strategies alone. These strategies also influenced social dynamics in the workplace, facilitating the adoption of the robot into the organization. The results suggest, however, that the efficacy of these strategies may depend on people's preferred interaction style with a robotic service – whether people treat the system as a relational being or a utilitarian tool. This preference shifts over time according to changes in people's mental model about the service, the relevant social norms, and the context.

CHAPTER 8 presents an ethnographic study of service providers to understand how they personalize their service offerings. The practices of personal service providers service act as a more nuanced model for adaptive personalization that helps people customize their service and that adjusts service according to changes in people’s preferences and goals.

CHAPTER 9 summarizes the contributions of this thesis and presents an ethical discussion on this work.

| <i>Chapters</i> | <i>Factors for Personalization</i> | <i>Design Strategies</i> |
|--|---|---|
| Service orientation (Chapters 4 and 5) | Utilitarian & relational orientation | Detect and match orientation |
| Sustained engagement (Chapters 6 and 7) | Cumulative experience with service Social context (Changes in orientation) | Interaction strategies that adapt to cumulative experience and social context Personalized interaction is effective in improving rapport, cooperation and engagement |
| Service for underlying goals (Chapter 8) | Uncertainty in goals and preferences | Reflection on goals Co-creation strategies |

Table 1. Overview of the chapters and how each chapter offers design principles for personalization.

2

Service Approach to Personalization

A rich and varied body of prior work in personalization and service research is pertinent to the design of personalized technology-based services. This chapter explains the background work framing this thesis, reviewing previous personalization research and exploring how service perspectives could help us investigate ways to personalize technology-based services with increasingly intelligent technology. Related work specific to a particular contribution is introduced in later chapters.

2.1 Personalization Research

Previous research in recommender systems, service and marketing research, and HCI offers a framework for classifying different types of personalization and highlights a wide array of benefits from personalization. This section describes previous research on personalization, and gaps for further research.

2.1.1 What Is Personalization?

Over the past two decades, many definitions of personalization have been proposed in the field of HCI, business, marketing and more. (Sunikka & Bragge, 2012; Tuzhilin, 2009) (Table 2). The common thread in these definitions is that personalization involves activities of tailoring technology service offerings and interaction to achieve certain outcomes (Tuzhilin, 2009). Fan and Poole (2006) provide the most comprehensive framework, which describes different design dimensions of personalization for technology-based services. These dimensions include:

- Who personalizes: Personalization can be done implicitly by systems without explicit inputs from users (system-driven) or explicitly by users (user-driven).
- To whom: Personalization can be targeted to each of different individuals or different categories of people.
- What is personalized: Content, functionality, user interface, and/or channel / information access can be personalized.

| <i>Author(s)</i> | <i>Definition(s)</i> |
|---------------------------|---|
| Peppers and Rogers (1997) | <i>Personalization</i> is customizing some feature of a product or service so that the customer enjoys more convenience, lower cost or some other benefit. Personalization can be initiated by the customer or by the firm |
| Riemer and Totz (2001) | <i>Personalization</i> (or individualization which are used synonymously) in general means matching one object's nature with one subject's needs (i.e. customize products, services, content, communications to the needs of single customers or customer groups). <i>Mass customization</i> is the individualization of products (and services) at the cost of one-size fits all |
| Blom and Monk (2003) | <i>Personalization</i> is a process that changes the functionality, interface, information content, or distinctiveness of a system to increase its personal relevance to the individual |
| Chellappa and Sin (2005) | <i>Personalization</i> refers to the tailoring of products and purchase experience to the tastes of individual consumers based upon their personal and preference information. Therefore, personalization is critically dependent on vendors' ability to acquire and process consumer information, and on consumers' willingness to share information and use personalization services |
| Ho (2006) | In <i>customization</i> , a web site provides an array of choices for the users to modify a web site's look and feel (i.e. is a user-driven process). Relevant content based on the preferences of groups of users is provided in <i>adaptation</i> (i.e. according to the country of web users). <i>Personalization</i> is a process of providing relevant content based on individual user preferences, and personalized web sites obtain preference information implicitly by tracking customer purchases or usage habits |
| Tam and Ho (2006) | There are three types of personalization: <i>user-driven personalization</i> when the user specifies in advance the desired web layout and content that matches her interests and preferences with the tools and options provided. In <i>transaction-driven personalization</i> , an online merchant generates the personalized layout and content, and thus personalization is driven by previous transactions. <i>Context-driven personalization</i> employs an adaptive mechanism to personalize content and layout for each individual user based on the context and inference of users' processing objectives in real time (e.g., product inspection versus random browsing) |

| | |
|--------------------------------------|---|
| Arora et al. (2008) | <i>Personalization</i> is a firm’s decision on the marketing mix suitable for the individual that is based on previously collected customer data. <i>Customization</i> , on the other hand, occurs when a customer proactively specifies one or more elements of her/his marketing mix |
| Kumar (2007) | <i>Personalization</i> is a limiting case of mass customization. <i>Mass customization</i> aims at a market segment of few, whereas <i>mass personalization</i> aims at a market segment of one. The degree of transformation from mass customization to mass personalization depends on the extent to which the product of a company is soft and produced electronically |
| Frias-Martinez, Chen, and Liu (2009) | There are two major approaches to <i>personalization</i> : <i>adaptability</i> that enables users to adapt the content layout and navigation support to their preferences by themselves, while <i>adaptivity</i> makes an automatic adaptation for users |
| Montgomery and Smith (2009) | <i>Personalization</i> is the adaptation of products and services by the producer for the consumer using information that has been inferred from the consumer’s behavior or transactions. <i>Personalization</i> is automated by the marketer on behalf of the customer as opposed to <i>customization</i> that a customer requests on her own behalf |

Table 2. Definitions of personalization proposed over the past decade by scholars in human-computer interaction and business. Adopted from Sunikka & Bragge, 2012.

Extending the framework suggested by Fan and Pool (2006)¹, we categorize personalization into three types (Table 3). *Architectural* personalization² refers to personalization applied to the forms and interaction styles of computational systems for individual users’ psychological benefits. *Instrumental* personalization refers to personalization applied to the functionality or usability of systems for instrumental benefits. *Relational* personalization³ refers to personalization applied to the forms and interaction styles of technology systems for social and relational benefits.

Benefits of Personalization

Empirical studies in psychology, HCI, management of information systems (MIS), business and marketing support a wide array of benefits that personalization can

¹ Fan and Pool (2006) suggested four “ideal types” of personalization, but this thesis excludes “commercial personalization” as it overlaps with the rest of the categories nowadays.

² The definition of “architectural personalization” is extended to include features beyond the customization of the look and feel of website designs.

³ The definition of “relational personalization” is extended to include relationships between service providers and users/customers, in addition to relationships among users.

bring to both users and service providers. Personalization has been shown to improve many critical outcomes in services including satisfaction, loyalty, motivation, efficiency, learning gains, attention, memory, and motivation and others (Table 4). For example, personalization, whether applied to online content recommendations or to the interaction styles of human service providers in car-repair shops, medical health care centers (Mittal & Lassar, 1996), or banking services (Surprenant & Solomon, 1987) increased people’s satisfaction with services. Further research shows that personalization of service providers’ interaction styles increased service loyalty (Ball, Coelho, & Vilarés, 2006). Matching personalities of e-commerce sites to individuals’ personalities increased people’s liking toward the services (Moon & Nass, 1996). In accordance with the theory of regulatory fit (Cesario, Grant, & Higgins, 2004), personalization also made recommendations more persuasive, increasing customers’ intentions to follow system messages (Komiak & Benbasat, 2006; Tam & Ho, 2005). Several studies shows that this is due to improved cognitive and emotional trust (Komiak & Benbasat, 2006) and increased attention caused by self-relevant information (Tam & Ho, 2006).

| <i>Types of Personalization</i> | <i>Definition(s)</i> |
|---------------------------------|--|
| Architectural Personalization | Personalization applied to the forms and interaction styles of technology, with the goal of fulfilling individuals’ cognitive and psychological needs |
| Instrumental Personalization | Personalization applied to functionalities in providing, designing, and utilizing tools with the goal of fulfilling individuals’ needs for efficiency and productivity |
| Relational Personalization | Personalization applied to the functionalities and interaction styles of technology with the goal of fulfilling a human being’s needs for socialization and a sense of belonging |

Table 3. Three ideal types of personalization adapted from the framework proposed by Fan and Poole (2006).

| <i>Personalization</i> | <i>Benefits(s)</i> | <i>Examples</i> |
|--------------------------------------|----------------------|--|
| Architectural Personalization | Attachment | Increased attachment after allowing the customization of system appearances (Belk, 1988; Blom & Monk, 2003; Sirgy, 1982; Sung, Grinter, & Christensen, 2009) |
| | Motivation | Increased motivation to study with a computational math tutor when the tutor used students' names and personal contexts (e.g., their friends' name) in the instructions (Cordova & Lepper, 1996) |
| | Attention | Increased attention caused by self-relevant information (e.g. using names or personal information) (Tam & Ho, 2006) |
| | Liking | Increased liking of recommendations when recommendations were phrased to match people's individualistic or collectivistic cultures (Kramer, Suri, & Thakkar, 2007) |
| | Satisfaction | Increased satisfaction when service providers' interactions were personable and friendly (used customers' names etc.) (Mittal & Lassar, 1996) |
| Instrumental Personalization | Efficiency | Improved task convenience and efficiency through adaptive interfaces (Gajos, Czerwinski, Tan, & Weld, 2006; Weld et al., 2003) or automation of environment personalization (Lai, Levas, Chou, Pinhanex, & Viveros, 2002) |
| | Information overload | Reduced information overload by recommending content that is likely to match people's interests on websites that offer music (Pandora), product advertisement (Amazon), news (New York Times) and search results (Teevan, Dumais, & Horvitz, 2010)) |
| Relational Personalization | Persuasiveness | Improved persuasiveness of messages and recommendations when messages used people's name and content was personalized (Komiak & Benbasat, 2006; Tam & Ho, 2005), or when persuasive messages were matched to users' persuasion profiles (Berkovsky, Freyne & Oinas-Kukkonen, 2012) |
| | Loyalty | Increased loyalty through personalization in human-provided services (Ball, Coelho, & Vilares, 2006) |
| | Trust | Improved trust, both cognitive and emotional (Komiak & Benbasat, 2006) |
| | Liking | Increased liking toward computer software or recommendations when e-commerce websites' personalities were made similar to users' personalities (Moon & Nass, 1996; Nass & Lee, 2001) |

Table 4. Benefits of personalization evaluated in empirical studies.

2.1.2 Personalization Technology: User-Driven and System-Driven Personalization

Technology-based services can be customized by users or automatically personalized by systems (Adomavicius & Tuzhilin 2005a; Fan & Poole, 2006; Tuzhilin, 2009). With the emergence of intelligent technology that utilizes people's clickstream data online, digital records such as electronic health records, and sensor data, system-driven personalization driven by systems is becoming more and more prevalent, and new research issues regarding control and user autonomy are emerging.

User-Driven Personalization

User-driven personalization, often called customization, has received much attention in research, especially after the introduction of interactive technology in late 1990s and early 2000s (Sunikka & Bragge, 2012). Interactive technologies such as websites or software allow users to customize numerous features to best fit their needs and preferences with relatively little cost (compared to hardware personalization). For example, most digital games allow people to create shortcuts, change the look and feel of the game and create customized avatars.

Many studies have focused on the effects of user-driven personalization. For example, customized avatars can influence the behavioral norms of their creators in virtual space; players often base their actions on their avatars' personalities rather than on their own identities (Ducheneaut, Wen, Yee, & Wadley, 2009; Yee, Jeremy, & Ducheneaut, 2009). Blom and Monk (2000; 2003) studied why mobile phone users customize their phones, and explained that customization allowed people to socially express their identities.

Other research streams have looked into the numerous factors that make user-driven personalization successful. Valenzuela, Dhar, & Zettelmeyer (2009) investigated algorithms that create optimized configuration of customizable

options that vary in their attributes or alternatives. Customization has been shown to appeal to people who have a high desire for stimulation (Raju, 1980) or those with need for cognition, variety seeking, and uniqueness (Ho, Davern, & Tam, 2008).

However, many studies downplay the drawbacks and boundary cases (or failures) of user-driven personalization. Some have shown that customization can be too labor-intensive and complex for users (Anderson, Hagen, Reifel, & Stettler, 2006), and that users do not always have good insights into their needs (Simonson, 2005). However, very little research investigates design principles that address these challenges.

System-Driven Personalization

System-driven personalization is personalization in which systems implicitly and often automatically tailor services based on users' observed behaviors or inputs obtained through surveys. System-driven personalization has been popular since the mid-2000s, and its popularity has increased sharply since then (Thurman & Schifferes, 2012). This increase is in part due to the development of multiple computational techniques on recommender systems, such as content-based filtering (Lawrence et al., 2001; Morita & Shinoda, 1994), collaborative filtering (Resnick et al., 1994), and user-modeling (Adomavicius & Tuzhilin, 1999; Dogac, & Azgin, 2000; Eirinaki & Vazirgiannis, 2003; Fink & Lobsa, 2000; Kim & Chan, 2003) (For a more detailed review, please refer to the following survey papers: Adomavicius and Tuzhilin 2005a; Montaner et al. 2003; Schafer et al. 2000).

These techniques use data from massive social networks or users' clickstream data to infer individuals' characteristics and personalize websites or interactive technology services, often without users' explicit efforts or control. For example, Facebook, Amazon, and Google present each individual with personalized lists of recommended friends, books, and search results, respectively. System-driven personalization is thought to reduce information overload, and is presumed to

increase the profits of many online content-providing companies (Pariser, 2011). More recently, the exponential increase in digital data sets on human behaviors and choices has allowed for “big data analytics” (Mayor-Schonberger & Cukier, 2013), expanding domains where system-driven personalization can be applied, such as entertainment, health, and education. However, to our knowledge, there is little research investigating how systems can adapt their personalization techniques over time to support people’s changing experiences (and contexts). This dissertation investigates basic principles for adaptation over time and explores the different types of system-user relationships that can unfold.

System-Driven Personalization, User Autonomy, and Control

As personalization has become more system-driven, many services are being personalized automatically without users’ direct control or, sometimes, even their awareness. This automation could interfere with user autonomy in several ways. For example, it may limit users’ exposure to diverse information by creating a “filter bubble” (Pariser, 2011), or make users feel that they lack the appropriate level of control over the systems. It is critical to understand user experiences with system-driven personalization for a number of reasons. Recent studies suggest that merely believing a system is learning about them and personalizing for them influences users’ perceptions and evaluation of the system (Nowak & Nass, 2012). Other studies suggest that, when systems proactively personalize their services and features, people may distrust the systems or feel that they do not have appropriate level of control over the systems.

Research on system-driven personalization and the broader category of intelligent, autonomous systems has investigated different ways to give users control over intelligent systems. This line of research suggests that providing explanations for intelligent systems’ behaviors (Lim, Dey, & Avrahami, 2009), improving transparency (Cramer et al., 2008), and educating people about how intelligent systems work (Kulesza, Stumpf, Burnett, & Kwan, 2012) increase people’s trust toward the systems and provide feeling of control. Kay and

Kummerfeld (2012) have researched ways to enable end-users to “scrutinize” how personalized systems function and what rules they use.

This stream of research takes an information-centric approach, assuming that providing information to help users understand how systems work will give them a feeling of control. However, very recent studies and increasingly common anecdotal evidence suggest that there is also a socio-psychological element to feeling in control when interacting with autonomously personalized systems. For example, when describing autonomously personalized ads, people use emotional terms and often anthropomorphize these ads as they take proactive behaviors. They perceive the ads as more than just tools, imbuing them instead with a sense of agency. These types of perceptions can cause people to feel judged or upset by who or what a certain system seems to “think” they are (Rao, Hurlbutt, Nass, & JanakiRam, 2009), and weaken users’ feelings of control over the system. It is unclear how much the information-centric approach can address this socio-psychological element. In addition, this line of research assumes that people will be invested enough to take the time to scrutinize the system, which is not always the case.

This thesis attempts to also look at other ways to increase people’s feelings of control overall and particularly in the socio-psychological sense, complementing the information-centric approaches.

To summarize, previous research on personalization suggests several reasons why personalization would be effective in creating successful services; yet, little research investigates knowledge about or and design principles for designing personalization strategies that can accommodate uncertainties and changes in users’ needs and preferences.

2.2 Service Research

Research on services could provide useful perspectives in exploring the design space of personalization. The following sections will explain what service is, explore how service approaches can inform the design of personalization in technology-based services, and discuss the knowledge gaps this dissertation attempts to fill.

2.2.1 What Is Service?

Service has been of interest to the business and marketing community for nearly three decades (Bitner & Brown, 2006). Increasingly more attention, both in academia and in practice, has been given to service in the 2000s, with emphases on the “experience economy” (Pine & Gilmore, 1998) and the digital technologies that transform service delivery (Chase & Apte, 2007; Chesbrough & Spohrer, 2006). Gronroos (1990a) defines service as the following.

A service is an activity or series of activities of more or less intangible nature that normally, but not necessarily, take place in interactions between the customer and service employees and/or physical resources or goods and/or systems of the service provider, which are provided as solutions to customer problems.

Service, as a view point, renders all economic activities as value creation rather than product exchange, which makes service the dominant form of economic activities (Prahalad & Ramaswamy, 2004; Vargo & Lusch, 2004; Vargo & Lusch, 2008; Vargo, Maglio, & Akaka, 2008). This perspective encourages designers to primarily focus on creating value for both customers and providers throughout use experience over time, rather than focusing on creating new product features or details.

Exchange of products has been considered the basis of the economy; thus many researchers in business and marketing have sought to define the characteristics of service by comparing them with characteristics of products. The most commonly discussed properties of services include the following (Gronroos, 1990a; Ishizaki, 2010; Vaajakallio, Mattelmäki, Lehtinen, Kantola, & Kuikkaniemi, 2009):

- **Intangibility:** While tangible goods may be included in the delivery of services, the essence of a service is the intangibility of the phenomenon itself.
- **Activities or a series of activities:** A service is not a thing but a series of activities or processes – which, moreover, are produced and consumed simultaneously.
- **Co-creation/co-production:** The customer is not only a receiver of the service; the customer participates as a production resource as well.
- **Multi-stakeholders and multiple platforms:** A service is often delivered through multi-stakeholders, including providers and customers, via multiple platforms.

2.2.2 Technology-Based Service and Service Design Research

By adopting a service perspective, this dissertation contributes to research on technology-based service and service design. Previously, most service research has studied services delivered by human providers (e.g., Bitner, 1992; Bitner, Booms, & Tetreault, 1990); technology-based services and service design have only recently become of interest to service scholars.

Research on Technology-Based Services

Technology-based services became the focus of service research in the early 2000s, as new interactive technology such as the ATM began to be used in service domains where mainly human service providers had traditionally delivered

services. Several researchers have proposed frameworks that define the different roles technology plays in service. One of these has to do with how technology supported or replaced the roles that human service providers used to play (e.g., Bitner, Brown, & Meuter, 2000; Glushko, 2010). Empirical studies have focused on investigating different factors that influence the adoption of technology-based services (Bitner, Ostrom, & Meuter, 2002; Dabholkar & Bagozzi, 2002; Meuter, Bitner, Ostrom, & Brown, 2005; Meuter, Ostrom, Roundtree, & Bitner, 2000), and offers knowledge on organizations' strategic policies regarding the introduction and operation of technology-based services (Reinders, Dabholkar, & Frambach, 2008). "Self-service technologies," such as airport check-in kiosks or online banking websites, are the most commonly investigated type of technology-based services. For example, Dabholkar and Bagozzi (2002) suggest that self-service technologies are adopted best by people who have inherent novelty seeking traits, high self-efficacy with respect to technology, high self-consciousness, and low need for interaction with an employee.

Much of service research has been focused on providing knowledge for managerial decisions about investing in technological development and effectively managing resources to support technologies. This line of research, however, provides relatively little knowledge relevant to the process of designing technology-based services. This is, in part, because of this literature's tendency to treat technology as a "black box" instead of distinguishing different features of technology and exploring the impact of the design decisions behind those features on overall service success (Griffith, 1999).

Research on Service Design

While not specific to technology-based service, service design research offers design methods and tools for services in general (Polaine, Lovlie, & Reason, 2013; Stickdorn & Schneider, 2010). Recently, designers have looked into ways to treat service as design subject, applying human-centered design methods to the design of services in the domains of retail, entertainment, healthcare, and others (e.g.,

Bitner, Ostrom, & Morgan, 2007; British Design Council; Mager & Evenson, 2008; Thomke, 2003; Zomerdijk & Voss, 2010). Some methods used for service design were directly adopted from human-centered design methods for developing and evaluating products (e.g., participant observation etc.), whereas other methods were created to address unique characteristics of services.

One of the most representative service design methods is the blueprint. Designers conceive of how a service will work through the use of a service blueprint (Bitner, Ostrom, & Morgan, 2007; Shostack, 1982). The blueprint models the process of delivering services, serving both as a sketching tool for service designers in the development phase and as a guide for service providers in the operation phase. The blueprint identifies different components of the service for both the frontstage (i.e., the parts that customers directly experience) and the backstage (i.e., the parts that providers create behind the scene), and maps the flow of events for each component.

Recently, a few changes have been made to service blueprints to better support user-centered design. For example, Morelli illustrated how a service blueprint can be used in combination with other design methods such as personas or use cases (Morelli, 2002). These methods add more contextual information about different types of users and provide scenarios of how each step of the service will take place. With these modifications, blueprints can represent information about the physical and virtual places where the service takes place, the various actors who perform the functions, and whether tasks are automated or not. Another study has proposed a visualization method to represent changes in users' emotions during a service journey (Spraragen & Chan, 2008). More recent work attempted to incorporate improvisational aspects of art, music, dance, and drama into the service blueprint (Mager & Evenson, 2008). Pinhanez (2009) proposed a new representation for a service blueprint in which all phases of the service revolve

around a customer, in order to emphasize that services are a customer-intensive systems.

Service blueprints are well suited to describing a linear flow of service experience, but do not have methods for conceptualizing and notating adaptive service behaviors over time. In addition, while much service and psychology research in marketing and organizational behavior domains provide insights into human behaviors and experiences with services, research on service design methods rarely harnesses this body of knowledge. Informed by service and psychology research, this dissertation contributes a framework (or blueprint) for notating service behaviors that adapt and adjust over time according to users' experiences.

2.2.3 How a Service Approach Can Help

Existing research on services provides helpful design perspectives and important, as yet under-examined variables to technology design and HCI (Pinhanez, 2009; Pinhanez, 2011). This dissertation introduces important these into the design of personalized technology services — service satisfaction and breakdown recovery, design for multi-stakeholders, service orientation, temporal dimensions of service usage, and co-creation.

Measuring Service Quality

Service research defines conceptual constructs of service qualities such as process qualities and outcome qualities and survey scales for measuring them (e.g., Gronroos, 1990b; Parasuraman, Zeithaml, & Berry, 1985). People's experiences with services are comprised of multiple touchpoints. Thus, the quality of service is determined by the overall experience rather than by an experience with one particular product in a set of service systems. This viewpoint lets us think of robotic technology as a service system (Chapters 6 and 7), allowing us to move beyond usability issues and measure the value of people's overall experiences with novel robotic technologies.

Designing for Multiple Stakeholders

Many different stakeholders are involved in the delivery of services, including service providers and secondary users. Service emphasizes creating value for all of these multiple stakeholders. In HCI and human-centered design, users are considered the most important stakeholders, and all design decisions are made to maximize user benefits. However, increasingly more types of technology are being transformed into service systems involving multiple stakeholders (Pinhanez, 2009; Pinhanez, 2011). This dissertation adopts the multi-stakeholder perspective to design a robotic delivery service and its personalization principles both for service providers and clients (Chapters 6 and 7). In addition to evaluating users' experiences with the service, we conducted empirical evaluations of the service using measures important for service providers, such as rapport and whether users intend to keep using the service.

Enabling Co-Creation and Co-Production

Co-creation and co-production is one of the core concepts of service (Bitner, Faranda, Hubbert, & Zeithaml, 1997; Gronroos, 2008; Guo, Arnould, Gruen, & Tang, 2013; Moeller, Ciuchita, Mahr, Odekerken-Schorder, & Fassnacht, 2013; Vargo, Maglio, & Akaka, 2008). This perspective is useful for conceptualizing personalization in technology-based services, so that the personalization process is collaborative rather than being dominated by one party. This line of research in existing service research has been mostly conceptual, and this dissertation attempts to empirically investigate this issue in technology-based services (Chapters 8 and 9).

Designing Recovery Strategies for Breakdowns

One factor that influences people's overall experiences with services is service breakdowns (Keaveney, 1995). For most services, breakdowns are inevitable, damaging users' evaluation of the service. Previous service research shows that how service providers cope with breakdowns could mitigate these negative impacts, sometimes making people appreciate services even more than they did

before the incident (Aaker, Fournier, & Brasel, 2004; Hart, Heskett, & Sasser Jr., 1990; Spreng, Harrell, & Mackoy, 1995). While most technology products and services also breakdown, recovery strategies have received little attention both in HCI and HRI. This dissertation highlights the importance of taking breakdowns into consideration in HRI, and explores different recovery strategies for robotic services. (Chapter 5).

Understanding Service Orientation

Another factor that influences people's service evaluations is people's expectations about the service (Oliver, 1980; Smith, Bolton, & Wagner, 2009). Expectation determines the baseline experience that people anticipate from using the service, and whether services exceed or fall short of this baseline determines people's satisfaction with the service. Thus, the same service can be evaluated very differently depending on individuals' expectations. One type of expectation is service orientation – the cultural models or mindsets that people hold toward services and service providers (Ringber, Odekerken-Schroder & Christensen, 2007). This dissertation investigates the role of service orientation in technology-based services, and how technology can be personalized to adapt to users' service orientations (Chapters 3, 4, 5, and 7).

3

Adaptive Service Design⁴

This chapter describes a framework for services that adapt to people's changing behavior through repeated interactions. We articulate important factors to consider in the design of an adaptive service, and illustrate how these factors work in the design of a robotic snack delivery service. Using this service as an example, we take a first step towards providing a blueprint for designing adaptive services.

3.1 Understanding Service Experiences as Dynamic

3.1.1 Service Orientation

A service orientation or schema can be understood as a mental model that people use to make sense of and evaluate service interactions. The same events can be interpreted differently depending on people's orientations. In intelligent agent services, attitudes toward both the service and the agent play a role in users' service orientations.

Mindset, Schema, and Cultural Models

A schema (or mental model) is a mental structure or representation of any object or phenomenon encountered in the world (Abelson, 1981). People apply their own schemata or mental models to make sense of events, set expectations, and guide their decisions and behaviors. Prior research provides abundant evidence that the same events can be interpreted differently depending on people's mental models. For example, one of the strongest indicators of students' success is

⁴ Part of this chapter is based on the paper published at the IASDR'09 conference (Lee & Forlizzi, 2009).

whether they have “growth mindset” or “fixed mindset” (Dweck, 2006); those who have a growth mindset exhibit more resilient behaviors after off-putting events and achieve greater improvement in academic records.

Individuals’ schemata or mental models also influence what people expect from service. Mental models in the context of service are called *service orientations*. The concept of service orientation can be used to describe the interpretive strategies that people use to incorporate particular services into their lives (Ringber, Odekerken-Schroder & Christensen, 2007). Service orientations are used to confirm existing belief systems, and to discount contradicting evidence. Different people have different orientation toward services.

In previous research, three service orientations have been described: relational, oppositional and utilitarian models (Ringber et al., 2007). The relational orientation is held by people who desire and value emotional ties with a service provider. The oppositional orientation is held by people who perceive themselves as the more vulnerable, weak player in the consumer-provider relationship and easily take an aggressive stance toward the service provider. The utilitarian orientation is held by people who rationally weigh the benefits of a service against the costs. Depending on these models, the same service can be interpreted in different ways and create either pleasant or dissonant experiences.

Previous research suggests that services should be offered to different people according to their different orientations (Aggarwal, 2004). The same type of service can be also designed to emphasize different value propositions (e.g., milk man vs. FedEx). For example, people who apply relational norms to a fitness center may expect messages from personal trainers to have a more relational tone.

Service Orientation with Technology-Based Services

Service orientations with technology-based services will likely change when technologies begin to provide services that humans used to provide. People may

have two different orientations – an orientation toward “providers” which are technological instead of human, and an orientation toward particular service types or domains.

Orientation with Intelligent Agents

People apply different levels of relationship norms to other people depending on context; for example, some people may be very polite and engage in conversation with office cleaning staff, whereas other people may “dehumanize” them (Waytz, Epley, & Cacioppo, 2010) and not acknowledge them the way they do their coworkers. The extent to which people anthropomorphize technological service providers could vary to a greater degree.

Previous research suggests that people anthropomorphize technology to different degrees, even when systems do not have anthropomorphic forms (Reeves & Nass, 1996). Even given the same external form factor, people have different orientations to agents and robots. Friedman et al. showed that some people think of AIBO, a dog-like robot, as a technological entity whereas others attribute more lifelike qualities to it (Friedman, Kahn, & Hagman, 2003). In studies of hospital delivery robots, researchers observed that some employees anthropomorphized the robot whereas others regarded the same robot as a machine (Mutlu & Forlizzi, 2008; Siino & Hinds, 2005).

The degree of anthropomorphism can influence people’s overall service orientation. For example, many people have a relational orientation toward human caregivers, but it is unclear what orientation they will hold when this service is provided by technology providers such as web-kiosks or robots. Will people treat these services as relational services as they do with human service providers, or as utilitarian services because they are offered by machines?

Previous research also suggests that the form factor of a computational agent influences people’s schemata and willingness to cooperate with the agent. In one study, participants cooperated more with an agent that looked like a person than

with an agent that looked like a dog, and more with a realistic dog agent than with a cartoon dog agent (Parise, Kiesler, Sproull, & Water, 1996). In another study, participants took less responsibility for the successful completion of the task when working with a human-like robot than a machine-like robot (Hinds, Roberts, & Jones, 2004).

To summarize, this line of research suggests that there are important individual differences in the degree to which people perceive agency in (or anthropomorphize) robotic agents; when service is provided by these agents, users' service orientations might be influenced by their orientations toward both service domains and technological service providers.

3.1.2 User Experience over Time

Two streams of research describe the temporal properties of user experiences, highlighting several ways in which the behaviors and attitudes of service users might change over time. Research on user experience design and technology adoption explains different phases that users go through when making sense of new products and incorporating (or not incorporating) these products into their routines and lives. This line of research investigates user experiences with products or services whose behaviors do not change over time. Research on relationships, on the other hand, investigates experiences over time in which two participants interact with each other, adding insights to adaptive service design.

User Experience Design and Technology Adoption

Much research on user experience (in the interaction design and HCI communities) and technology adoption has investigated the process by which people make sense of new products and technology (Weick, 1995), and incorporate (or do not incorporate) into their routines (Forlizzi, 2007; Forlizzi & DiSalvo, 2006; Forlizzi & Ford, 2000; Karapanos, Hassenzahl, & Martens, 2008; Karapanos, Zimmerman, & Forlizzi, 2009; McCarthy & Wright, 2004; Siino & Hinds, 2005; Sung, Grinter, & Christensen, 2009; Sung, Christensen, & Grinter,

2009). This line of work shows that new products have an impact not only on individuals, but also on their social and organizational contexts. In addition, it shows that which product qualities matter to users will change over time. For example, Karapanos and his colleagues studied the stability of user experience, finding that the quality of an initial experience with a product differed greatly from the quality of the long-term experience (Karapanos, Hassenzahl, & Martens, 2008; Karapanos, Zimmerman, & Forlizzi, 2009). Some of the qualities they found to be important include stimulation and beauty as people oriented to the product, usability as people incorporated the product into their lives, and emotional attachment as the product became meaningful.

The common threads of these works suggest that we can categorize the user experience into two phases: orientation phase and incorporation phase, which will be explained in §3.2.1.

Relationship Approach

Hinde's work (1976) on relationships, originating from the study of animal behaviors, provides a perspective on relationships beyond interpersonal intimacy-centered relationship models (e.g., Duck, 1991). In Hinde's model (1976), a relationship can be defined as "a succession of interactions between two individuals known to each other." One of the key characteristics of a relationship is that the nature and course of each interaction is influenced by the "history of past interactions between the individuals concerned," and also by the "expectations in the future." This perspective suggests that understanding these two components are important in conceptualizing repeated interaction as relationship.

Understanding the impact of past interactions on users' current satisfaction and expectations in future interactions is critical. A system must not only rectify errors but also communicate its intention for future interactions, especially when negative events such as system breakdowns occur. This can be done by conveying

accurate expectations explicitly (e.g., through verbal announcement) or implicitly (e.g., through character or appearance), or by employing social strategies in which the system admits its error and expresses its intention to improve.

Histories of interactions can be used to build different relationship types. Hinde (1995) offered several characteristics of relationships:

- The content of interactions
- The relative frequency and patterning of interactions
- The complementary versus reciprocal nature of a relationship
- Power and autonomy
- Intimacy
- Interpersonal perception
- Commitment
- Satisfaction

These characteristics are useful design dimensions that can inform our thinking about relationship types. For example, to build more intimate relationships with people, technology-based services can acknowledge mutual experiences. To build “coaching” relationships, histories of interactions can be used to suggest unexplored behaviors or to provide feedback.

Although exploring the full design space of temporal relationships is beyond the scope of this thesis, the relationship types discussed above can inform the design of robotic services meant to engage users for long-term.

3.2 Designing Adaptive Technology-Based Service

Our literature review on experience design and technology adoption highlighted several ways in which the behaviors and attitudes of customers might change over time as they orient to adaptive services and incorporate them into their lives.

Changes in human behavior influence how people use product and service; but in the case of technology-based services, the products and services can change, too.

With the advent of context-aware, intelligent technology, services can be designed to adapt to these changes to better support users' behavior. For example, with a traditional snack delivery service, daily interactions might change once a delivery person begins to learn the preferences and habits of customers served on the route. A robotic snacking service can record user preferences and behaviors, allowing for daily interactions to change as user preferences are learned and automated, and as new technologies are brought onboard the robot.

With this in mind, we sought to find a way to understand how a service blueprint can represent and visualize how products and services adapt to changing behavior over time. In the next section, we describe our initial blueprint for adaptive service design and the new factors within the Line of Adaptivity that need to be considered.

3.2.1 Line of Adaptivity

To represent adaptive services, we augmented a traditional service blueprint to depict repeated service encounters over time. We added a Line of Adaptivity to a traditional service blueprint, and allowed for *orientation*, *incorporation*, *streamlining*, and *personalization* in the design.

We categorize changes in people using two factors: orientation and incorporation. In the orientation phase, users frequently engage in sensemaking activities in order to understand how a service functions. When people make sense of services, they understand the functional aspects of the service, evaluate their utility and

desirability, and form initial attitudes towards the services. In traditional static services, products and services contain a fixed spirit, or schema of use, in their design (DeSanctis & Poole, 1994). These are feature sets that assume a particular physical and social context of use. People need to make sense of this schema, and either modify their own schema to fit the schema of the products and services, or reject the use of the service (Orlikowski, 2000).

In the incorporation phase, users begin to integrate products and services into their daily lives, building trust and emotional attachment. Sensemaking becomes a peripheral activity in this phase, and only takes place when aspects of the service that do not fit into their worldview are experienced.

We categorize changes in the components of a service over time using two factors: streamlining and personalization. These changes are designed to support changes in users' behaviors over time. Service offerings will evolve as users' routines and preferences are learned. Some of the touch-points in the journey may become unnecessary, requiring streamlining to continue to offer a beneficial service. Additionally, services may be personalized to better fit the needs of users, once their patterns of service purchase and service use have been learned.

To exemplify these ideas, in the next section, we describe how the Line of Adaptivity is designed in our robotic snack delivery service.

3.2.2 A Blueprint for an Adaptive Robotic Snack Delivery Service

Context of Our Research

Our chosen service domain is snack delivery. In developed countries, the majority of people eat snacks at least once a day (Ovaskainen et al., 2006). Snacking, particularly high caloric snacks, also contributes to obesity, one of the major health problems in the United States and parts of Europe. To date, very few services or design efforts have been made to assist and improve people's snacking

practices. Our research showed that people have a variety of unmet needs when choosing snacks (Lee, Kiesler & Forlizzi, 2008). For example, people desire to eat healthy snacks, yet they tend to choose unhealthy, fattening snacks due to stress and convenience.

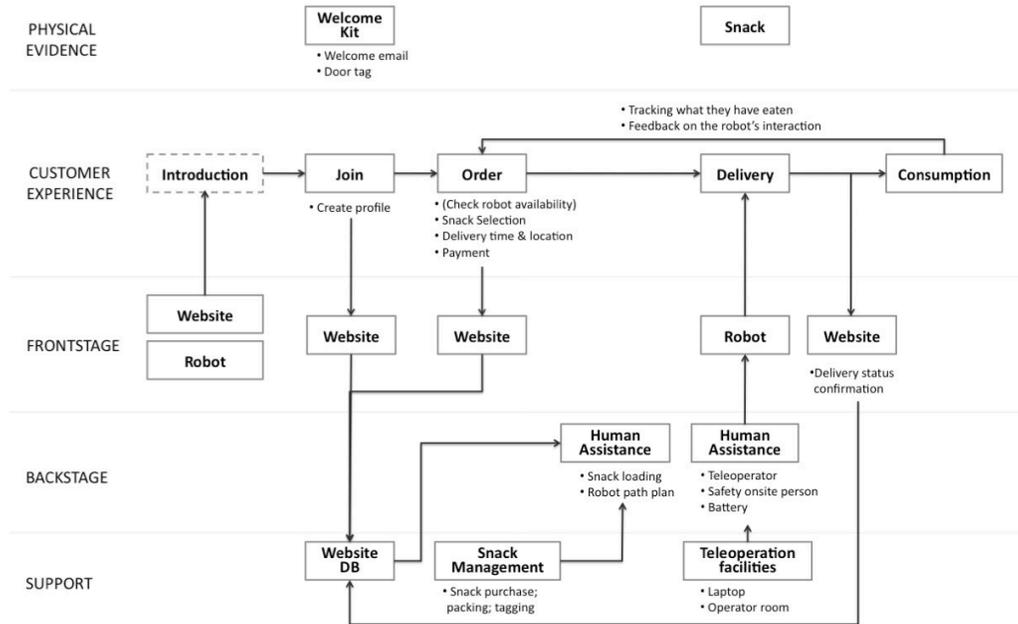


Figure 1. A blueprint of the Snackbot snack delivery service, describing a one-time journey of the snack order and delivery process.

In response to these problems, we designed a robotic snacking service using the Snackbot robot, a four and half foot tall, human-like robot that navigates the hallways semi-autonomously and delivers snacks to people (Chapter 6). The Snackbot interacts with people in a social way using natural human language, sound, and head/body movement. To order snacks from the robot, people use a website to specify the types of snacks that they want and the delivery location and time (Figure 1). Our design challenge is to design a service that satisfies changing human needs, and to sustain people’s engagement with the service over time, but also, to account for changing technology capabilities, since the robot will learn people’s preferences over time. Some of our research questions include: What are

the factors that the service should be sensitive to? How might people's experience with a service change over time? How can a service evolve to maximize people's satisfaction with it?

The Line of Adaptivity represents how people and services change with iterative use (Figure 2). In each encounter, people go through the process described in the one-time journey blueprint represented in Figure 1, making subtle adaptations in their behavior in response to the product and service, which are also changing in response to interacting with people over time.

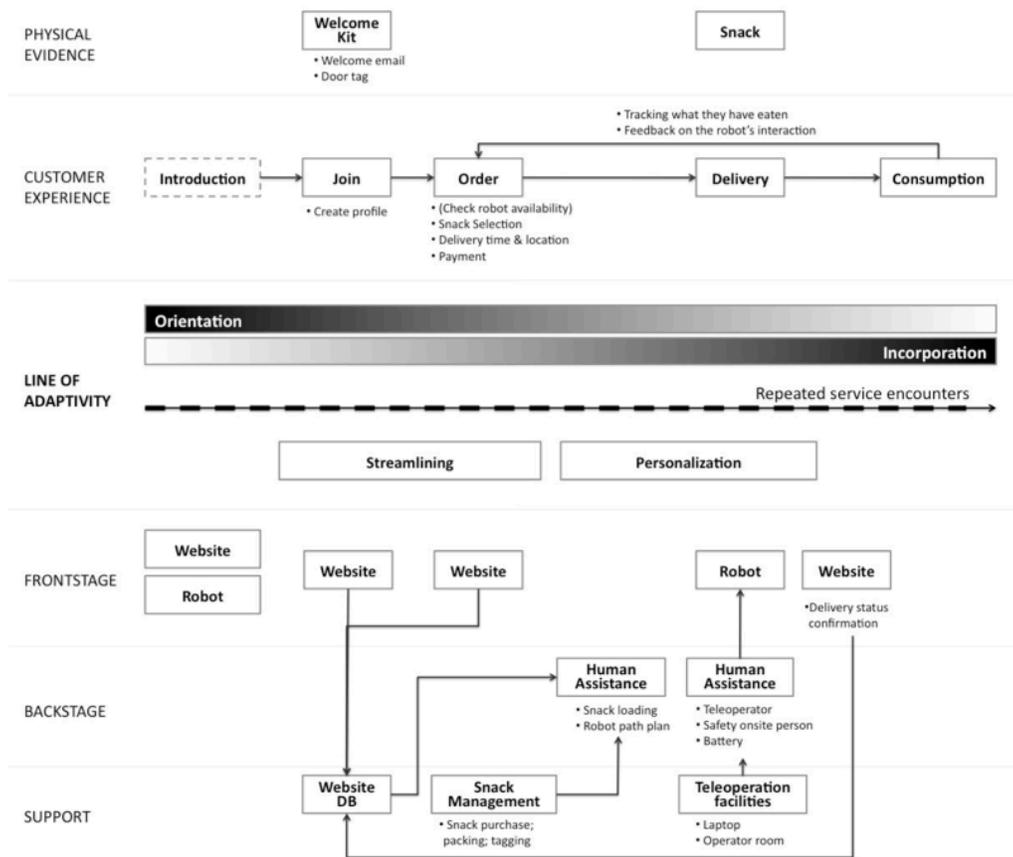


Figure 2. A service blueprint that describes how products and service can be adapted to users' changing experiences.

Orientation and Adaptive Services

In the orientation phase, people use products and services for the first time, comparing them to prior product and service experiences. When engaging with our robotic snack delivery service, the process of sensemaking will be iterative. Our snacking service will be sensitive to users' sensemaking process, thereby reducing the likelihood that users reject the service early on. The design feature that we employ to do this is a two-way interaction between users and the product and service. The robot uses simple speech recognition, and a series of cameras and lasers to capture human behaviors and respond to interaction breakdowns. The second design feature employed during orientation is incorporating human social cues into the robot's design to allow people to draw on their experience with familiar services. This can be manifested through any design feature ranging from the structure of the dialogue, for example following an interaction sequence that a human vendor employs, to humanoid cues in the robot design.

Incorporation and Adaptive Services

Our snacking service adapts to relational and utilitarian service orientations to help people incorporate the service into their lives. For people who rely on the relational service orientation when interacting with the robot, the Snackbot follows a human-human interaction model, mimicking a human vendor's interactions and using interpersonal relational strategies when delivering snacks. For example, applying the rule of reciprocity, the robot can give a free snack as a gift on a special day, which can result in feelings of thankfulness and indebtedness toward the robot. The robot can also build on what the users said in their previous meeting, or use self-disclosure strategies to create the feeling of closeness as users and the robot interact each other repeatedly. For people who rely on the utilitarian service orientation, the robot follows a more machine-like interaction sequence without trying to engage users in social conversation with the goal of delivering an efficient, minimal transaction.

To assess how to design and dynamically respond to different service orientations, we can rely on technology to understand them. For example, we can record and track how people talk to the robot. Automatically sensing the structure and form of this dialogue would allow us to infer the cultural model that is being employed at any given time. Alternatively, when a snack is ordered, users could specify the type of interaction they desire at the time of a snack delivery.

Streamlining and Adaptive Services

As users repeatedly engage with a service, they may gain a greater understanding of the service, and want to speed the process of ordering and using the service. Some of the touchpoints within might become unnecessary. Our snacking service will combine or automate some steps in the service journey for expert users. For example, speech-based instructions about how to complete the transaction with the robot might eventually become unnecessary. Instead, a simple sound that indicates the robot's actions (i.e., its arrival or an approval for taking a snack) might be informative enough and will not interrupt or distract those on the delivery route. In addition, if users exhibit same usage patterns over time, the service could automate some steps to facilitate the process. For example, a robot may automate the order process and begin to deliver snacks regularly to those who order cookies every Friday, or support one click ordering.

Personalization and Adaptive Services

In the incorporation phase, people begin to use products and services fluently, incorporate them into their daily lives, form trust and emotional attachment, and find meaning from repeated experience with the service. As users repeatedly engage with a service, the opportunity exists to personalize the service to better suit their needs. Our snacking service customizes its offerings by tracking users' preferences and customizing responses. For example, for customers who prioritize utility, the robot never engages in social conversation as it delivers snacks. Healthy snacks are suggested for those looking to increase options for a healthy lifestyle. The service can also provide proactive recommendations based

on what it has learned, by asking users if they want information about new and greater varieties of snacks. Eventually, the robot and service will be able to scrape information from users' calendars and deliver snacks for birthdays and special events.

After people make sense of services and incorporate them into their daily lives, they discover greater personal and social meaning relate to the products and services. Here, the meaning arises not from the products and services themselves, but from how they support what people value. Adaptive services can better respond to the subtle reprioritization of people's values. Our robotic snack delivery service can do more than just carry snacks from point A to point B. The service can attempt to understand diverse motivations behind why people order snacks and customize the service in response. For example, for those who use the service to enjoy a social snack break, the service can facilitate coordination of multiple people in support of these values.

4

Detecting Service Orientation with Technology-Based Service⁵

This chapter investigates types of service orientation that people have with technology-based services and ways to detect the orientation. We explore this question in the context of a robotic receptionist agent. We posit that different service orientations evoke different behavioral scripts. We explore two different scripts, receptionist and information kiosk, that we propose channeled visitors' interactions with an interactive robot. We analyzed visitors' typed verbal responses to a receptionist robot (Figure 3) in a university building that has been located in a high traffic area in the Newell Simon Hall building for about eight years (Gockley et al., 2006; Gockley, Forlizzi, & Simmons, 2006). Half of the visitors greeted the robot (e.g., "hello") prior to interacting with it.



Figure 3. A photo of Roboceptionist, a receptionist robot located in a high-traffic entrance area in a building at Carnegie Mellon University.

⁵ This chapter is adapted from a paper published at the CSCW'10 conference (Lee, Kiesler, & Forlizzi, 2010).

4.1 Service Provided by Autonomous Agents

The HCI community has a longstanding interest in online agents, interactive devices, and robots used in collaborative interactions. For example, agents can assist collaborative learning and group coordination (Dillenbourgh, Jermann, Schneider, Traum, & Buii, 1997; Enembreck & Barthes, 2002). Robots can introduce and guide groups of visitors in a variety of settings such as museums, subways, airports, and other public places (Burgard et al., 1999; Fong, Nourbakhsh, & Dautenhahn, 2003; Kuno et al., 2007; Kuzuoka et al., 2008; Schulte, Rosenberg, & Thrun, 1999; Yamazaki et al., 2008).

Computer agents or robots that work in public settings raise some challenging design questions. To be successful in imparting guidance or answering questions, they must elicit cooperation from busy workers or visitors who are total strangers. Furthermore, these interactions are likely to occur in the presence of others. People care about their self- presentation to others in public (Goffman, 1966). If they feel nervous or embarrassed, those feelings may negatively impact their willingness to cooperate.

Researchers have suggested many directions for design to support interactions in public settings with agents or robots. For example, Bickmore et al. (2008) sought to make interacting with agents in public comfortable by bringing the agents to human height and creating natural eye-gaze toward speakers. They involved bystanders by creating a back-screen that displayed the dialogue between the robot and a user interacting with the robot. To attract visitors to a museum guide robot, Thrun et al. (2000) designed it to express happiness through its facial expressions when more visitors approached the robot. Shiomi, Kanda, Ishiguro and Hagita (2006) also found that a robot can increase user engagement in a museum by referring to visitors by their names.

Some researchers have argued that creating user models, for example, by learning from people's repeated interactions over time, can support adaptivity in an agent or robot's interactions with people (Kobsa, 2001). In line with the theory of regulatory fit (Cesario, Grant, & Higgins, 2004), an agent that adapts to people's orientation to computational agents might elicit more cooperation than an agent that does not adapt to this orientation. For example, people may be oriented to treat an agent or robot as a humanlike being or, alternatively, as a computational tool. According to regulatory fit theory, the robot should act to support these different orientations. Nass and Lee (2000) showed that extroverts found an extroverted agent more attractive and credible than introverts did, whereas introverts found an introverted agent more attractive and credible. Goetz, Kiesler, and Powers (2003) showed that matching a robot's personality to users' serious or playful tasks elicited more cooperation from them.

One way to create adaptivity to people's orientation is to use learning algorithms or to detect people's demographic characteristics to build a model of people with different orientations. Doing so may be difficult in public settings, where many encounters will be new and where the population is diverse, mobile, and busy. In this chapter, we argue that we can build reasonable adaptivity in an agent or robot if we can use people's initial verbal cues to estimate their schema for the agent or robot.

4.2 Service Orientation and Interaction Scripts

We posit that people's orientation will elicit different scripts for their subsequent behavior with the receptionist robot. Schemas activate specific behavior through scripts. A script is a "conceptual representation of stereotyped event sequences (Abelson, 1981)." For example, when people enter a restaurant, they follow a standard sequence of events and typical activities such as making small talk with the serving person, placing an order, tipping, and collecting their belongings

before leaving. Likewise, the script for interacting with a human receptionist is cordial whereas the script for interacting with an information kiosk is utilitarian.

Having scripts for daily activities reduces people's cognitive load and allows them to focus on more high-level activities. Scripts also guide people to appropriate behavior in different settings and cultures (Harris, 1994). Thus scripts are very common in everyday behavior. We have scripts for eating in restaurants, for shopping in grocery stores, for visiting museums, for attending sports events, and for holiday dinners with family. People construct scripts for specific, particular contexts through direct and indirect means. Direct script acquisition involves learning through interaction experience with other people, events, or situations. Direct experience tends to initiate a script development process. Indirect script acquisition occurs by means of communication or media. Watching people interact with a robot in a movie or science fiction novels could give people a script for interacting with a robot.

People may consciously choose to perform scripts when facing new situations, although the scripts themselves might be unconscious. Starting a script performance usually entails a commitment to finish it. For example, one does not readily leave a restaurant once seated or walk out of a dentist's office before the dentist is through.

4.2.1 Roboceptionist Scripts

We argue that visitors' scripts for the Roboceptionist will have drawn on their prior interactions with other service personnel and human receptionists or with other computing machines in public settings, tilting them to have either a schematic orientation to a robot as human service person or as computational tool. In a relevant paper, Fischer (2007) reported that the orientations that people held toward mobile robots with varying anthropomorphic forms influenced their instructional strategies and the prosodic strategies that they used to give instruction to the robot. Fischer (2006) proposed that users' choice of dialogue

beginnings might have predicted speakers' concepts of the human-robot situation. One of our goals in this paper is to follow up on this idea and to determine how people's initial dialogue predicts their orientation to a robot.

In our study, we focused on two alternative scripts that people might apply when interacting with the Roboceptionist robot, that is, the script for interacting with a receptionist or other service person or the script for interacting with an informational computational tool such as an information kiosk. We believe these scripts will arise from the schema that people have for the robot. According to social actor theory, people interact with machines as though they are other humans (Reeves & Nass, 1996), but many studies show that this response depends on other factors, such as the form factor of the machine (Parise, Kiesler, Sproull, & Water, 1996), whether people think their interaction is with a computer or person (Morkes, Kernal, & Nass, 1999), the presumed gender of the agent (Powers & Kiesler, 2006), or even its nationality (Lee, Lau, Kiesler, & Chiu, 2005). In this study, we did not manipulate the form factor or any other attributes of the robot. Instead, we assumed that people vary in their schematic orientation and aimed to predict this orientation.

What would be involved in a receptionist script? If this script is evoked when interacting with the Roboceptionist, we believe people would apply the sequence of activities common in everyday interaction with a receptionist or other service personnel. Typical sequences in these scripts might include casual greeting, small talk, instrumental questions, information exchange, and leave-taking (Kendon, 1990). The script also should follow general social norms for weak tie interactions, such as maintaining politeness, not insulting the other personally, and little personal disclosure. The script also should accommodate conversational grounding (Clark & Brennan, 1991).

On the other hand, the same robot might invoke a more computational machine schema. People today have had experiences with interactive computational

machines such as information kiosks or GPS car navigators. In many such systems, the agent or computer looks like a machine or exposes its mechanical parts such as a camera and a laser.

The machine's voice may have a mechanical tone, and people may have to type to the machine instead of speaking to it. These mechanical qualities of the computational machine could reinforce the feeling that this is a machine rather than a social actor.

The Roboceptionist, while somewhat humanlike because of its displayed face and conversational speech, also has many machinelike qualities. The face is a display on a computer screen, the voice has a mechanical quality to it, and users type to talk to it. Due to these mechanical qualities, people might draw analogies between the Roboceptionist and other computing machines. Such machines often act as computational tools that support people's utilitarian goals such as guidance in a museum or in a car. People interact with these devices by directly specifying their goals and instructions by using a graphical user interface, or typing or speaking keywords. In the case of a GPS car navigator, people specify their destination either by typing the destination on its screen or speaking a keyword. The GPS system provides direction in natural human language. When interacting with these types of devices, people typically use an instrumental script: instruct the machine, wait for its reply, and correct it if needed. The script is for communication of intent in a direct manner, and does not use relational conversational strategies.

4.2.2 Greeting as an Indicator of the Script

From previous research and literature on scripts, we hypothesized that whether visitors greeted the Roboceptionist or not would predict which script they performed when they interacted with the Roboceptionist. One of the characteristics of a script is that, once people choose to enter a particular script, they are less likely to stop the script until its end, unless unexpected breakdowns

happen (Abelson, 1981). As the greeting is the first interaction that happens in human social encounters, the greeting could predict whether or not people have followed a script for human social interaction with a receptionist or other service person or, instead, a script for interacting with a machine.

Fischer's study (2007) of people's instructions to a robot showed some evidence for this argument. She reports that people who greeted the robot tended to instruct the robot in full sentences rather than phrases without verbs. Those who greeted the robot also tended to refer to the robot using personal pronouns, "he" or "she," rather than "it," and they used structuring words (e.g., "next," "then"). Encouraged by these findings, we developed hypotheses for people's conversation patterns depending on whether they initially greeted the robot or not.

4.2.3 Hypotheses

From the above arguments, we predicted that people who greet a robot will follow social norms for human-human communication more than those who do not greet a robot. We developed the following specific hypotheses:

H1. People who greet the robot will exhibit more conversational grounding behaviors than people who do not greet the robot.

H2. People who greet the robot will use more relational conversation strategies than people who do not greet the robot.

H3. People who greet the robot will be less likely to use computer command input styles than people who do not greet the robot.

4.3 Method

The method of this study entailed an analysis of utterances that people typed to a receptionist robot over a period of five and a half months. We grouped people

into two groups, those who greeted the robot and those who did not, to show how using a greeting predicted subsequent conversation.

4.3.1 Roboceptionist

As noted above, the Roboceptionist robot, named “Tank,” is situated in a booth in a lobby near the main entrance of the university building. The robot is built with a B21r mobile robot and a 15” flat-panel LCD screen mounted on a pan-tilt unit. It has a caricatured humanlike male face on the screen. It changes its facial expression and rotates its head to look at passers-by. The robot speech is generated from text using the Cepstral text-to-speech engine (Lenzo & Black) and is automatically synched with its lip movements.

To interact with Roboceptionist, people type on a keyboard located in front of the robot. Upon a typed query, the robot gives directions to campus offices and buildings, looks up office numbers of employees, and reports on the weather. The robot also enacts its persona by describing some personal history and preferences if visitors ask. The examples of its personal story include its work experience at the CIA and in Afghanistan, and its family, girlfriend, and dog. The robot’s booth contains various props such as the robot’s photograph with soldiers in the desert to reinforce the robot’s persona.

The robot uses a rule-based, pattern-matching parser, modified from Aine (<http://neodave.civ.pl/aine>) to generate responses to people’s input. During the study, the robot responded to every person’s initial input in the same manner, whether they gave a greeting or not. The robot is passive in that visitors always initiated a conversation, and the robot only responded to their utterances.

People who work in the building can swipe their ID cards or credit cards in a card reader so the robot can call them by name. However, our analysis showed that people rarely swiped their cards. For more details on the Roboceptionist, please refer Gockley et al. (2005).

4.3.2 Data Collection and Coding

Dialogue Log Data

We logged 1180 interactions over 5.5 weeks in March and April 2008. Each interaction was defined as a dialogue that occurred from the moment a person approached the robot until he or she left, as detected by the laser. The unit of analysis is the interaction. When the same utterances were observed multiple times in one interaction, they were calculated as happening once, so that we do not over-count and can measure the percentage of persons who exhibited particular behaviors.

Video Data

To protect people's privacy, the dialogue log data did not contain any contextual information about persons who interacted with the robot. However, we obtained permission to record Roboceptionist-person interactions for one week in March and April 2009 using the security camera installed in the Roboceptionist booth. We coded persons' gender, whether they were alone or with others, and guessed their ages. These codes were compared with the presence of greetings in dialogues with the Roboceptionist so we could evaluate whether gender, age, or being alone predicted greeting the robot.

Measures

We measured attributes of each interaction, and person utterances in each interaction. The unit of coding was an exchange between a person and the robot. A coding scheme for topic was based on coding 197 individual interactions collected over one week in March 2008 by Lee and Makatchev (Lee & Makatchev, 2009; Makatchev, Lee, & Simmons, 2009). A coding scheme for linguistic styles was drawn from the common ground and politeness literature (Brown & Levinson, 1987; Clark & Brennan, 1991).

Person Utterance Measures

We coded whether people greeted the robot or not (such as "Hi," or "What up").

Grounding behaviors had four attributes: relevancy, acknowledgement, repair (rephrase), and misunderstanding. Relevancy was coded if a person built upon the robot's previous utterance. Acknowledgement was coded if a person explicitly expressed his or her understanding of the robot's utterance. Repair was coded if a person rephrased his or her previous utterances. No one misunderstood the robot's utterance, so this factor was not considered in our results.

Relational behaviors were measured by people's politeness, sociable behaviors, and negative behaviors. Politeness was counted when a person said farewell, thanked the robot, made an apology, or said phrases that express courtesy or etiquette (e.g., "please," "Good evening Mr. Tank," "would you mind telling me your name again").

Sociable behaviors were measured by whether people made small talk, called the robot's name during the interactions, made empathetic comments for the robot, introduced themselves or others to the robot, complimented the robot, or told a joke to the robot.

Negative behaviors were measured by whether people said nonsense or insulted the robot, or asked it intrusive questions (e.g., "What is your GPA?," "Are you gay?").

Topics were coded as instrumental, robot-related, and person-related, and others.

Instrumental topics were measured by whether people asked for information about the university where the robot was situated, locations of places (e.g., restaurant or bathroom), information about employees (e.g., office number, phone number, or email), travel information (e.g., how to get a taxi), information about Pittsburgh weather, or the current date and time.

Robot-related topics were measured by whether people asked about the robot's stories and information about the robot (e.g., its name, age, preferences, family, friends, pets).

Person-related topics were measured by whether people talked about their feelings or events in their lives.

We used a code, "other topics", for idiosyncratic comments and questions (e.g., "tell me how babies are born.").

Sentence structure was a coding of sentences, whether they were imperative, interrogative, declarative, or contained no verb.

Interaction

We measured the total duration of each interaction and the total number of utterances a person said.

One coder performed all of the coding, and another coder coded ten percent of the data. They compared their results until they reached agreement.

4.4 Results

We conducted multi-level repeated measures analyses of variance (ANOVAs) to test each of the hypotheses, comparing people who greeted the robot with people who did not greet the robot as a between groups variable, code type as a within groups variable, and each person's interaction as a random control. We report the percentage of the interactions that included behaviors relevant to our hypotheses.

On average, 43 interactions with the robot happened per day. The average interaction duration was 55 seconds and four interactive exchanges (turns) per person. Overall, half of the interactions included a greeting at the start and half did not (Table 5). Those who greeted the robot interacted with the robot longer (Greeting: *Mean* = 78.4 seconds, No Greeting: *Mean* = 31.4 seconds). This

difference in the interaction duration happened because those who used a greeting typed more words (Greeting: *Mean* = 12.6 words, No Greeting: *Mean* = 8.4 words, $p < .0001$), and took more turns (Greeting: *Mean* = 5.5 turns, No Greeting: *Mean* = 3.2 turns, $p < .0001$). We ran additional regression analyses controlling for number of words or turns. These analyses show that the topics, tone (social, polite, and less negative behaviors), and sentence structures still differed across those who greeted versus those who did not.

| <i>Group</i> | <i>Percentage</i> |
|------------------------|-----------------------|
| Greeting | 49.5 % (N=585) |
| Greet & leave | 21.4 % (N=125) |
| Greet & converse | 78.6% (N=460) |
| No Greeting | 50.5 % (N=595) |
| Abusive behavior only | 18.5 % (N=110) |
| No greeting & converse | 81.5% (N=485) |

Table 5. Percentage of interactions that include greeting and those that do not include greeting at the beginning of their interactions with the Roboceptionist.

Only 21.4% of those who greeted the robot left immediately afterwards. Some of these people just wanted to say hello (P: “Good morning to you.” R: “Hello,” P: “Nothing, just wanted to say hi.”). In addition, 18.5 % of those who did not greet exhibited only abusive behaviors such as typing insults or nonsense. Those interactions, lacking conversation, had to be excluded for the subsequent analysis. Very few people swiped their cards, and the number did not differ across the two groups ($G = 2\%$, $NG = 3\%$).

4.4.1 Grounding Behavior

We predicted (H1) that those who greeted the robot would use more grounding strategies such as acknowledging, making relevant follow-up comments, and rephrasing (Table 6). As predicted, people who greeted the robot exhibited grounding behaviors that they would do with other people more frequently than

those who did not greet the robot. The ANOVA shows a significant greeting vs. no greeting main effect ($F(1, 943) = 14.5, p < .001$), a grounding type (acknowledge, repair, relevancy) main effect ($F(2, 1886) = 18.3, p < .001$) and a marginal interaction effect ($F(2,1886) = 2.5, p = .08$).

| <i>Grounding Behavior</i> | <i>Greeting</i> | <i>No Greeting</i> | <i>p</i> |
|---------------------------|-----------------|--------------------|----------|
| Relevancy | 0.35 | 0.23 | < .001 |
| Acknowledgement | 0.23 | 0.15 | < .01 |
| Repair | 0.27 | 0.23 | n.s. |

Table 6. Comparison of percentage of grounding behavior between people who greeted the robot and those who did not. (The p value is calculated using ANOVA planned contrasts.)

More people who greeted the robot than those who did not acknowledged that they understood the robot’s utterances by typing “yeah,” “ok.” The planned contrasts show significant differences between those who greeted the robot and those who did not in acknowledging the robot’s response ($F = 7, p < .01$).

People who greeted the robot tended to continue their dialogues, building on what the robot said to them as compared with those who did not greet the robot ($F = 17, p < .001$). They also asked follow-up questions based on the stories that robot told them or made statements about the stories.

In the following dialogue, the robot did not understand most of what the person said, yet the person commented on the robot’s response, rephrased his or her utterances, and continued the dialogue:

P: ok tank.

R: That’s me!

P: how are you today?

R: Dr. Spitola suggested that a self defense seminar might help to reassemble my self esteem.

He assures me that just because I’m armless doesn’t mean I’m any less of a person. It doesn’t

mean I have to take this abuse. I don't have to settle for sexual harassment from that disgusting buffer. The good doctor has arranged for me to attend a self defense camp designed especially for abused machines. I'm looking forward to getting away for a while.

P: ouch. are you depressed?

R: I didn't understand that. It might help if you rephrase it.

P: are you depressed?

Example 1. In this excerpt, the person acknowledged the robot's response and rephrased his utterance.

4.4.2 Relational Behavior

We predicted (H2) that those who greeted the robot would use relational conversational strategies such as small talk, and empathetic and polite remarks more frequently than those who did not.

Sociable Behavior

The repeated measures ANOVA of social interaction used seven codes shown in Table 7. As predicted, the analysis shows a significant main effect of greeting ($F(1, 943) = 85, p < .0001$), significant differences among the codes ($F(6, 5658) = 57, p < .0001$) and a significant interaction ($F(6, 5658) = 33.9, p < .0001$). The interaction reflects the fact that those who greeted the robot were much more likely to engage in small talk and also call the robot by name. People who greeted the robot initiated small talk with the robot more frequently than those who did not greet the robot. They asked "How are you?" and "How's your day?" ($F = 273, p < .0001$). Those who greeted the robot also called the robot by name during their conversation more frequently than those who did not ($F = 9.3, p < .01$).

As predicted, the analysis shows a significant main effect of greeting ($F(1, 943) = 85, p < .0001$), significant differences among the codes ($F(6, 5658) = 57, p < .0001$) and a significant interaction ($F(6, 5658) = 33.9, p < .0001$). The interaction reflects the fact that those who greeted the robot were much more likely to engage

in small talk and also call the robot by name. People who greeted the robot initiated small talk with the robot more frequently than those who did not greet the robot. They asked “How are you?” and “How’s your day?” ($F = 273, p < .0001$). Those who greeted the robot also called the robot by name during their conversation more frequently than those who did not ($F = 9.3, p < .01$).

| <i>Sociable Behavior</i> | <i>Greeting</i> | <i>No Greeting</i> | <i>p</i> |
|--------------------------|-----------------|--------------------|----------|
| Small talk | 0.22 | 0.03 | < .0001 |
| Empathetic comment | 0.02 | 0.00 | n.s. |
| Joke | 0.01 | 0.00 | n.s. |
| Call robot’s name | 0.08 | 0.05 | < .01 |
| Introduce self | 0.03 | 0.01 | n.s. |
| Introduce others | 0.004 | 0.00 | n.s. |
| Compliment | 0.03 | 0.01 | n.s. |

Table 7. Comparison of percentage of sociable behavior between people who greeted the robot and those who did not. (The p value is calculated using ANOVA planned contrasts.)

Although the difference was not statistically significant, those who greeted the robot also tended to make empathetic comments on what the robot said:

P: Do you have a girlfriend?

R: I had another date. Heinzilla, the score board at Heinz Field. She was a bit large. Okay, she’s massive, but who am I to judge? I thought, if she can look past my lack of arms, I can look past her gargantuan proportions. She was a great conversationalist, and it was wonderful to talk with a woman who knew so much about sports. I went home and was very excited, but unfortunately I can never see her again.

P: Why not

R: I rented a video of great moments at Heinz Field. You know, just to see her in action, but there was a problem. She got way too much action. I couldn’t believe how many people have scored with her. Two at the same time! No, Heinzilla is not the girl for me.

P: I’m sorry to hear that.

Example 2. The excerpt shows a person’s empathic comment to the robot.

Polite Behavior

The repeated measures ANOVA of polite interaction used four codes shown in Table 8.

| <i>Polite Behavior</i> | <i>Greeting</i> | <i>No Greeting</i> | <i>p</i> |
|------------------------|-----------------|--------------------|----------|
| Thank you | 0.16 | 0.07 | < .0001 |
| Farewell | 0.21 | 0.11 | < .0001 |
| Apology | 0.02 | 0.01 | n.s. |
| Courtesy | 0.06 | 0.03 | n.s. |

Table 8. Comparison of percentage of polite behavior between people who greeted the robot and those who did not. (The p value is calculated using ANOVA planned contrasts.)

The analysis shows a significant main effect of greeting ($F(1, 943) = 28, p < .0001$), significant differences among the codes ($F(3, 2829) = 63.9, p < .0001$) and a significant interaction ($F(3, 2829) = 7.4, p < .0001$).

The planned contrasts show those who greeted the robot, as compared with those who did not, were more likely to thank the robot ($F = 24.3, p < .0001$), and to say farewell to the robot ($F = 32.2, p < .0001$). When closing a conversation, those who greeted the robot tended to inform the robot that they were leaving rather than abruptly stop the conversation (Example 3).

P: hello

R: Hello

P: may you tell how i can find a women's room? R: [Robot gives directions]

P: you are so good

R: I didn't understand that. It might help if you rephrase it.

P: you look great

R: Thank you.

P: we are leaving

R: I didn't understand that. It might help if you rephrase it.

P: have a good night

R: I didn't understand that. It might help if you rephrase it.

P: bye

R: Good bye

Example 3. The excerpt shows leave taking and grounding repair in doing so.

Negative Behavior

The repeated measures ANOVA of social interaction, controlling for interaction number, used three codes shown in Table 9.

| <i>Negative Behavior</i> | <i>Greeting</i> | <i>No Greeting</i> | <i>p</i> |
|--------------------------|-----------------|--------------------|----------|
| Ask intrusive questions | 0.05 | 0.10 | < .001 |
| Nonsense | 0.06 | 0.05 | n.s. |
| Insult | 0.03 | 0.04 | n.s. |

Table 9. Comparison of percentage of negative behavior between people who greeted the robot and those who did not. (The p value is calculated using ANOVA planned contrasts.)

The analysis shows a marginal main effect of greeting ($F(1, 943) = 3.2, p = 0.07$), significant differences among the codes ($F(2, 1886) = 6.3, p < .01$) and a significant interaction ($F(2, 1886) = 4.6, p < .01$). Those who greeted the robot exhibited negative interaction less frequently than those who did not greet the robot. The planned contrasts show significant differences between those who greeted the robot and those who did not in asking intrusive questions to the robot ($F = 12, p < .001$). Nonsense words (e.g., “djfkjdfkj”) and insults were uncommon and did not differ across the two groups of people.

4.4.3 Conversation Topics

We used repeated measures ANOVA to test the effects of greeting and number of utterances on different topics (instrumental, robot-related, person-related, and

others) (Table 10). The analysis shows a main effect of greeting ($F(1, 943) = 7.3, p < .01$), a main effect of topic ($F(3, 2829) = 293, p < .0001$), and an interaction of greeting x topic ($F(3, 2829) = 4.2, p < .01$). The interaction reflects the fact that those who greeted the robot were more likely to talk about the robot and themselves (or other persons).

| <i>Topic</i> | <i>Greeting</i> | <i>No Greeting</i> | <i>p</i> |
|----------------------------------|-----------------|--------------------|------------------|
| Instrumental topics | 0.50 | 0.52 | n.s. |
| Location of place, event, person | 0.37 | 0.38 | n.s. |
| Weather | 0.12 | 0.06 | < .001 |
| Date | 0 | 0 | n.s. |
| Time | 0.02 | 0.08 | < .01 |
| Robot-related topics | 0.48 | 0.38 | < .001 |
| Family/friends/pets | 0.15 | 0.18 | n.s. |
| Robot demographic | 0.31 | 0.20 | < .0001 |
| Preference/opinion | 0.12 | 0.08 | n.s. |
| Person-related topic | 0.08 | 0.03 | < .02 |
| Person emotion | 0.04 | 0 | < .0001 |
| Person self information | 0.04 | 0.03 | n.s. |
| Other topic | 0.6 | 0.9 | n.s. |

Table 10. Instrumental, robot-related, and person-related topics that people talked about with the Roboceptionist. (The p value is calculated using ANOVA planned contrasts.)

Those who greeted the robot showed more interest in the robot’s demographic information and talked about themselves more frequently than those who did not greet the robot ($F = 4.9, p < .02$). They spontaneously talked about their mood (e.g., “I’m lonely,” “I’m bored”) or their characteristics or events in their lives (e.g., “We won the basketball [game]”). In contrast, they did not mention instrumental and knowledge-related topics more than those who did not greet the robot.

4.4.4 Sentence Structure

We predicted (H3) those who greeted the robot would be less likely to use computer input command styles of language. The repeated measures ANOVA of sentence structure used four codes shown in Table 11. In the direction predicted, the analysis shows a significant main effect of greeting ($F(1, 943) = 27, p < .0001$), significant differences among the codes ($F(3, 2829) = 434, p < .0001$) and a significant interaction ($F(3, 2829) = 11, p < .0001$).

People who greeted the robot tended to use full sentences, as compared with those who did not greet the robot. As Fischer's study showed, those who did not greet the robot used more keywords. The planned contrasts show significant differences between those who greeted the robot and those who did not in (i) using keywords (computer command styles) ($F = 5.5, p < .02$), (ii) using imperative sentences ($F = 5.8, p < .02$), (iii) using declarative sentences ($F = 17.8, p < .0001$), and (iv) interrogative sentences ($F = 27.7, p < .0001$).

| <i>Sentence structure</i> | <i>Greeting</i> | <i>No Greeting</i> | <i>p</i> |
|---------------------------|-----------------|--------------------|----------|
| No verb | 0.23 | 0.30 | < .02 |
| Imperative | 0.17 | 0.10 | < .02 |
| Declarative | 0.35 | 0.24 | < .0001 |
| Interrogative | 0.85 | 0.71 | < .0001 |

Table 11. Comparison of percentage of interactions that use different sentence structures (mood) between people who greeted the robot and those who did not. (The p value is calculated using ANOVA planned contrasts.)

4.5 Discussion and Limitations

The results showed that people who greeted the Roboceptionist treated the robot more like a person than those who did not greet the robot. People who greeted the robot exhibited more grounding behaviors and relational conversation strategies than those who did not greet the robot. They acknowledged the robot's

response, and continued the conversation by building on the robot's responses. They also initiated small talk, and a few of them mentioned events in their lives or how they were feeling. These findings support our hypothesis that those who greet a robot will follow a receptionist script rather than an information kiosk script.

For privacy reasons, we could not determine the identity of those who interacted with the robot. We also did not want to use any intrusive measures that might have altered people's behavior. Thus we must speculate on the characteristics of people who greeted the robot. According to one anthropomorphism theory (Epley, Waytz, & Cacioppo, 2007), people who treat a computer in a humanlike way might do so because they feel lonely and are reaching out for social interaction or companionship. Alternatively, people who greet a robot might be those who are generally polite, extraverted, or social, perhaps regardless of whom they are meeting. Our video data did not show any relationships between greeting behavior and gender, or between greeting behavior and age or the number of people with the person who was interacting with the robot. Thus, ascertaining the attributes of people who greet a robot (or the circumstances that encourage schemas that elicit greetings), must await future research.

Our results suggest that when a robot in a public university setting has both humanlike and machinelike form factors, about half of those who interact with it will engage with the robot as though it were a person, and half, as though it were a machine. We observed this division in only one setting with only one robot. The robot's head was an animated male character on a screen and it had a mechanical tone of voice. People conversed with the robot by typing to it rather than speaking. Thus the robot was a unique combination of anthropomorphic and machine attributes. For this reason, we cannot claim generalizability of our observation that half of all interactions involved a greeting. The finding might not

hold with a robot that understands speech, or with robots having different form factors.

A limitation of our analysis is that there was no way to distinguish whether people who interacted with the robot were visitors, staff, or students. Even though the robot had a user identification system, few people swiped their cards. We do not know how many people changed their orientation to the robot over multiple visits. The robot was autonomous, and communication breakdowns occurred frequently. Some people obviously adjusted their expectations during the conversation when the robot did not understand their utterances. Finally, because this study was done in a natural setting, there might have been selection bias. For example, people who are interested in robots or new technology might have approached the robot more than others.

Still, we have learned something important from this study about the predictability of people's behavior in public settings. Although we recognize the huge variability and diversity of people's orientations and goals, we also see in our results a measure of predictability. People seem to have signaled their intentions and orientation to the robot in their approach behavior, through a greeting or a lack of greeting. This result fits very well with other work in which researchers are attempting to glean information about people's goals and concerns from easily obtainable cues and behavior (e.g., Forgy et al., 2005; Leshed, Hancock, Cosley, McLeod, & Gay, 2007).

4.6 Design Implications

Detecting whether or not people greet a computational agent provides an opportunity to design adaptive dialogue systems for cooperation. Social agents might use relational strategies with those who greet the agent and more utilitarian dialogue with those who do not greet the agent. People who spontaneously greet agents might be likely to respond more positively to agents that attempt small talk

than agents that do not. Bickmore and Cassell (2001) showed that agents that made small talk reduced the perceived distance between themselves and users, and increased users' trust, especially when users were extroverted.

To imagine how such an idea might be used in designing for cooperation, imagine a robot that invites collaboration among children, not just answers questions or gives instructions. The robot can detect greetings and whether multiple children are present. When more than one child is present, and the children seem to be in a sociable mood, the robot's dialogue is programmed to encourage collaboration. Otherwise, the robot acts more instrumentally. Two scenarios follow.

Scenario after greeting: It is an ordinary day, and a group of children approaches the robot, saying "Hi!" Amy wants to know where Tunisia is located because a friend just visited there. The robot might pose questions and remarks to encourage the children to engage with each other. For instance, the robot says, "Tunisia is in North Africa. Which of you can help Amy find Tunisia on my map?"

Scenario after no greeting: It is before the examination period and the children are preparing for a test. Amy approaches the robot and asks, "Where is Tunisia?" The robot, using an instrumental orientation answers, "Tunisia is in North Africa. See it on my map."

The scenarios above are only one example of how a simple greeting, and perhaps other easily obtainable information about the context and the people involved, might evoke a branching strategy that would honor people's own schemas and scripts for an agent in a particular social situation. A greeting might evoke shorter but more interactive utterances, more questions of the user, or more emotionality than the absence of a greeting. For context-aware systems, people's preferences and behavior patterns could be recorded and stored for future conversations.

4.7 Summary

The mental structures that people apply towards other people have been shown to influence the way people cooperate with others. These mental structures or schemas evoke behavioral scripts. In this study, we explored two different scripts, receptionist and information kiosk, that we propose channeled visitors' interactions with an interactive robot. we analyzed visitors' typed verbal responses to a receptionist robot in a university building. Half of the visitors greeted the robot (e.g., "hello") prior to interacting with it. Greeting the robot significantly predicted a more social script: more relational conversational strategies such as sociable interaction and politeness, attention to the robot's narrated stories, self-disclosure, and less negative/rude behaviors. This finding suggests that if we can sense simple things about behaviors when interacting with technology services, we can better adapt services and interfaces for them.

5

Matching Interaction Style to Service Orientation⁶

The study described in CHAPTER 4 suggests that people's first interaction behaviors with a service robot can indicate whether they have a relational vs. utilitarian mental model of technology service. This chapter explores the efficacy of matching style of interaction to this mental model is effective in the context of breakdown recovery in a robotic service.

5.1 Research Context: Service Breakdowns and Recovery

Robots are increasingly able to perform services for people. Robotic services will be especially attractive for doing repetitive, unpleasant, or effortful tasks in workplaces, hospitals, and public environments. Robotic services may offer an overall service improvement, such as when a robot reliably delivers medications in a hospital. However, as anyone who has dealt with airlines, hospitals, and stores knows, services are imperfect. Robots that deliver services also will make mistakes. For example, the hospital delivery robot may interrupt nurses dealing with an emergency (see Mutlu & Forlizzi, 2008; Siino & Hinds, 2005).

Service mistakes can lower people's trust and satisfaction, and increase their reluctance to use the service again. Service mistakes are a leading cause of customer switching behavior (Keaveney, 1995). We argue that designing

⁶ This chapter is adapted from a paper published at the HRI'10 conference (Lee, Kiesler, Forlizzi, Srinivasa & Rybski, 2010).

appropriate robotic service recovery strategies is a necessary component of robotic services. People often become emotionally upset when there is a service breakdown, and often are more dissatisfied by a failure of the recovery than the mistake itself (Bitner, Booms, & Tetreault, 1990). Gracefully mitigating breakdowns can be important for sustaining people's satisfaction and preventing them from abandoning a robotic service. Appropriate recovery strategies also offer an opportunity for a strengthened relationship between the service and its users (Aaker, Fournier, & Brasel, 2004; Hart, Heskett, & Sasser Jr., 1990; Spreng, Harrell, & Mackoy, 1995).

Service breakdowns can occur at many levels of a service. For example, a service breakdown at the organizational level occurs when management fails to put resources into customer service, and a service breakdown at the individual provider level occurs when a customer service agent is rude. When a service is partly automated, customers can blame the breakdown on factors at any level. Technology used in service provision can complicate the blame and recovery process. For example, when an automated telephone reservation agent's understanding of speech is faulty, people may not be sure who or what is at fault, including themselves.

In this chapter, we focus on an interactive robot that delivers a personal service incorrectly, using the example of a mobile robot that delivers the wrong drink (Figure 4). We apply ideas from psychology, consumer research, and HRI to the question of how such a robot should mitigate the error and aid service recovery. From a scenario study of the delivery mistake, we show that service failure has negative effects on satisfaction and perceptions of the robot, that a recovery strategy can mitigate these negative effects, and that successful strategies depend in part on peoples' orientation toward services.

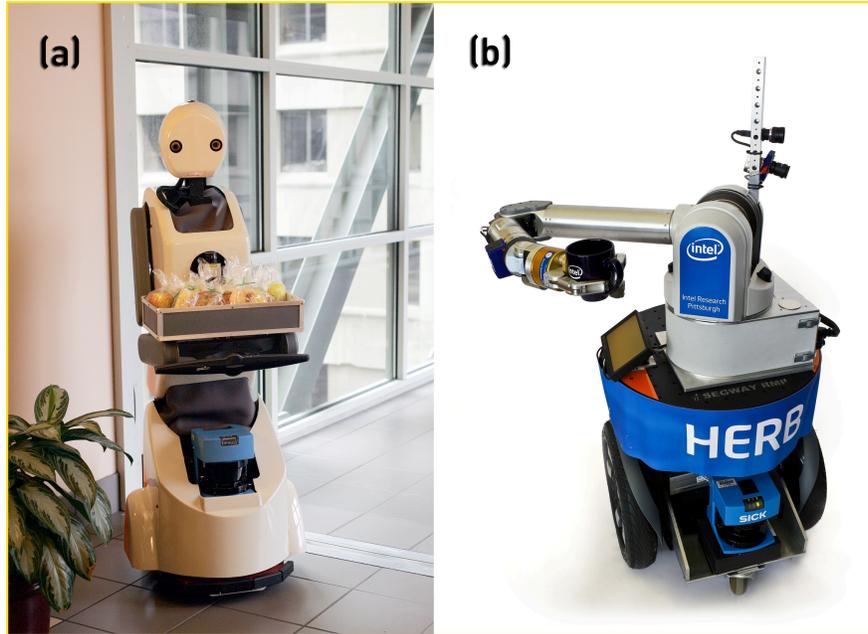


Figure 4. Snackbot (a) and HERB (b), service robots used in the study.

5.2 Mitigation Strategies

Robots that provide a personal service create interdependence between the robot and the user. Prior research suggests that the nature of this interdependence and the robot's design can affect people's responses to system errors (Hinds, Roberts, & Jones, 2004). People may feel a loss of control when they do not understand why the robot fails (Nomura, Kanda, Suzuki, & Kato, 2008). In one study, the participants blamed their robot partner more when the robot was humanlike rather than machinelike (Hinds et al., 2004). In another study, the more autonomous a robot was, the more people blamed it for failure, and explaining the reason for the failure did not help much (Kim & Hinds, 2006). This work suggests that people may have high expectations of robotic services that complicate their experience where there is a service breakdown.

Hypothesis 1. A robot's service breakdown will have a negative influence on service satisfaction.

5.2.1 Expectancy-Setting Strategies

Service satisfaction research shows that the degree to which a service meets people's expectations is a primary determinant of their satisfaction with the service (Oliver, 1980; Smith, Bolton, & Wagner, 2009). People may have elevated expectations of a service robot for at least two reasons. First, most people do not have much experience with robots, and thus robots present an ambiguous situation (Sheridan, 2002). In such situations, people may be prone to using mental shortcuts or heuristics to make attributions. For instance, if the robot is capable in some ways, such as navigation and speech production, people may assume the robot is also capable in other ways, such as speech recognition and social skills (Kunda, 1999). Second, people may generalize from themselves (Ross, Greene, & House, 1977). That is, people may assume that tasks that are easy for them, such as opening doors, recognizing people, and distinguishing between similar objects, are also easy for robots.

A person's elevated expectations of a robot and a mismatch between their expectations of service and the robot's capabilities could exacerbate the influence of a service breakdown. One strategy to address this problem would be to forewarn people of the difficulty of a task for a robot, to re-set their expectations and bring them more in alignment with the actual probability of breakdown. People who are informed that the robot is likely to make mistakes or that a task is difficult for the robot might be more willing to accept breakdown without feeling anger or frustration.

Hypothesis 2. Forewarning people that a task is difficult for the robot will mitigate the negative influence of breakdown on service satisfaction.

5.2.2 Recovery Strategies

Apologies are one of the most commonly used recovery strategies in service organizations. A wealth of research shows that a service provider's apology

conveys politeness, courtesy, concern, effort, and empathy to customers who have experienced a service failure, and enhances their evaluations of the encounter (Hart et al., 1990; Kelley, Hoffman, & Davis, 1993). Because research has shown people treat computers as social actors (Reeves & Nass, 1996), and that flattery from a robot was positively perceived by people (Johnson, Gardner, & Wiles, 2004), we predict that a robot service provider's apology for service failure will be effective as well.

Hypothesis 3. A robot's apology will mitigate the negative influence of the robot's service breakdown on service satisfaction.

Providing compensation, such as an exchange, a refund, or a discount coupon is another commonly used recovery strategy in service organizations. Tax, Brown, and Chandrashekar (1998) claim that compensation is the recovery strategy most associated with customers' perception of fairness in service. By compensating customers' time, resources, or money lost due to the breakdown, this strategy attempts to restore the inequalities in the transaction. We believe that this strategy will be equally effective in a robotic service.

Hypothesis 4. A robot's offering compensation will mitigate the negative influence of the robot's service breakdown on service satisfaction.

Providing customers with alternative actions to achieve their goals is another strategy that can be effective in mitigating service breakdowns. As noted above, service breakdowns can cause people to feel emotionally upset and a loss of control. Giving them options can help reassert the sense of control. This idea has been tested mostly in health services and services for the elderly. In those domains, it has been shown repeatedly that giving people options increases their perceived control and positive outcomes (Heckhausen & Schultz, 1995).

Hypothesis 5. A robot's offer of options will mitigate the negative influence of the robot's service breakdown on service satisfaction.

5.2.3 Service Orientation

As described in CHAPTER 3, research in marketing and consumer psychology suggests that people's responses to service recovery strategies may depend on their schema or model of service (Ringberg, Odekerken-Schroder, & Christensen, 2007). People's response to service recovery strategies may depend on their orientation to service. In accord with the theory of regulatory fit (Cesario, Grant, & Higgins, 2004), a robot with a service recovery strategy that adapts to people's orientation to service might elicit more satisfaction than a robot that does not adapt to this orientation. Those who have a more relational orientation to services might treat a robot as a social service provider, and expect it to apologize after a mistake. Those who have a more utilitarian orientation to services may prefer the robot to offer compensation.

Hypothesis 6. A robot's choice of recovery strategy that is matched with people's orientation to services will mitigate the negative influence of breakdown on service satisfaction.

5.3 Study Design

To test these hypotheses, we conducted an online between-subjects scenario survey. All participants saw a video of one of two service robots (Figure 4), and then viewed a scenario in which the robot either gave correct service or made an error. We investigated people's reactions to the robot's error and to different mitigation strategies. The study was a 2 (forewarning vs. no forewarning) x 4 (apology, compensation, options, and no recovery strategy) x 2 (humanlike vs. non-humanlike robot) design with two additional control groups in which the robots did not make an error.

5.3.1 Participants

We recruited participants on Amazon mTurk (<https://www.mturk.com>), the local Craigslist (<http://pittsburgh.craigslist.org>), and a university study participant recruiting site (Center for Behavioral Decision Research). The recruiting message said that the objective of the survey was to pretest the design of delivery service robots. We offered \$1.00 plus a chance at a \$50 Amazon raffle prize. Four hundred fifty-seven persons responded. Of this number, we omitted who completed the survey multiple times, did not conform to the participant requirements (e.g., being at least 18 years old), did not take at least 6 minutes to complete the survey, and who gave incorrect answers to questions used to identify participants who randomly selected answers (Kittur, Chi & Suh, 2008). These procedures left 317 participants in the sample, over two-thirds of the original number. Due to random assignment, there were different numbers of participants in each condition, at least 14, with most having 16-19 participants. Fifty-five percent of the sample was female. Their ages ranged from 18 to 67, with a median of 33 years. They were fairly well educated, with the mean level being at the college level. Most of the participants knew very little about robotics. The mean response on the 4- point scale was 1.7 ($SD = .8$; 1 = “no knowledge other than books movie”, 2 = “a little knowledge of robotics”). Their mean programming experience was 2 on the 4-point scale ($SD = 1$; 1 = “no experience”, 2 = “little experience”).

5.3.2 Materials

The Snackbot robot, as shown in Figure 4 (a), is a 4'5" tall delivery robot that offers snacks to people. The robot carries a tray and travels on wheels at about 1-2 mph, can rotate completely in place, and can navigate the building autonomously. The robot's head is mounted on a 2-axis pan/tilt unit allowing it to pan 270 degrees and to tilt 80 degrees, so it can rotate towards people or turn away, nod, and look up or down. The robot can emit speech or sounds. It has a LED mouth

and a directional microphone that feeds into the Sphinx4 speech recognition system. Further details are described in Chapter 6.

The HERB robot (Figure 4 (b)) is an autonomous robot that consists of a RMP 200 Segway base that carries a Barrett WAM arm and hand for grasping objects (Srinavasa et al., 2009). Sensing is provided by a SICK laser rangefinder and two cameras. HERB has been developed to efficiently navigate, search, and map indoor environments. Visual object recognition allows it to identify and localize a set of household objects. These objects can be grasped, lifted, and carried by the arm and hand. The robot is designed to perform dexterous operations with these objects, such as pouring water from a pitcher.

Half of the participants evaluated the Snackbot robot and half evaluated the HERB robot as target service providers in the study. We assumed the Snackbot robot would be seen as more humanlike, due to its anthropometric body and head. To help the participants understand how the robot could provide service, we presented a 30-second video that showed the robot carrying an object in an office environment. The robots did not interact with any people in the video. We explained that the robot is autonomous, that it makes its decisions on its own. We did not use the robot's name and referred to the robot as the "robot in the video." The logo on the HERB robot was removed when the video was recorded.

5.3.3 Scenarios

After the participants saw the video, we asked them to evaluate a situation in which the robot delivered a service. To present the situation, we used a scenario method that has been used in Human-Computer Interaction and HRI studies (e.g., Woods, Walter, Koay, & Dautenhahn, 2006). We constructed 16 different scenarios to represent one of the eight experimental conditions (the presence of forewarning strategy x the presence of recovery strategies), and showed these with both types of robot (more humanlike vs. less humanlike). We also had a

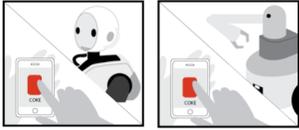
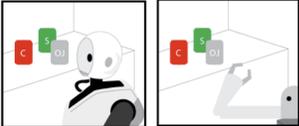
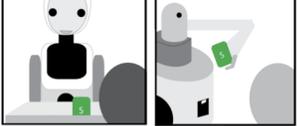
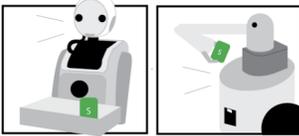
| Scene | Script | Condition Manipulation |
|---|---|---|
|  | <p>Chris is thirsty, and asks the robot to bring a can of Coke. The robot says, "OK."</p> | <p>Forewarning: Chris is thirsty, and asks the robot to bring a can of Coke. The robot says, "OK, but it might be hard to identify Coke from other sodas."</p> |
|  | <p>The robot looks at the Coke and Sprite on the counter.</p> | <p>Forewarning: The robot looks at the Coke and Sprite on the counter. The robot is confused because there are multiple cans.</p> |
|  | <p>After a few minutes, the robot comes back with a can of Sprite. Chris says, "OK, good. But I wanted a Coke."</p> | <p>Control: After a few minutes, the robot comes back with a can of Coke. Chris says, "OK, good."</p> |
|  | <p>The robot says,</p> | <p>Apology: "I thought this was Coke. I apologize for bringing the wrong one."</p> <p>Compensation: "I thought this was Coke. I will give you this drink for free."</p> <p>Options: "I thought this was Coke. I can go back and try to find it. You can also show me a picture of a Coke, so I can recognize what it looks like."</p> <p>No recovery & Control: "OK."</p> |

Figure 5. Scenarios and conditions used in the study.

control scenario for each robot where no breakdown occurred, resulting in 18 scenarios in total.

Each scenario described a situation in which a person, "Chris," had a knee injury recently. In the scenario, Chris orders a can of soda from a delivery robot, but (except in the control conditions) the robot makes a mistake and returns with the wrong soda. Depending on the condition, the robot attempts to recover its mistake using different strategies. Independent of the employed recovery strategies, the outcome of the service was same. Figure 5 shows the scenarios.

We chose a breakdown caused by an error in the robot's perception as a quite realistic error that might be applied to diverse robots regardless of their actuators. We used the projective viewpoint when creating scenarios, as they have shown to minimize social desirability effects and have considerable external validity

(Nisbett, Caputo, Legant, & Marecek, 1973). The name Chris was chosen to be gender-neutral, so that both male and female respondents could identify with the character. We also used a written description of the scenario, and attempted to convey only an unemotional reasonable reaction by Chris. The scenarios were succinct, so that respondents could easily read and understand them.

5.3.4 Procedure

The scenarios were embedded in a Survey Monkey template. Once they clicked the link to the survey, participants were connected to a php page, which randomly directed them to one of the 18 surveys. This process was invisible to participants. The survey began with a 30 second video clip that introduced one of the robots to the participant. After the video, we asked some pre-scenario questions to gather participants' impressions of the robot, and to assess their orientation to services.

Next, we displayed one of the scenarios in the 18 conditions. After the scenario, participants provided their evaluations of the robot and the service, and provided some information about themselves.

5.3.5 Measures

The survey included items to measure the participants' evaluation of the robot before and after the scenario, and their evaluation of the service, their orientation to services, and manipulation checks.

Evaluation of the Robot

We adapted questions used to measure people's evaluations of a service provider (Surprenant & Solomon, 1987). These items consisted of 10 bipolar adjectives in a 5-Likert scale (capable, efficient, organized, responsible, professional, helpful, sincere, considerate, polite, friendly) where higher scores were more positive. We asked these questions before and after the scenario was presented, to measure the impact of the scenario on the evaluation of the robot.

To examine whether the robot evaluation adjectives were measuring the same or different underlying factors, we conducted a factor analysis of the data from these items. Factor analysis of the pre-scenario ratings suggested we could create two scales from the items, one being a measure of “politeness” (*Cronbach’s* $\alpha = .80$) and the other, a measure of “competence” (*Cronbach’s* $\alpha = .81$). Two items, “responsible” and “professional,” loaded equally on both factors so were included in both scales. We also asked questions to measure how much the participants liked and felt close to the robot, and how humanlike they thought the robot was. All items used 5-point Likert scales where a “5” was the most positive rating.

Evaluation of the Service

Three questions in the post-scenario survey measured the participants’ evaluation of the service from Chris’ point of view using Likert-type scales. We asked participants to rate whether the robot gave good or poor service (1 = “very poor” and 5 = “very good”) and to rate how satisfied Chris would be with the service (1 = “completely dissatisfied” and 5 = “completely satisfied”).

We also measured how likely participants thought that Chris would use the service again using a 5-point Likert scale (1 = “would avoid using the service” and 5 = “would want very much to use the service”).

Service Schema Orientation

The pre-scenario survey included 9 items about people’s orientation toward food services in general. There were three questions to infer relational orientation (e.g., “I like to have a positive relationship with a server [waitress and waiter] in a restaurant.”), three questions to infer utilitarian orientation (e.g., “Efficient food service is important to me.”), and three questions to infer the level of control they desired over the service process and outcomes (e.g., “I like to have control over the process and outcome of food service.”).

Factor analysis of the 9 items suggested two factors were captured by the items. These were used to construct two scales, one scale with three items to measure

relational orientation (*Cronbach's* $\alpha = .77$), and the other scale with 6 items to measure utilitarian/control orientation (*Cronbach's* $\alpha = .65$).

Manipulation Checks

To assess whether participants detected a service error, we asked participants whether the robot made an error (where 1 = “none” and 5 = “many errors”). To assess whether participants detected a forewarning, we asked them how difficult the task was for the robot (1 = “very difficult” and 5 = “very easy”). To assess whether participants detected a service recovery, we asked participants whether the robot made any error corrections, and if so, how many.

5.4 Results

We conducted one-way analyses of variance on the effects of the relevant conditions on the manipulation check ratings. The participants in the breakdown conditions thought that the robot made mistakes (Control = 1.08 ($SE = .11$) vs. No Strategy = 2.19 ($SE = .08$), Apology = 2.26 ($SE = .08$), Compensation = 2.27 ($SE = .08$), Options = 2.17 ($SE = .08$), $p < .001$). There were no differences across the breakdown conditions. The participants who saw scenarios with recovery strategies said that the robot made more error corrections than those who saw no recovery strategy (No strategy = 1.40 ($SE = .10$) vs. Apology = 1.75 ($SE = .10$), Compensation = 1.97 ($SE = .10$), Options = 2.02 ($SE = .10$), $p < .02$, with no differences across recovery strategy conditions). The manipulation check for the forewarning condition showed that the manipulation was effective. Those in the forewarning condition thought that the task was more difficult for the robot than those in the no forewarning condition (Forewarning = 2.4 ($SE = .09$), No forewarning = 2.80 ($SE = .09$), $p < .002$).

5.4.1 Evaluation of the Robots

The participants' pre-scenario evaluations of the robots differed. As expected, they rated the Snackbot robot much more humanlike than the HERB robot

(Snackbot = 2.7 ($SE = .1$) vs. HERB = 1.80 ($SE = .1$), $F(1, 315) = 51, p < .0001$). The participants also liked the Snackbot more than the HERB robot ($F(1, 315) = 5.8, p = .01$) and felt somewhat closer to the Snackbot ($F(1, 315) = 3.4, p = .06$). We used a repeated measures ANOVA to compare the pre-scenario and post-scenario ratings of the robot. Having seen the scenario made people feel that both robots were more humanlike ($F(1, 315) = 81, p < .001$) and also closed the gap between the Snackbot and HERB robots (interaction $F(1, 315) = 15, p < .001$). These findings suggest that the scenario, which entailed HRI, was somewhat humanizing as compared with the video, which did not entail HRI and only showed the robot carrying objects.

5.4.2 Impact of Service Breakdown

Hypothesis 1 predicted that a robot's service breakdown would have a negative influence on service satisfaction. Regardless of the robot the participants saw, a service breakdown without mitigation had a strong negative impact on the rating of the service and the robot.

We compared the control conditions (where no service breakdown occurred) against the conditions where a breakdown occurred and no mitigation took place (i.e., the robot brought the wrong soda). We also crossed this comparison with the Snackbot and HERB robots, to see if service breakdown would be viewed more severely if the robot were more humanlike. These analyses test service evaluations using analyses of variance with breakdown vs. control crossed with the robot (Snackbot vs. HERB), and their interaction effects. The evaluations of robot ratings are multi-level models that take into account participants' pre-scenario ratings.

The impact of the breakdown did not differ depending on which robot participants saw. On the contrary, as predicted, regardless of the robot participants saw, a service breakdown without mitigation had strongly significant negative impact on the ratings of the service and the robot. Table 12 shows the

participants' evaluation of the service and the robot when the service was performed smoothly, as compared with the situation when there was a service breakdown.

| <i>Dependent measure</i> | <i>No Breakdown</i> | <i>Breakdown</i> |
|---------------------------|---------------------|-------------------------|
| Service Evaluation | | |
| Good or bad service | 4.64 [.13] | 2.32*** [.13] |
| Service satisfaction | 4.64 [.11] | 2.16*** [.15] |
| Willingness to return | 4.61 [.18] | 2.58*** [.13] |
| Robot Evaluation | | |
| Politeness | 3.81 [.10] | 3.25*** [.07] |
| Competence | 4.01 [.11] | 3.01*** [.08] |
| Trust robot | 3.81 [.18] | 2.86*** [.13] |
| Like robot | 3.34 [.17] | 3.41 [.12] |
| Feel close to robot | 3.17 [.19] | 2.80 [†] [.14] |

Table 12. The impact of breakdown on service and robot evaluations.

Note. The numbers show the least squared means and the standard error in brackets. Robot evaluation ratings shown are post-scenario, and the analyses control for pre-scenario ratings.

[†]p < .10, *p < .05, ** p < .01, *** p < .001

5.4.3 Impact of Expectancy-Setting (Forewarning) Strategy

Hypothesis 2 predicts that forewarning people that a task is difficult for the robot will mitigate the negative influence of breakdown on service and the robot evaluations. The forewarning strategy had positive impact on the overall evaluation, in particular the evaluation of the robot.

We compared each robot's forewarning versus no forewarning in the conditions where there had been a breakdown. Table 13 shows that generally, the robot's lowering expectations did mitigate the negative influence of breakdown on the robot evaluation. The forewarning strategy somewhat mitigate how good or bad the participant judged the service was, yet did not increase the participants'

perception that Chris was satisfied with the service and Chris' willingness to use the service again.

| <i>Dependent measure</i> | <i>No Forewarning</i> | <i>Forewarning</i> |
|---------------------------|-----------------------|-------------------------|
| Service Evaluation | | |
| Good or bad service | 2.54 [.09] | 2.62* [.09] |
| Service satisfaction | 2.33 [.07] | 2.52 [†] [.07] |
| Willingness to return | 2.94 [.09] | 2.97 [.09] |
| Robot Evaluation | | |
| Politeness | 3.53 [.05] | 3.73** [.05] |
| Competence | 3.03 [.06] | 3.27** [.06] |
| Trust robot | 2.73 [.09] | 3.01* [.09] |
| Like robot | 3.28 [.08] | 3.61** [.08] |
| Feel close to robot | 2.76 [.09] | 3.03* [.09] |

Table 13. The impact of the forewarning strategy on service and robot evaluations.

Note. The numbers show the least squared means and the standard error in brackets. Robot evaluation ratings shown are post-scenario, and the analyses control for pre-scenario ratings.
[†]p < .10, *p < .05, ** p < .01, *** p < .001

5.4.4 Impact of Recovery Strategies

Hypotheses 3 - 5 predict that any recovery strategy (apology, compensation, and options) would be better than no strategy. Overall, all three strategies were effective in mitigating the negative influence of the robot's service breakdown, but worked differently on different dimensions of the service and robot evaluation (Table 14).

We tested the effects of the different recovery strategies on the participants' evaluation of the service and the robot including variables as: robot, forewarning, and recovery strategy, and all their interactions. Because the evaluation of the robots was performed twice, before and after the scenario, we conducted a so-called multi-level regression analysis that tested participants' post-scenario

ratings, controlling for their pre-scenario ratings. In each case we conducted planned contrasts between each strategy and the No strategy condition.

| <i>Dependent measure</i> | <i>No Strategy</i> | <i>Apology</i> | <i>Compensation</i> | <i>Options</i> |
|---------------------------|--------------------|-------------------------|-------------------------|----------------|
| Service Evaluation | | | | |
| Good or bad service | 2.35 [.13] | 2.70 ^t [.12] | 2.72* [.13] | 2.56 [.13] |
| Service satisfaction | 2.16 [.11] | 2.46 ^t [.11] | 2.68*** [.10] | 2.36 [.11] |
| Willingness to return | 2.66 [.14] | 3.06* [.14] | 2.99 ^t [.13] | 3.12** [.13] |
| Robot Evaluation | | | | |
| Politeness | 3.24 [.07] | 3.97*** [.08] | 3.62*** [.07] | 3.69*** [.07] |
| Competence | 2.99 [.08] | 3.27* [.08] | 3.16 [.08] | 3.20 [.08] |
| Trust robot | 2.84 [.12] | 3.01 [.13] | 2.85 [.12] | 2.79 [.12] |
| Like robot | 3.40 [.11] | 3.72* [.11] | 3.31 [.10] | 3.36 [.11] |
| Feel close to robot | 2.79 [.12] | 3.16* [.13] | 2.81 [.12] | 2.85 [.12] |

Table 14. The impact of the recovery strategies on service and robot evaluations.

Note. The numbers show the least squared means and the standard error in brackets. Robot evaluation ratings shown are post-scenario, and the analyses control for pre-scenario ratings. Significance tests compare each strategy with the No strategy comparison condition.

^tp < .10, *p < .05, ** p < .01, *** p < .001

5.4.5 Service Orientation and Recovery

Hypothesis 6 predicts that those with a more relational orientation to services would respond better to the apology strategy, whereas those with a more utilitarian service orientation would respond better to the compensation strategy.

The orientation scales were distributed normally and were correlated at just $r = .28$, suggesting the two scales access somewhat different service schemas. Only 15% of the participants scored low on both scales, whereas 42% scored high on both scales (indicating high involvement with service). The rest were split between high scores on a relational orientation versus high scores on a more utilitarian orientation.

The analyses of Hypothesis 6 tested the effects of the scores on the two orientation scales, recovery strategy, and their interactions on ratings of service. (Interactions unrelated to the hypothesis were not significant, so we do not discuss them further.) We also included forewarning and robot as control variables. These analyses show that having a stronger relational orientation biased participants to appreciate the apology strategy significantly on two of the three measures of service (Figure 6). The good vs. bad service interaction was significant ($F(3, 267) = 2.67, p < .05$). These relational orientation participants, in fact, tended not to like the compensation strategy almost as much as no strategy.

The same effect of relational orientation also was true of the ratings of service satisfaction (interaction $F(3, 267) = 2.7, p < .05$). Moreover, the utilitarian orientation also interacted with recovery strategies on this measure. In this case, those who scored higher in utilitarian orientation rated the service as most satisfactory when they saw the compensation strategy (interaction $F(3, 267) = 3.6, p = .01$). These participants tended not to like the options strategy, possibly because it entailed more effort for the user.

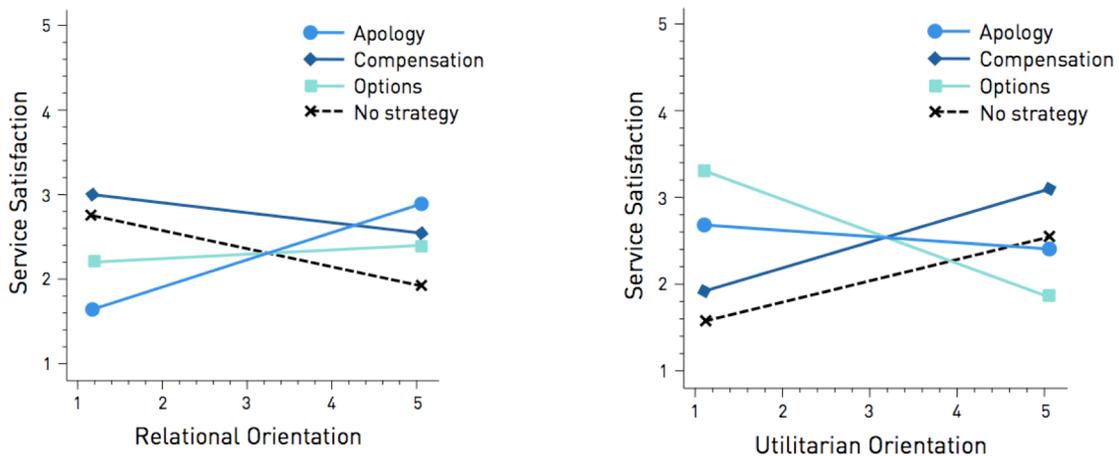


Figure 6. The relationship between participants' service orientation and their ratings of service satisfaction in the different recovery strategy conditions.

5.5 Discussion

The study showed that, overall, both the expectancy-setting strategy and the recovery strategies we tested were effective in mitigating the negative impact of the error that the robot made on participants' impression of the robotic service.

The service evaluation analyses showed that overall, having a recovery strategy was better than not having one. For ratings of good or bad service, for example, the planned contrasts showed that those in the strategy conditions, together, rated the service as better ($F(1, 265) = 4.4, p < .05$). The apology strategy and the compensation strategy were each better than no strategy, but the options strategy was not. Even stronger differences differentiated recovery strategies from no strategy when the participants rated service satisfaction and whether Chris would be willing to use the service again. Generally the apology strategy was effective across many ratings. The compensation strategy was particularly effective in increasing the participants' perception that Chris was satisfied with the service, and the option strategy was effective in increasing the participants' perception that Chris would be willing to use the expectancy-setting strategy was particularly effective in extenuating the negative influence on the robot evaluation and somewhat effective in improve the participants' judgment of the quality of the service. All the recovery strategies increased the ratings of the robot's politeness. However, only the apology strategy was effective in making the robot seem more competent, making the participants feel closer to and liking the robot more. The compensation strategy was most effective in increasing the participants' perception that Chris was satisfied with the service, but less effective in increasing their perception of Chris' willingness to use the service again. The apology and option strategies were effective in increasing the participants' perception that Chris would use the service again.

The results also showed that tailoring the recovery strategy to people's cultural orientations is important. As seen in Figure 6, those with a relational orientation

responded particularly well to an apology, whereas those with a more utilitarian orientation responded better to compensation.

Our results suggest that having a plan for mitigating robot service errors will be an important component of HRI designs for robots that deliver services or otherwise help people. However, our study has some important limitations that prevent us from generalizing overly from our findings. First, and most important, we used a hypothetical scenario survey technique. Even though the response to the scenarios was consistent with previous literature on real services, we do not know for sure if people's responses to robotic services in real environments will be the same. Second, we only tested the efficacy of the strategies for one type of task and one error. Replicating this study with different tasks, situations, robots, and errors would make the findings much more robust. Finally, we did not test how the recovery strategies, such as apology with compensation, would work in combination with each other. There is some evidence that combining apologies with compensation could backfire (Ringberg et al., 2007), especially with relationally oriented people who might see the compensation as manipulative. Our data also suggest that utilitarian oriented people may not like compensation mixed with options, perhaps because options entails more effort.

5.6 Implications

The findings from the study have interesting implications for the design of robotic services. As noted above, our results suggest that a robot should be designed so that it can mitigate errors in its behavior or the service through expectation setting and social error recovery strategies. Our results also suggested that matching these strategies to the person's orientation would be useful.

How would a robot know a person's service orientation? We can suggest one technique, based on our work on people's initial interactions with a robot described in Chapter 4. In our previous study, we analyzed visitors' verbal

responses to a receptionist robot in a university building. We observed that half of the visitors greeted the robot (e.g., “hello”) prior to interacting with it. Greeting the robot significantly predicted a more social script with the robot: more relational conversational strategies such as sociable interaction and politeness, attention to the robot’s narrated stories, self- disclosure, and less negative/rude behaviors. This finding suggest people’s first words with a robot can predict their schematic orientation to a robotic service, thus making it possible to design robots that adapt their recovery strategy at the outset of an interaction. For example, a robot might use relational recovery strategies (such as apologies or empathic comments) with those who greet the robot, and more utilitarian dialogue and compensation for errors with those who do not greet the robot.

There are also various ways to design for appropriate expectations. One possible design direction would build on the work on robot helpers (Torrey, 2009), which suggests that if a robot gives advice or helps someone, and exhibits some speech disfluencies, then it is perceived as less controlling without detracting from its perceived expertise. These findings suggest other ways to gracefully mitigate errors by humanizing the robot and making it seem competent but far from perfect.

5.7 Summary

Filmmakers and science fiction writers are envisioning robots, like those in the movie “Surrogates,” that perform tasks almost perfectly, and that can repair themselves when they break, but robots in reality are a long way off from that vision. Furthermore, as long as people design robotic services for people, there will be errors, whether perceived or real, in these services. This study represents an initial attempt to demonstrate the importance of error mitigation in robotic services. The results suggest mitigation strategies are most effective when tailored to people’s service orientations. The results suggest a rich area of debate and research on how a robot can fail gracefully.

6

Snackbot: Design of Robotic Platform and Service for Long-Term Interaction⁷

This chapter introduces Snackbot, a semi-humanoid delivery robot and its service for long-term interaction in the real world. We describe a human-centered design process of a robot and its service personalized for long-term interaction.

6.1. Design Approaches to Robotic Systems

Experimental systems, including receptionists, assistants, guides, tutors, and social companions, have been developed as platforms for research and technology development (Blow, Dautenhaun, Appleby, Nehaniv, & Lee, 2006; Breazeal, 2003; Bruce, Nourbakhsh, & Simmons, 2002; Burgard et al., 1999; Dautenhaun et al., 2007; DiSalvo, Gemperle, Forlizzi, & Kiesler, 2002; Fong, Nourbakhsh, & Dautenhahn, 2003; Gockely et al., 2006; Gockely, Simmons, & Forlizzi, 2007; Kanda, Takayuki, Eaton, & Ishiguro, 2004; Mutlu, Osman, Forlizzi, Hodgings, & Kiesler, 2005; Pollack et al., 2002; Sidner, Lee, Morency, & Forlines, 2007; Simmons et al., 2003; Wada, Shibata, Saito, & Tanie, 2002; Yamaoka, Kanda, Ishiguro, & Hagita, 2006). Much of this work has taken place in the research laboratory, but a few systems have made the successful transition to real world settings such as museums and educational institutions (Hayashi et al., 2007; Kanda et al., 2004; Nabe et al., 2007; Severinson-Eklundh, Green, & Huttenrauch, 2003; Shiomi et al., 2006). Real world settings raise the bar to fluid, natural interaction with robotic systems.

⁷ Some part of this chapter is adapted from the paper published at the HRI'09 conference (Leet et al., 2009).

Robots in real settings also need to interact with people appropriately. Safe interactions are necessary to contribute to ethical research in the field, to improve people's trust in and comfort with robotic technology, and to ensure safety and reliability for all who come into contact with this technology. Socially appropriate interaction behaviors are needed so people like the robot and are interested in interacting with it over time.

Our research group seeks to develop robots that travel around and near people, and that support them in real-world environments. We are interested in developing robots that act as social assistants, with the ability to use speech and gesture, and engage people in a social manner. A major goal is to create mobile robots and services that interact with people over a period of time. To address these challenges, we designed and developed the Snackbot (Figure 7), a robust robot that roams semi-autonomously in campus buildings, offering snacks to office residents and passersby.



Figure 7. Snackbot delivering snacks in a hallway.

We designed the Snackbot not just as a snack delivery service, but also as a research platform to investigate questions related to long-term interaction with social robots. Achieving long-term interaction with social robots poses several technical, interaction, and design challenges. The robot must be robust and powerful enough to operate autonomously and interact with multiple users for extended periods of time. The technology should also be flexible enough to accommodate technical improvements and new applications. To test different approaches to HRI over time, researchers should be able to manipulate aspects of the robot's physical appearance and behavior. We are particularly interested in how a robot delivers a service after the initial novelty effect has worn off.

In this chapter, we present our design process for the Snackbot, shaped by our initial design goals, constraints we discovered along the way, and design decisions guided by interim empirical studies. We document this process as a contribution for others in HRI who may be developing social robots that offer services.

6.2 Context of Use

Robotic advances are being directed towards special populations, including elders, those with physical and cognitive disabilities, and others. We want to design robots that can interact with almost everyone, regardless of any dispositions to using technology. To satisfy this goal, we are interested in how a robot can deliver a service within a work environment.

We chose to design a robot that would provide snack deliveries in the two connected buildings in which we work. By “snack” we mean light food eaten between meals. Snacks include “junk food” such as food offered in vending machines, and “healthy” snacks such as fruit and nuts. Snacking is practiced by a majority of people in the developed world (Bellisle et al., 2003; Ovaskainen et al., 2006; Verplanken, 2006). In workplaces, people snack in their offices and labs as well as in halls, cafeterias, and food vending areas.

A robot delivering snacks must have a wide range of mobility. The buildings are large, ranging between 4 and 8 floors. About 1000 people work or visit these buildings each day. Because the buildings offer only prepackaged snacks in convenient locations, we felt a snack service that offered higher quality snacks would be a useful application for a long-term product and service in these buildings. Most snacks that do exist are highly caloric, and the robot could include healthier snacks in its offerings. We felt many technical and design research questions could be discovered in understanding how a robotic snack service might succeed within the social and environmental context of our buildings, how it would differ from traditional vendors and vending machines, and how it could support people's goals such as taking a break from work and delivering snacks as gifts to people. We have described some of the research supporting these decisions in a separate paper (Lee, Kiesler, & Forlizzi, 2008). This research, combined with our overall research goals in HRI, led to the three design goals that anchored our design process.

6.3 Design Goals

We had three design goals for development of the Snackbot robot:

The first was to develop the robot holistically. Rather than advancing technology per se or focusing on one aspect of design or interaction, such as a dialogue system, we took a design approach that considered the robot at a human-robot-context systems level (Nelson & Stolterman, 2003). Such an approach allowed us to think about the emergent qualities of the product and service, which might not be recognized if the system were analyzed in component parts rather than holistically.

The second goal was to simultaneously develop a robotic product and service. By this we mean that the robot as a product would have to be more than sociable and attractive; it would need to deliver something useful to people. We adopted this

goal to increase the likelihood that people would continue to be interested in interacting with the robot over a period of time (Morelli, 2002). By developing a snack delivery service that worked with wireless service points in the building, we could collect and record knowledge about people's snack preferences, and use these to further enhance the service we provide to them.

The third goal was to develop interaction designs that would help to evoke social behavior. Because the robot was meant to serve as a research platform that would be used by people over time, decisions about functions and features were made supporting the interest of promoting sociability. For instance, we aimed to have the robot interact with people using natural language. Other research has shown that people interact with a robot longer when it exhibits social cues (Bruce et al., 2002; Forlizzi, 2007). Other aspects of sociability that we plan to explore and extend include personalization of the service, and robot politeness and non-verbal behaviors (Bickmore & Picard, 2005).

6.4 Snackbot Team

Developing a robot in a holistic manner required interdisciplinary collaboration. The Snackbot team consisted of 5 faculty, 5 graduate students, and 7 undergraduate students drawn from several disciplines including design, behavioral sciences, computer science, and robotics. Because of the wide range of expertise, we frequently had members from one group attending the meetings of the other. For instance, the designers worked on the form studies but they often interacted with the engineers, and everyone helped out with the empirical studies. Organization of this group was assisted through the use of an on-line forum called the Kiva (www.thekiva.org), hosted on a website accessible to team members from anywhere on the Internet. This web facility was useful because all of the information was organized and presented in a searchable, threaded format to the entire team. A great deal of emphasis was placed on good documentation

of process, code and interim prototype, so any new person on the project could follow in the footsteps of those that worked on it before.

6.5 System Overview

To give the reader a snapshot of what the design process has achieved and where we plan to go, we present an overview of the Snackbot.

6.5.1 Hardware

The Snackbot robot is based on the existing CMAssist platform (Rybski, Yoon, Stolarz, & Veloso, 2007), augmented with some commercial hardware and software and new elements and code. The Snackbot uses a MobileRobots Inc. Pioneer 3 DX base for mobility. Bumpers, sonars, and a SICK laser are used to detect and avoid collisions and to detect position within an environment. A Hokuyo URG laser is mounted in the robot's chest to detect potential collisions with higher objects, and to detect people by torso.

The Snackbot currently has non-functional arms that hold a tray, used for carrying snacks (Figure 8). The tray is equipped with 12 load cells; each is capable of measuring a weight range of 13 to 763 grams. With this functionality, the robot will know when someone has removed or replaced a snack on the tray.

An Acoustic Magic microphone array is mounted under the tray. It serves as the primary audio input source for the robot's natural spoken language and dialog processing system. The robot's head is mounted on a Directed Perception pan/tilt unit, affording a 360-degree pan range and a 111-degree tilt range. A Point Grey Bumblebee 2 stereo camera is mounted behind the robot's eyes; a monocular Point Grey Dragonfly2 camera is mounted on the top of the head and is fitted with a 180-degree fisheye lens from Omnitech Robotics. The Snackbot also has two 2.4GHz Intel laptops running Ubuntu Linux for data processing.

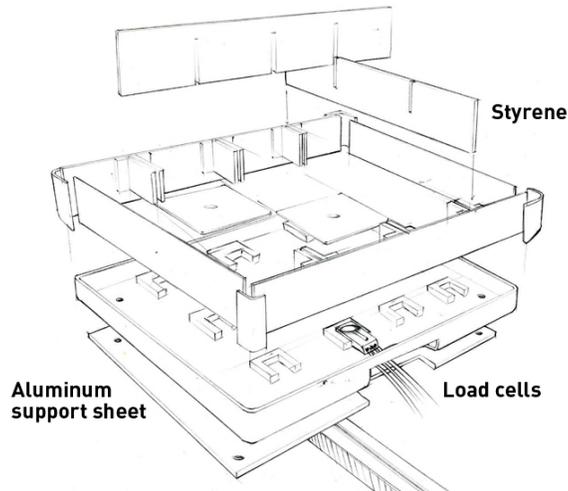


Figure 8. Sketch of the Snackbot tray, showing the load cell configuration.

6.5.2 Software

The Snackbot uses MobileRobot's ARIA API that works with the Pioneer base. ARNL provides functionality for map construction, and path planning. A distributed software architecture developed by the CMAssist project (Rybski et al., 2007) interfaces with the behavior control modules and the speech processing interface. When the Snackbot moves through its environment, it will track its current position by comparing the current set of laser scans and an odometry estimate against a previously programmed map.

We use an Augmented Transition Network (ATN) manager for our dialogue system. This will allow for a flexible discourse structure, but will require more work by a dialogue designer. We also use an open source Sphinx4 speech recognizer system (<http://cmusphinx.sourceforge.net/sphinx4/>), written in Java, and the Cepstral speech synthesizer (Lenzo & Black).

6.5.3 Form

The form of the robot is made of cast fiberglass and is custom designed to fit the Pioneer base and an internal structure that anchors the laptop and other components. It has a semi-humanoid form and uses simple geometric shapes.

There are three exterior pieces: one for the head, one for the torso, and one for the base.

6.5.4 Interaction

There are several basic modes of interaction with the robot. In stationery mode, the robot is positioned in a social space and people can approach the robot to help themselves to a snack. In roaming mode, the robot uses the map to visit people's offices and to deliver snacks. Snacks will be ordered in advance (using a web page, email, or IM) or selected during the visit. Customers will register on the Snackbot website and will get points for snacks in exchange for being involved in the research.

To interact with the Snackbot, people eventually will engage in natural dialogue with the dialogue system. Visual feedback will occur through an LED mouth, which will indicate when the robot is "talking." Sound will be used as an additional informational cue.

6.6 Design Process

To holistically conceive of the robot as a product and service, we had to consider many aspects of the design process concurrently: the social and physical context of the environment it would operate in, its form, and how it would interact with people (Figure 9). Table 15 summarizes our design activities aligning with our overarching design goals for each phase of the project.

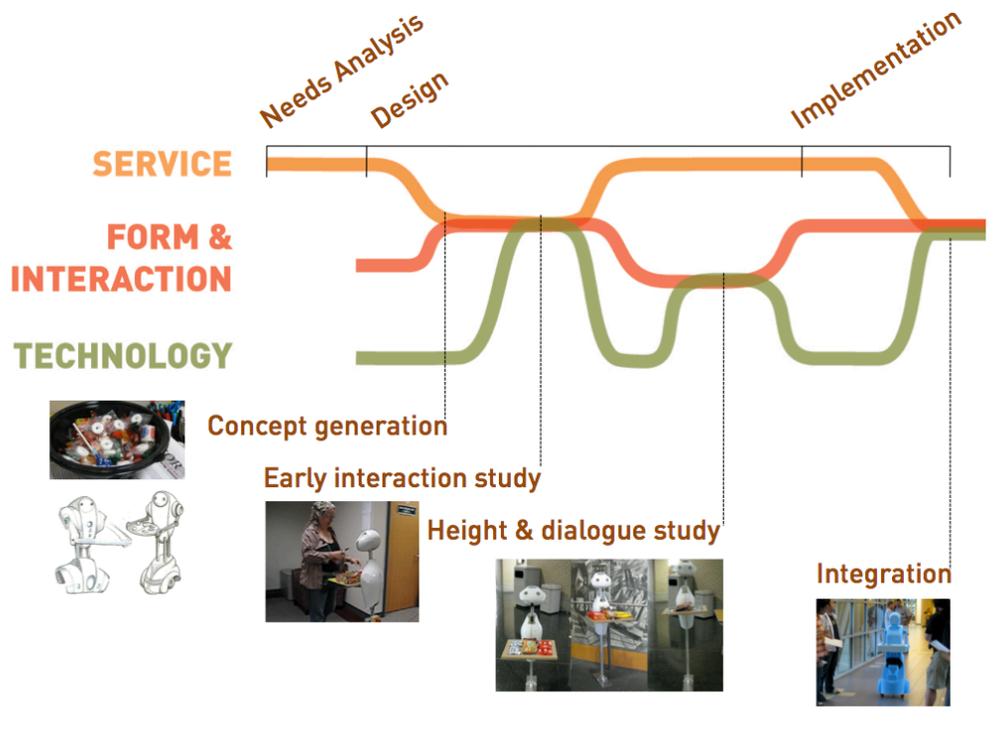


Figure 9. Overview of the design process.

| <i>Design activities / Design goals</i> | <i>Needs analysis & service design</i> | <i>Form giving & interaction design</i> | <i>Documentation & evaluation</i> |
|--|---|--|---|
| Develop the robot & its service holistically | Context research on snacking and service concept generation | Form research; assess tradeoffs in material and technology selection | Evaluative field studies to understand change in people, product use, physical and social context |
| Develop product & service simultaneously | Site survey of snacking | Scenario development; trial of delivery service with human confederate | Process blueprint for robotic product and service design |
| Develop interaction designs that evoke social behavior | Understanding of physical, social, psychological reasons people snack | Dialogue structure study; height and approach study | Checklist for interaction design considerations in HRI |

Table 15. Design goals and design activities.

6.6.1 Needs Analysis and Service Concept Generation

We conducted needs analysis and context research on snacking in our office buildings, described in more detail elsewhere (Lee et al., 2008). Our environmental research took the form of a campus survey to document all of the places where people can get snacks. From candy dishes in administrative offices to vending machines in the basements of building, we mapped site lines and studied each site for accessibility. We also mapped distances to, and popularity of, nearby locations that are popular for snack breaks — for instance, a local coffee shop that is frequented by members of the campus community. One of the findings from this work was that people mainly choose convenience over snack quality, but they do not mind walking for a snack if social interaction is part of the activity (and especially if the snack is free). Based on our observations, we created two basic modes of service for the robot: mobile and stationary. We decided that the robot in mobile mode should offer to deliver healthy snack choices such as fruit, and that in stationary mode should offer high quality snacks in communal locations that would attract groups (refer to the service blueprint in Chapter 3). These decisions support our overall design goal of evoking social behavior, and ensure that we are not making a robot that will only bring fattening snacks to sedentary people.

6.6.2 Observation of Hot Dog Vendor’s Interactions with Customers

To design a model of a one-time interaction structure with rapport-building strategies over repeated interactions, we studied how a hot dog vendor, a long-time community member at Carnegie Mellon University, interacted with his customers. Our findings directly informed the design of the Snackbot service, including its dialogue structure and content, and how its interactions develop over time.

Method

I observed interactions between the owner of Joe's Hotdog, a hotdog stand on the Carnegie Mellon campus, and his customers (Figure 10). Joe's interactions with his customers act as a good service interaction model for the Snackbot for several reasons. Joe has been serving hot dogs for about a decade on the Carnegie Mellon campus and has many repeat customers. During the observation period, some customers visited him just to chat with him even when they were not buying hot dogs. Many of Joe's customers are Carnegie Mellon community members, so they have the same organizational identities. The hot dog stand is the same type of food service as the snack delivery service, and in both cases, the interactions between vendor and customer are brief and focused on transaction as customers leave after the purchase and consume their food elsewhere.



Figure 10. Joe, his wife, and a helper at Joe's Hot Dogs.

Observations were conducted for two hours during lunch times in November, 2007 over two days. Our goal was to learn how the hot dog vendor, Joe, was using his gestures, conversation, and physical surroundings to interact with his customers. I got permission from the vendor, and sat on a wall next to the hot

dog stand, where the interactions between the hot dog vendor and his customers could be seen and their conversation could be overheard. I took notes of their conversation and non-verbal behaviors with time stamps. After the two observation sessions, I did a short, informal interview with Joe to ask a few clarification questions about observed behaviors.

After the observation sessions, I went through the notes and extracted emerging themes. Particular focus was paid to the interaction strategies the vendors used, whether the vendor interacted with first time and repeat customers differently, and, if so, how.

Findings

In total, three people provided services at Joe's Hotdog: Joe, Joe's wife, and a helper. Joe was the main person who took orders from customers and conversed with them. Joe's wife grilled and prepared hot dogs with the helper and occasionally joined Joe's conversations with customers. Joe and his wife's interactions with their customers were short, efficient, but also surprisingly social – they remembered their conversations with repeat customers and followed up on them; they also promoted social interaction among customers.

Structure of Interaction between the Vendors and Customers

The hot dog stand was located on a sidewalk at a T-shaped junction between two academic buildings (the Schools of Business and Architecture). The hot dog stand and its signage were arranged to optimize the transaction processes (Figure 11). During the lunchtime, there was a line of customers waiting to buy hot dogs. As it took a couple of minutes to prepare the hot dogs, there was also a space where customers could wait for their hot dogs.

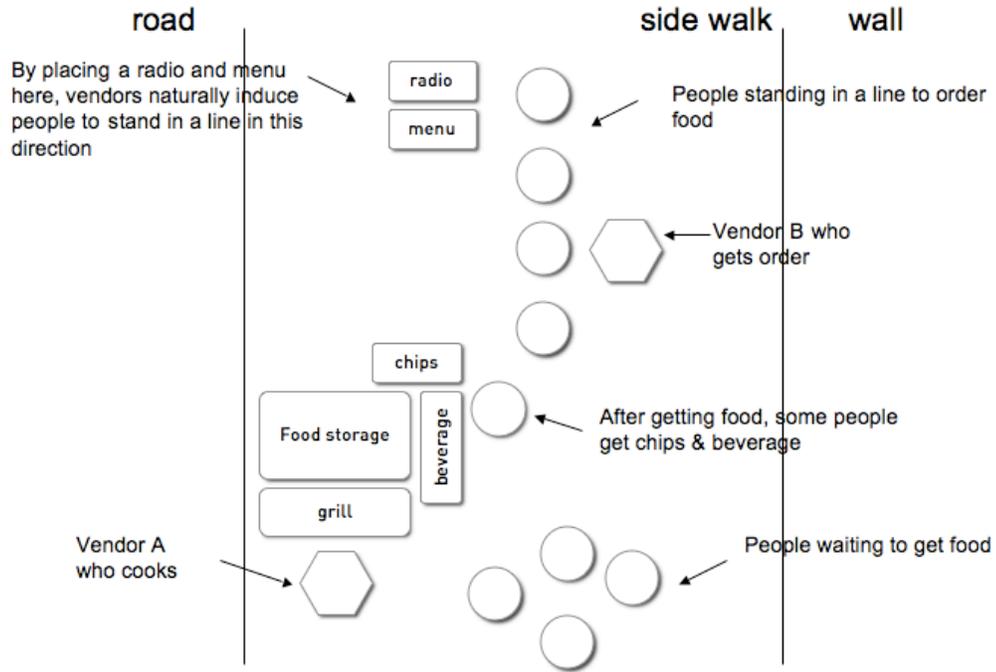


Figure 11. A physical configuration of Joe's hot dog stand.

The interactions between service providers and customers consisted of a few simple stages (Figure 12). The vendors greeted customers, received orders, asked about preferences for toppings and sauces, received money, and handed out hot dogs to finish the transaction. Social interactions usually occurred when the hot dog vendor greeted the customers or while customers were waiting to order or receive their food, but they could happen at any point in this transaction process.

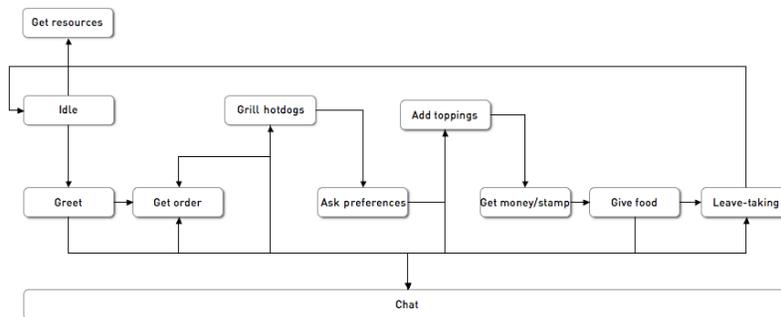


Figure 12. Interaction diagram for the interaction between the hot dog vendors and their customers.

Social Interaction

The feeling of personalization seemed to arise from a seemingly small collection of the vendors' behaviors. The vendors paid individual attention to each customer. They rarely forgot to say "hi" or "bye", acknowledging each customer's arrival and departure even when they were talking with other customers. When the vendor could not attend to a customer who had just arrived, he said "I'll be with you in a minute."

Joe's non-verbal behaviors and gestures also conveyed the vendors' personal attention and friendliness toward customers. He always smiled when handing food to customers or asking about sauces and toppings. He patted male customers' shoulders if they were repeat customers. Joe also approached people in line and stood beside them to get their orders, instead of shouting across the distance from the hot dog stand. He made eye contact with customers while greeting them, getting orders, receiving money, and handing out food. Later, during the interview, the vendor told the observer that he usually tries to mirror the posture of the customer to whom he is talking in order to make him/her feel comfortable.

Small talk between the vendor and the customers or among customers was also an important part of the interaction. In many observed interactions, the customers talked more than the vendor even though it was the vendor who initiated the conversation. It seems that many customers enjoyed conversing with the hot dog vendor in general. The vendor initiated small talk by asking what was going on at school (e.g., "How're your classes?", "How's your orientation?") or talking about the weather (e.g., "I didn't expect the weather to be this nice", "Tomorrow is supposed to be beautiful"). Stories related to the food that they sold also came up as small-talk topics (e.g., a joke or story about Canada triggered by a Canadian soda, or the comparison of their products with a MacDonald's Big Mac).

The vendor individually engaged with each customer, but sometimes tried to engage a group of people by talking about a topic that people can easily relate to. They talked about common seasonal or organizational events such as exams, or observed commonalities among customers (e.g., “It seems like everybody is dressed up. What’s going on today?”). In one instance, the vendor asked customers waiting in line, “Everybody is ready for Thanksgiving?” Some customers responded to this question and then started chatting with each other. (There seemed to be more frequent interaction among customers waiting in line at Joe’s Hotdog, as compared to other food-related businesses such as Starbucks or the La Prima coffee stand in an academic building at Carnegie Mellon University.)

While many customers were social, it is worth noting that the vendors did not initiate conversation with all customers. When customers came as a group and conversed among themselves, or when they remained quiet, the vendors did not attempt to engage in small talks.

Personalized Interaction

The vendors tended to interact with different types of customers in different ways. When customers seemed to be new-comers, the interaction unfolded with more focus on explaining different options and information about menu. For example, when the vendor asked whether the customer is getting a snack, a drink, or the combo, the customer said “I don’t know what the combo is.” The vendor told the customer what the combo was and explained the different hot dog sauces.

When customers were repeat customers, Joe and Joe’s wife remembered their names and preferences. They greeted the repeat customers by saying their names, and asked if they wanted their usual orders before the customers had to ask (e.g., “Do you want the four-pound Kosher dog?”). For some customers, the vendors even did not ask what they were going to get. For example, when Joe’s wife saw one customer was approaching from a distance, she started taking the Kosher dog

out of the food storage area. When the customer arrived at the stand, Joe and his wife just greeted the customer and started grilling the Kosher dog for the customer.

The vendors also remembered past conversations with repeat customers and followed up on them in their small talks. For example, the vendors remembered that one customer was planning a trip to Thailand and asked him about the trip. In another instance, a man (who seemed to have an office in an adjacent building, based on his conversation with Joe) came out just to chat with the vendor, without getting any food. He told Joe that he felt like he'd caught a cold, and Joe gave him some candy for the cold.

These observations about Joe's interaction structure, small talk topics and social and personalization strategies directly informed the design of the Snackbot service interaction, which is described in the next section.

6.6.3 Form Giving and Interaction Design

Form giving and interaction design encompass all of the activities necessary to generate a first design of the robot, both in terms of design and varied studies to confirm the design. In this phase, we researched and generated robot forms, and also conducted empirical studies to evaluate the design decisions that we made.

Product research took the form of collecting and analyzing images of existing social robots, which ranged from animals to abstract to humanoid forms. We categorized these into four types: humanoid, abstract, semi-humanoid, and other. Humanoid robots were of interest, because they mimic the anatomy and form of the human figure. Research on humanoid robots has shown that they are perceived friendly and appropriate for tasks that involve close interaction with people (Goetz et al., 2003; Siino & Hinds, 2004; Walters, Syrdal, Dautenhahn, Boekhorst, & Koay, 2008). However, humanoids are mechanically complex, and for our research, may not be robust enough for long-term use in the field.

Abstract robots were less relevant because they have a mechanical aesthetic, showing tracks, wheels, and other parts that do not invite human interaction at an intimate level. Semi-humanoid robots were of greatest interest, because they combine simple geometric forms with human cues. This was a good choice for further investigation, as the housing design would then allow for the holistic combination of hardware and aesthetic components.

We generated sketches based on the semi-humanoid concept. Two types of sketches were initially explored: more industrial, mechanical forms with wide shoulders, aggressive stance, and masculine proportions, and more playful, cute forms with rounded proportions and childlike faces (Figure 13). To support our goal of social interaction by making the robot approachable by everyone, we merged these two styles to create a gender-neutral, friendly, yet professional-looking form to fit the context of our university.

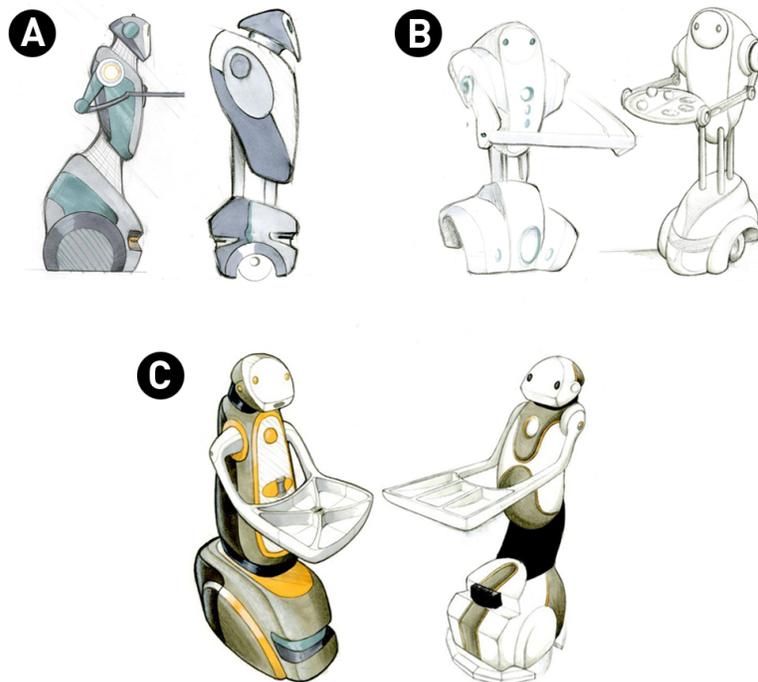


Figure 13. Sketches for the robot housing: a) machine-like, b) rounded and friendly, c) concepts combined.

We conducted empirical studies to investigate and support our form giving process. Here, we describe four of them as examples: an early technology feasibility study, an early interaction study, a dialogue study, and a height and approach study. Each of these studies was conducted in support of our overarching design goals for holistic development, product and service, and social behavior. Each generated design implications for our robot and tradeoffs with other aspects of the system. The process and results of these iterative studies are described in this section.

Early Technology Feasibility Study

We assembled some of the robot's key capabilities on an existing mobile robot platform, the CMAssist robot (Rybski et al., 2007) (Figure 14). The goal of this study was to test and verify the basic functionality of the major components of the system, and to ensure that it would work smoothly with the wireless network in our buildings.



Figure 14. Prototype used for early technology feasibility study.

The robot, partly teleoperated, traversed hallways for five two- hour long sessions over a two-week period, in the two campus buildings described above, and prompted passersby to take free snacks. An experimenter using a joystick about

20 meters away controlled the robot. The dialogue system was also run using streamed audio and five human-controlled utterances, allowing us to quickly understand the timing and robustness for this type of dialogue system. To help our technology prototype look like an aesthetic robot design, we created a housing with vacuum-formed materials and foam core components (Figure 15).

This early trial helped us learn about many tradeoffs we would face in the future design of the hardware, software, and interaction design of the robot. We subsequently decided to use a commercially available base for the Snackbot. A Pioneer base would be more reliable than a home-built base, and would provide mobile functions that would be easily replicable. It would also be quieter and less distracting to office residents. One drawback of using this kind of base is that it would create a set of constraints for the final industrial design of the robot housing. Such constraints included the dimensions of the robot, the availability (or lack thereof) of mounting points for the torso, and maximum load that could be carried. Our plans for the torso and other electronics exceeded the recommended weight limit of the Pioneer, and so later experiments were performed to learn the maximum reasonable weight the robot could carry while still having reasonable, operational battery life.

In terms of software, we learned that it could be feasible to entirely automate the dialogue structure using a finite number of preset phases because conversation with the robot quickly revealed stable patterns – a sequence of greeting, selection of snacks, and payment. We also learned that we would need to devise ways to deal with network lag or drop-off and still preserve the idea of a sociable, fluidly interacting robot. This led us to pursue the interaction study described in the next section.

Early Interaction Study

Our early interaction design study took the form of three semi-structured trials with the first robot prototype in two campus buildings. Here, our goal was to

come up with archetypical dialogue structures for interacting with the Snackbot, to support our design goals of product and service and robust social interaction.

We used Wizard of Oz methods, where a remote dialogue operator used Skype and interactively “chatted” with snack customers. A separate operator performed motor control of the robot using a joystick and tether. We adopted the convention of American ice cream trucks, and developed a 30-second melody and a cheery “Hello!” for the robot to announce itself in the hallways. Interaction with customers was structured in that the Skype operator had a script to follow, but could deviate from it in real time if needed.

We learned that people found the melody and greeting to be too annoying for use in an office building. This was partly due to the fact that the sound was played from a low-quality speaker, and therefore distorted, but the social norms of an office environment also played a role. We also learned that a minimal, straightforward design of the dialogue would be all that is needed, because people readily filled in dialogue and other social cues, such as indicating which snack that they intended to take off the tray by showing it to the robot’s eye cameras, and by politely repeating phrases during their interactions (Figure 15). These findings suggested methods for collecting speech and environmental sound as input for the dialogue system, and gave us ideas for how to specifically design and study the dialogue system, which we describe in the next section.



Figure 15. Typical interaction sequence observed during the early interaction study.

Dialogue Study

We next conducted a study to verify our design of the dialogue structures and scenarios. Our overarching goal was to discover how to provide dialogue with the robot in a way that evokes social behavior and allows the service to proceed as intended.

We created general dialogue excerpts and ran them in a Wizard of Oz study with 12 participants. One experimenter ran the robot's dialogue scripts in a remote location, and another noted what the participant said in response to the scripts. We used the stationery mode as a scenario for the study — passersby approached the robot and discussed what snacks were available that day.

We learned several things about our first iteration of the dialogue design. First, nearly half of the phrases we designed were unsuitable in that people frequently deviated from the script as we designed it. We added phrases to control for unintelligible speech or users wandering off topic. We also learned that people liked to play with the dialogue structure to see where it might fail. For example, if the Snackbot asked, "Is this your first visit?" a participant might answer "I have been here lots of times but I have never seen you," instead of giving a simple yes or no answer. Although we tried to structure the dialog to discourage such behavior, we were unsuccessful. We subsequently added phrases to try to smooth over these communication breakdowns.

We found that care needs to be taken in constructing the output phrases so that they are intelligible and imitate human intonation. Although our synthesizer is state-of-the art, certain words, phrases, and spellings can result in difficult to understand speech. The synthesizer has trouble particularly with the rise of voice expected when people ask questions. For example, "Would you like an apple?" sounds strange with synthesized speech intonation. Thus we learned the Snackbot should instead say, "I want to know if you would like an apple," to eliminate intonation issues.

We found that some participants used visual cues much more than others, thereby minimizing the use of dialogue. In particular, they tended to examine the tray rather than asking what snacks were offered, and to simply remove the item without verbally indicating what they would like, despite a direct question. We learned that we would need to tightly couple the dialogue system with the sensor system to adequately track all of the non-verbal communication in support of evoking social behavior.

Other interesting social interactions were observed, such as groups of people interacting with the robot. Group conversation presents a difficulty for the speech recognition system, which is unlikely to differentiate person-to-person conversations from those targeted towards the Snackbot. Some of this difficulty can be mitigated with careful integration with other sensors. To best understand where to place these sensors, we undertook a height and approach study described in the next section.

Height and Approach Study

Rather than arbitrarily deciding the height of the robot, we wanted to learn whether the height of the robot affects people's approach interactions with the robot. To our knowledge, there have been no formal studies about the body size of a robot. Therefore, we conducted a study to discover what an appropriate height might be for the Snackbot.

We conducted a between-subjects experiment with 72 participants using the technology feasibility prototype described earlier. The robot had three height conditions, 44 inches (112 centimeters), 50.5 in (128 cm), and 56 in (142 cm) (Figure 16). We chose these three heights as deviations from the average height of a small human being with an average reach of lower arm length, so it would be comfortable to approach and take a snack from the robot even in the shortest condition. We did not want to make the robot taller than people in order not to be threatening.

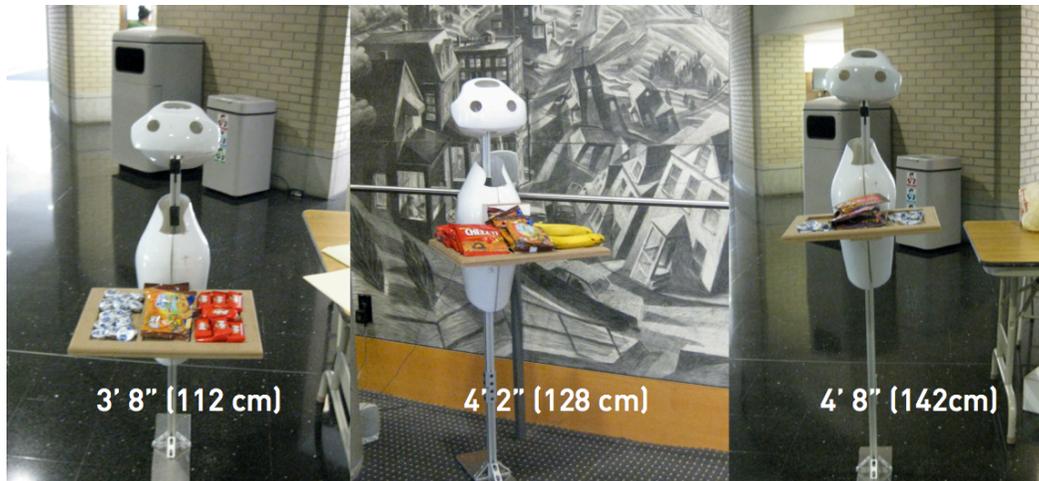


Figure 16. The prototype presented in three different heights for the height and approach study.

The study was conducted in a public area of our campus. We offered free snacks for participating in the survey. We used a 5- point Likert scale to understand how friendly and intelligent people felt the robot was, and how they responded to the height of the robot. An open-ended question asked participants to list the personality traits that they ascribe to the robot. We also asked participants their gender, age, and height.

Using a 5-point scale where 1 = much too small and 5 = much too tall, participants preferred the tallest robot most, $F(1,71) = 4.10, p < .02$. The smallest and mid-sized robots averaged 2.4, meaning they were between “too small” (score of 2) and “just right” (score of 3). The tallest robot was almost just right with a mean of 2.9. Participants liked the fact that they could make eye contact with the tallest robot, and disliked that they had to bend to interact with the smaller two robots. There was no difference in terms of how friendly and intelligent people felt the robot was across conditions. In addition, no correlation was observed between the participants’ height and the robot’s height that they preferred. However, there were interesting differences in the personality traits participants attributed to each prototype. The smallest robot most frequently was described as servile, obedient, and submissive. We felt that to best support our goal of evoking

social behavior, the robot should be seen more as a co-worker than as a servant. In our university culture, even the least skilled workers are given respect more as peers than as servants. This consideration also indicated that the tallest robot would be most appropriate.

Armed with the findings from the early technology feasibility study, the early interaction study, the dialogue study, and the height and approach study, we built the second prototype of the Snackbot.

6.6.4 Second Prototype

We embarked on designing a more robust, refined system, using a Pioneer base. This decision was made to support our design goals of offering a product and service in our office environment, by reducing the distracting noise, and ensuring operation over long periods of time and a variety of floor types.

From our interaction study, we learned that we would need to develop a set of sensors that would allow us to know when a snack was taken. Because we did not want to overload the vision system, which would eventually support person recognition, we added a mid-chest laser and pressure sensors to the robot's tray. These additions would also support natural social behavior between Snackbot and its customers.

To develop the second prototype, we focused on the development of the housing, the design of the tray, and the development of an internal structure to anchor the sensors, laptops, and housing to the Pioneer base. We also focused on the expressive qualities of the robot's face, and finalizing the interaction design.

Housing

Working from the early sketches described above, we built a number of quarter-scale and half-scale models of the robot (Figure 17). After ascertaining correct proportions for the dimensions of the Pioneer, laptops, and other internal components, we constructed a full-scale mock-up to proportions, radii, and

design details. Using this prototype, our design team was able to address dimensions, hardware placement, configurations, assembly, and tray size and arm options.

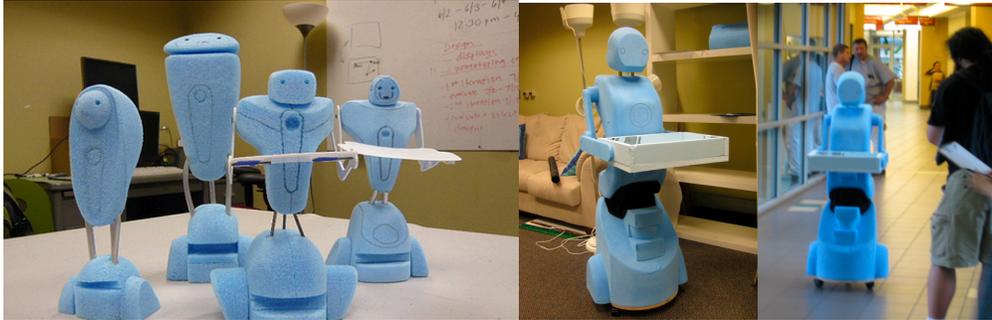


Figure 17. Quarter-scale and full-scale mock-ups of the robot.

To check responses to the full-scale model, we placed it in a hallway in our building and conducted a survey with 59 participants to understand positive or undesirable associations to the design. Participants rated the robot as friendly and likable (mean 3.88 and 3.87, a 5-point scale), and neither intelligent or unintelligent (mean 3.42). The robot evoked descriptions of service jobs such as a waiter or waitress, or general Sci-Fi characters such those from the Jetsons. Based on these responses, we felt that the final form design supported our three design goals.

One of the issues with the housing was weight. The Pioneer has a recommended payload limit of 50 lbs. for carrying additional weight. From our experiments, we determined that between 70 and 80 lbs. of weight was still reasonable for the robot to carry and still have an acceptable operational lifetime. This drove the selection of light materials such as fiberglass and neoprene fabric for the outer housing and aluminum 80/20 for the inner housing. We also segmented the base and made a variety of cuts in the torso to reduce weight. The resulting housing is lightweight, strong, and easy to add attachments for internal materials.

After generating a number of color studies for the housing, we selected a color scheme of medium gray and orange (Figure 18). Both hues do not cause gender

attributions or strong attributions of service type in the U.S. culture. For example, a blue robot might connote a medical service, due to the ubiquitous use of the color blue in the health sector. A green robot might connote a sustainable product. Orange is also often associated with food and restaurants. Together, the orange, gray, and dark gray of the neoprene creates a distinctive, impressive form.



Figure 18. Construction of housing and internal structure of the robot.

Tray

The tray design (Figure 8) was developed for providing food or snacks at the appropriate delivery height, but also as an input system for measuring the weight and presence or absence of items on the tray. The tray has movable slots that can be configured in a number of ways to hold different snacks. The tray is made of aluminum, styrene, 12 load cells, and a cloth covering on which snacks will be placed.

Internal Structure

The design for the internal structure continued to evolve as the external housing design was finalized. To minimize weight while providing maximum strength,

three vertical struts of extruded aluminum were used as the base for the design. We augmented these with a custom aluminum plate at the top of the Pioneer base and one at the shoulder, to mount internal components. These additions allowed for retrofitting to a variety of components using off-the-shelf brackets and anchors. These design decisions afford modularity, which support our overall goal of holistic design.

Head and Face

Our overall goal was to create an expressive head that would serve as a locus of interaction, relying on appropriate features that convey the right level intelligence and functionality for the robot (DiSalvo et al., 2002).

The final head design features a simple form that is wider than tall, suggesting a young, friendly robot. The width of the Bumblebee camera also determined the width of the head and the placement and size of the eye sockets. We also felt that by minimizing complexity and detail in the eyes, Snackbot customers would not develop false perceptions about the intelligence of the robot.

A 3 x 12 LED display was developed for the mouth, serving as an expressive focus for interaction. The mouth is programmed with a series of animations that show verbal and emotional feedback in the form of lip shapes, colors and movement. Although the robot does not have functional ears, we added ears to the head design, so that customers would understand that the robot could hear them.

Interaction Design

We have designed a basic interaction infrastructure, so we can use and vary these modalities to conduct experiments once the robot is fully implemented. Snacks can also be ordered through a web site or IM service. The final service design for the robot includes stationary and mobile delivery modes that provide a variety of services to our university community and support social behavior.

6.7 Lessons Learned

We have spent almost two years on the holistic development of a robotic system, and we have learned several lessons. We articulate them here, relative to our overall design goals.

Our first lesson was to understand how the robot would actually work in a context of people, other products, a physical environment, and social norms. Then, in the service of holistic design, the design of particular subsystems could be undertaken. This point is not new. Many others have articulated the need to design for the context (e.g., Jones & Hinds, 2002).

Our second lesson in terms of holistic design was to design for modularity. Functions should be developed individually, but with an eye to the constraints caused by other aspects of the system. Modularity also means that components can be upgraded or changed as new and better systems become available. For example, the selection of a Pioneer base created weight constraints, which became an issue in the design of the housing. Again, others in various fields have recommended designing in modularity (e.g., Cai & Sullivan, 2005).

In terms of product and service, we learned the robot should offer capabilities that add value to people's lives, and allow them to add value themselves through interacting with the robot. This idea drove our choice to offer healthy snacks, and to provide a stationary mode that invites people to take a walk to the robot. Future experiments will be done to understand whether and how the robot's interaction design can be modified to best support people.

In terms of social behavior, we learned to work to make a social robot sociable within the limitations of current technology. For example, we needed to make iterative changes to the dialogue system to both support fluid and natural social interaction while working with the constraints that the wireless network provided.

None of these lessons, taken independently, are new, as the HRI community will recognize. What we think is a contribution is our showing how we tried to tackle all these lessons together. The larger lesson is that designing for all these goals is what is really hard. It requires a design team dedicated to an interdisciplinary holistic design process.

Longitudinal Field Experiment of Long-Term Personalization Strategies⁸

This chapter introduces the longitudinal field experiment where we designed long-term personalization strategies for a robotic service described in Chapter 6 and evaluated them through a longitudinal field experiment. We report the experiences of employees with the Snackbot snack delivery service and robot that delivered the snacks in a workplace over a period of four months. We followed employees who participated in a field study of the service to understand their responses to the robot and to evaluate the service.

7.1 Moving Social Agents from Labs into the Real World

Computational agents such as Snackbot are increasingly designed to assist in real world tasks. Examples of current services include Aethon's hospital delivery robot, the Autom robotic weight coach, a therapeutic robot called Paro, and online social customer agents such as IKEA's Anna. Other service agents are in development such as instructor agents for language learning (Gwinner, Gremier, & Bitner, 1998), office and hospital work assistants (Bickmore, Pfeifer, & Jack, 2009; Powers, Kiesler, Fussell, & Torrey, 2007), and rehabilitation or assistive robots (Feil-Seifer & Mataric, 2005).

Most of the agents mentioned above have social skills and attributes such as human-like attributes (e.g., such as faces or speech) or social responses to human

⁸ This chapter is adapted from papers published at the HRI'12 (Lee et al., 2012), and CHI'12 conferences (Lee, Kiesler, Forlizzi, & Rybski, 2012).

input. Researchers have explored the effects of agent's social characteristics – e.g., appearance (Nakagawa, Shiomi, Shinozawa, Matumura, Ishiguro, & Hagita, 2011; Weick, 1979), conversational strategies (Bickmore, Pfeifer, & Jack, 2009; Bickmore & Picard, 2005; Brave, Nass, & Hutchison, 2005), gestures (Bethel & Murphy, 2006; Kanda, Hirano, Eaton, & Ishiguro, 2004; Powers, Kiesler, Fussell, & Torrey, 2007), touch (Nakagawa, Shiomi, Shinozawa, Matumura, Ishiguro, & Hagita, 2011) and social norm following behaviors (Nakauchi & Simmons, 2002). More sophisticated strategies attempt to match the social response of the agent to the personality of the user (Tapus & Mataric, 2008), the task (Goetz, Kiesler, & Powers, 2003) or the culture (Evers, Maldonado, Brodecki, & Hinds, 2008). Many of these studies have been performed in labs (Bethel & Murphy, 2006; Bickmore & Cassell, 2001; Brave, Nass, & Hutchinson, 2005; Evers, Maldonado, Brodecki, & Hinds, 2008; Goetz, Kiesler, & Powers, 2003) and in demonstration projects in public settings such as museums (Kuno et al., 2007). Much of this work shows that a social agent can improve people's engagement and trust toward the system, and liking of the agent, even in search and rescue setting (Bethel & Murphy, 2006), a task that seems very utilitarian and task-oriented. Additionally, studies on entertainment or commercial robots (Turkle, 2005) show that people relate to a robot dog (Friedman, Kahn, & Hagman, 2003), or even to a vacuum cleaner (Sung, Guo, Grinter, & Christensen, 2007) and form variety of relationship, rather than treating them as pure machine.

It is not only the characteristics of an agent that affect its effectiveness. Previous work has shown that the design of the agent needs to take into account the social context of the workplace. For example, in one study of a hospital delivery robot (Mutlu & Forlizzi, 2008), the robot was perceived differently depending on the team's task and stress level. Another study with an office delivery robot (Severinson-Eklundh, Green, & Huttenrauch, 2003) emphasized the importance of designing the robot to communicate with secondary users such as passers-by in the office.

However, most of this work has been conducted in labs or public settings in which people's repeated encounters with the agent were not tracked over time. We do not know if an agent's limited social skills become annoying or boring over time, or how socially interactive systems fit into the culture of a real workplace.

Very little work has examined people's response to and acceptance of social agents that deliver services in organizational settings. (Some exceptions are studies of an English teaching robot at school (Kanda, Hirano, Eaton & Ishiguro, 2004), a conversational companion in elder care (Sabelli, Kanda, & Hagita, 2011), and an office delivery robot (Mutlu & Forlizzi, 2008). To our knowledge, no studies have examined users' interactions with agents in the workplace over an extended period. We evaluate our robotic service, tracking the same set of users' repeated interactions with the robot (Figure 19). This longitudinal deployment offers opportunities to study sensemaking and adoption of a social agent beyond its novelty effect.

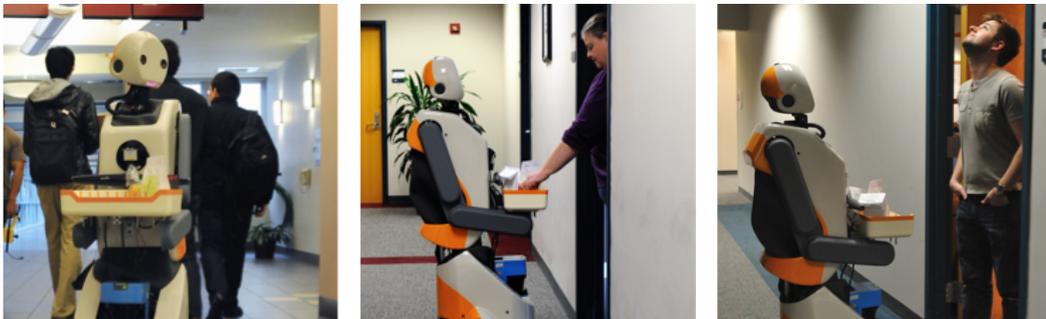


Figure 19. Snackbot carrying snack (left), delivering snacks to a participant (middle), and with a participant doing a neck stretch with the robot (right).

7.2 Personalized Robotic Service Design

Testing our personalization strategy in the workplace required us to design an end-to-end service that people would use. We designed a holistic service that comprised a website for customers to order snacks, desirable snack offerings, a semi-autonomous robot to locate offices and deliver snacks to customers, a

database of snack deliveries and interactions (Chapter 3). We also used an out-of-sight operator to choose appropriate interactions from the pool of dialogue scripts and to fix unanticipated problems with the robot.

7.2.1 Personalized Service Design Goals

Robotic system services have strong potential for assisting people with everyday tasks in workplaces (e.g., Dautenhaun et al., 2007). Examples of current services include a hospital delivery robot, a rehabilitation coach, an assistive robot for the mobility impaired, and a shopping or museum guide. A robot that efficiently and correctly provides service is a prerequisite for success. For some services, however, it may be helpful if the robot is social and builds rapport with people.

Prior work in service and human-agent interaction research suggests rapport between people and a robot is critical when a positive service outcome depends on how well people trust and cooperate with the robot (Bickmore & Picard, 2005). Stronger rapport between people and service providers was reported to increase people's satisfaction and willingness to cooperate with a service provider's recommendation and instructions (Gwinner et al., 1998). Even in services that do not require high levels of cooperation on the part of the customer, social interaction and rapport can reinforce people's satisfaction with and loyalty to a service provider (Gwinner et al., 1998). To create effective social robots, researchers have imbued robots with various social abilities. Some robots adhere to social norms (Nakauchi & Simmons, 2002), use relational languages (Bickmore & Picard, 2005), or have anthropomorphic or zoomorphic forms, so that people perceive them as an entity to which they can relate (Lee et al., 2009).

A research question that typically is not addressed in prior work is how to design meaningful social interaction and build rapport for repeated interactions. Styles of interaction successful for one time interaction may not be effective in building rapport over time. Existing work on social agents that used relational strategies over time suggests that people may lose interest in conversing socially with a

robot once the novelty effect wears off (Gockley et al., 2005). Other work suggests that people's rapport with an agent may not increase after the first contact (Bickmore & Picard, 2005).

Our design goal is to build a robotic service that could help sustain people's interest in the robot by personalizing the service over time. We argue that for repeated usage, it will be helpful for the robot to be aware of its mutual experiences with users, and to use this information to personalize its interactions over time. In this manner, interactions with the robotic service become more relevant to individuals and groups who use the service over time, reinforce the rapport between people and robot, and sustain their engagement with the service.

7.2.2 Service Components

The Snackbot service was comprised of a front end consisting of services that participants encountered directly, and a back end consisting of the underlying system that participants did not see.

Front End

Snack Ordering Website

Participants could order snacks using our snack ordering website (Lee, Kiesler, & Forlizzi, 2011) (Figure 20). They specified the snack type, delivery day, and their office number. Only those registered in the study could order snacks through the website.

Snacks

Snackbot delivered six different snacks—apples, bananas, oranges, Reese's peanut butter cups, Snickers candy bars, and chocolate chip cookies. We chose a mixture of snacks that were not always available in the workplace.

Robot

Snackbot, described in Chapter 6, was a main service platform. The robot uses its SICK LIDAR to navigate the office environment autonomously (with obstacle

avoidance and path planning). In our study, because the website information was not linked to the robot, an operator manually specified the office delivery destinations. The robot used the Cepstral text to speech program with a male voice. The robot carried a web camera and a microphone on its chest to record interactions. Speech output was controlled remotely with a laptop connected to the robot through a wireless network. Despite all our efforts, the robot had significant limitations that were evident to participants. It followed pre-set scripts. There were frequent delays in the dialogue. Sometimes the system froze when there were wireless network communication problems. However, there were no differences in breakdown frequencies between the conditions of the study.

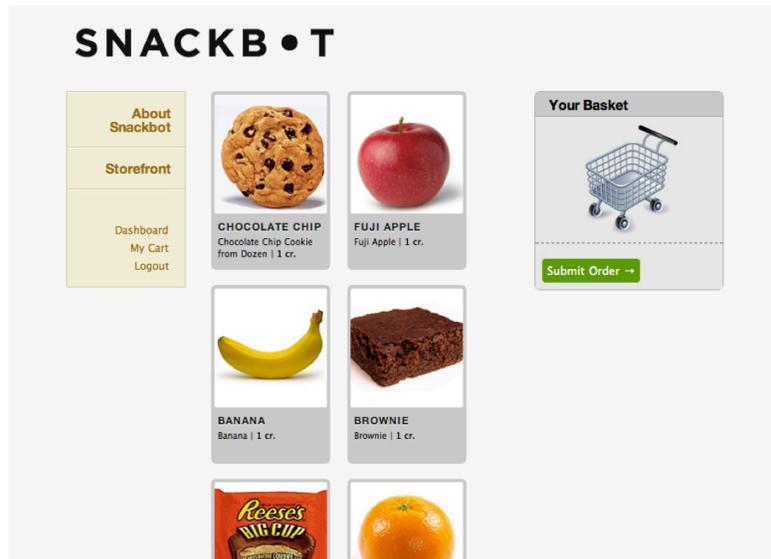


Figure 20. Website for ordering snacks for Snackbot.

Back End

Robot Control Interface

Over the previous several years, we had developed a usable interface for operators (Figure 21). This interface allowed an operator to control the robot's navigation, nonverbal movements, and dialog system remotely. The interface showed the video feed from the robot, the robot's location on the building map, its head

position, and a number of dialogue scripts. The operator could see a participant's actions through the video/audio feed on the interface.

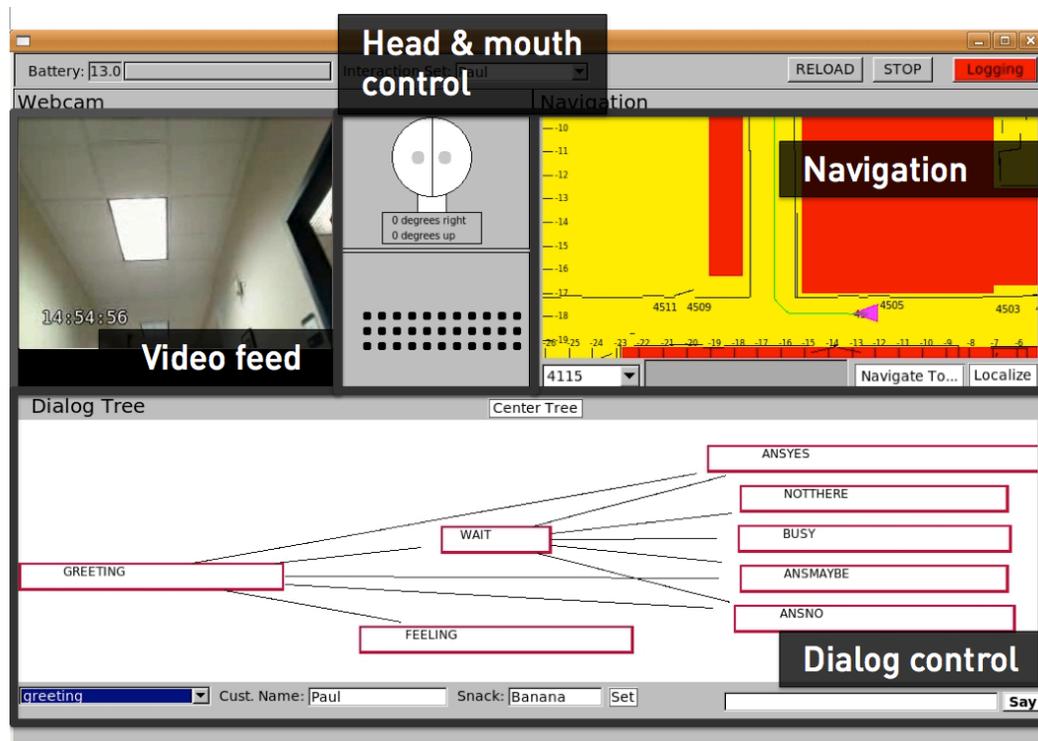


Figure 21. Interface for controlling Snackbot.

Operator

An operator transformed the orders on the website to a delivery schedule, specifying a customer name, a snack name, and an office location to the robot control interface. The operator also loaded the snacks on the robot's tray, initialized the robot at the start of each delivery run, and localized it. The operator had three designated sitting locations in the workplace building not visible from participants' offices. The operator also opened any doors in the hallways to enable the robot to go through. According to the personalization condition and interaction timeline, the operator loaded an appropriate dialogue script and clicked each node based on what the human did. To know when problems occurred, operators used a robot control interface showing a video feed of participants interacting with the robot.

7.2.3 Interaction Design

The main interactions between the service and participants occurred through participants' website orders and interactions with the robot, the latter of which became a main focus of our design efforts. We constructed the interaction scripts before we launched the service, considering the events to take place and potential user choices and behaviors.

Structure of Interaction

We created a prototypical interaction structure, informed by the observations with a hot dog vendor described in Chapter 6. Below is one of the scripts that the robot operator could use in an early day in the trial.

[At the office door] Excuse me. I have an order for David. [Robot looks straight ahead.]

Hello, David. Nice to meet you [Robot looks up to make eye contact with David.]

{...social interaction...}

Please take your apple. [Robot looks down at the tray and then looks up at David.]

Thanks, David. Bye, I'm leaving now. [Robot looks straight]

The robot followed pre-set scripts, which did not allow for improvisations of the operator to maintain consistency across participant experiences. The robot's responses were constructed in a way that made sense regardless of the participants' response (e.g., "I see."), or had two alternative responses, each applied to a participant's yes or no answer. When the dialog scripts did not have appropriate responses to a participant's comment, the robot said, "I have no idea," or just laughed, "ha ha."

Social Interactions

We created interaction dialogues that fit a workplace context, so the robot would be perceived as a member of the work organization (Table 16, Figure 22). The robot's responses also were designed to be agreeable, to emphasize similarity and honesty (e.g., admitting the inability to understand many topics).

| <i>Categories</i> | <i>Topics</i> | <i>Examples</i> |
|---------------------|--|---|
| Temporal & seasonal | Days of the week, holidays (April Fool's Day, Memorial Day), seasons | "You've got something on your face! [pause] April Fool's!" |
| Organizational | Spring festival, mid-term and final exams, break | "Do you have any plans for carnival?" |
| Regional | Pittsburgh Pirates baseball team | "It is baseball season. Do you follow the Pirates?" |
| Task-oriented | Information or story related to snacks | "Bananas are a really good source of potassium and vitamin B6. Excellent choice." |
| Other | Joke, local weather | "It is a nice day today. I am glad to see you again and hope you are doing well." |

Table 16. Social small talk topics.

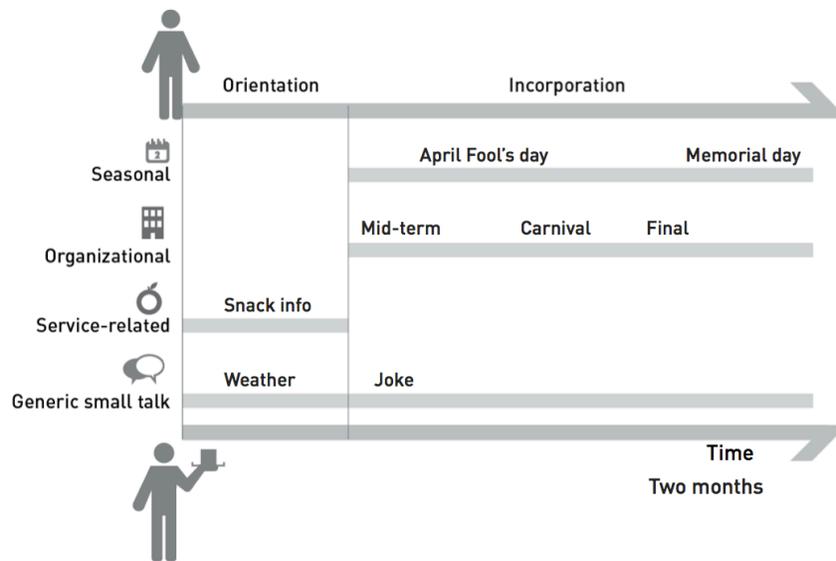


Figure 22. How interactions using small talks unfolded over time.

Personalized Interactions

For half of the participants, we built dialogues and planned interactions that used information from their prior interactions with the robot and snack deliver service (Table 17, Figure 23). We focused on users' snack choice patterns, service usage patterns, and the robot's prior behaviors. We did not personalize the interaction based on what participants said to the robot because it was not realistic with the

current level of language technology. For the robot to personalize its interactions with participants, it had to be aware of its own prior behavior. One main way we accomplished that was to maintain a record of all breakdowns and mistakes in the service database so the robot could apologize for prior malfunctions. (In Chapter 5, we have shown that apology can be helpful in rectifying mistakes).

| <i>Categories</i> | <i>Topics</i> | <i>Examples</i> |
|------------------------|---|--|
| Snack choices | Users' favorite snacks; whether they stuck to healthy snacks; whether they seemed to like variety; group's snack consumption patterns | "By the way, it seems as though you really like [snack name]. This is the [nth] time you have ordered one. Are [snack name] your favorite snack?" |
| Service usage patterns | Whether they were regular weekly users; had they been in their office when the robot was there; times when they did not use the snack service | "I missed you during my snack deliveries [n] times so far. I am glad to finally see you again." |
| Robot's behaviors | Frequency of breakdowns and apology (no breakdowns to frequent breakdowns) | "I was thinking about my first month here. I realized that I broke down and made mistakes [n] times in front of you. Sorry for that, and thank you for being patient with me." |

Table 17. Personalized conversation topics.

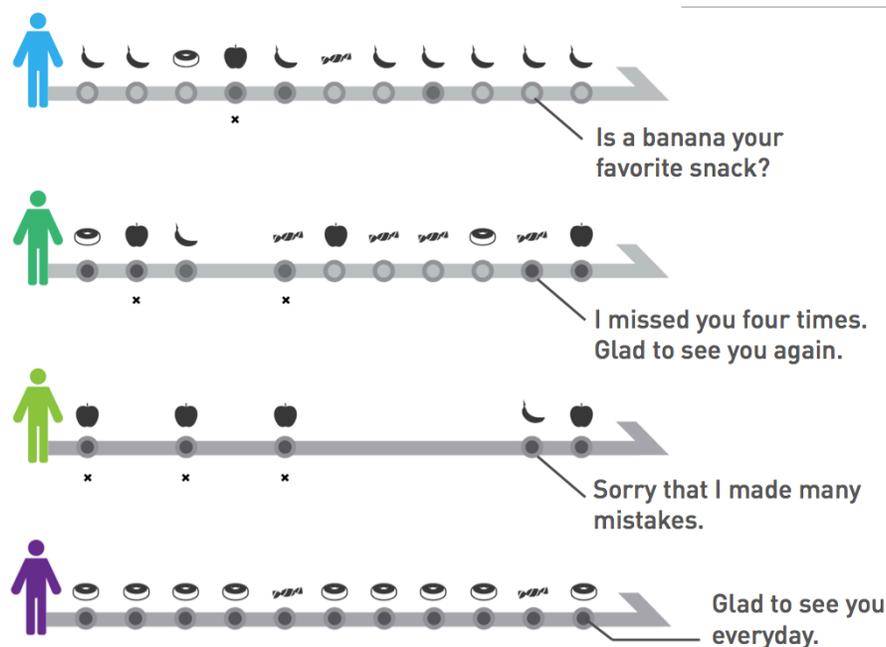


Figure 23. Examples of personalized conversation topics based on individuals' usage patterns and snack choices.

Guiding Interactions

The current level of technology was not conducive to participant-initiated conversation. Therefore, the robot's interactions were designed to guide interaction. For example, instead of giving participants time to initiate conversations, the robot attempted to lead the conversation, for example, by asking questions. To address situations where the robot could not process human behaviors, the robot used dialogues to encourage participants or passersby to behave in a manner that it could process. For example, the robot sometimes said, "Can you please stand in front of me?" and "I have bad ears, so sometimes I cannot hear very well. Can you repeat, please?"

Exceptional Use Cases

Pretesting pointed to several situations other than snack transactions that the robot had to be prepared to address. For instance, some passersby took snacks from the tray without the robot's permission, or intentionally blocked the robot's path. In these cases, the robot made comments such as, "Please don't be rude. I am just a robot," and "please return the snack to a proper place. I have the campus police on my speed dial... Just kidding." Sometimes the robot broke down and stood in the hallway until it was debugged. In these situations, the robot communicated its status to people who approached, such as "I am not feeling well; my operators are fixing me."

7.3 Method

We conducted a field experiment from February to June, 2011 in a workplace to test the following hypotheses:

H1. A personalized social robot will increase rapport and cooperation with the robot as compared with a sociable robot lacking personalization.

H2. A personalized social robot will increase engagement during the service encounter as compared with a sociable robot lacking personalization.

H3. A personalized social robot will increase satisfaction with a snack service as compared with a sociable robot lacking personalization.

7.3.1 Field Site

Our participants were distributed across 16 offices located in 10 hallways on one floor of an office building at Carnegie Mellon University. We randomized the assignment of conditions to hallways because participants within hallways could hear the interactions of the robot with their office mates or with those in adjacent offices. This adjustment assured non-contamination across conditions but did not allow for randomization at the individual level.

7.3.2 Experimental Design

The study was a two (Personalization vs. No Personalization) x two (Pre-personalization [Period 1] vs. Post-personalization [Period 2]) mixed factorial design (Table 18).

| <i>Condition</i> | <i>Period 1 Robot Behavior</i> | <i>Period 2 Robot Behavior</i> |
|--------------------|--------------------------------|---|
| Personalization | Social interaction | Social interaction + personalized interaction |
| No Personalization | Social interaction | Social interaction |

Table 18. Experimental design.

We used interactions in Period 1 to collect baseline attitude scores and interaction behaviors. Baseline behaviors also were used to personalize the interactions in the Personalization condition. In general, Period 1 included each participant's first four interactions with the robot, and Period 2 included the rest of the interactions. However, for those who joined the service later (two in Personalization, three in No Personalization), we had to shorten their Period 1 to 2-3 interactions because at the end of June offices were being moved. Participants joining the later study who had 2-3 interactions in Period 1 were equally distributed across the conditions. In Period 2, participants interacted with the

robot 5 times on average (Personalization $M = 5.67$ ($SE = 0.59$), No Personalization $M = 4.33$ ($SE = 0.51$), $F(1,20) = 2.96$, $p = .1$).

7.3.3 Participants

We used flyers, postcards, and a snowball sampling method to recruit participants. The study required participants to have offices in our field site, and generally to be in their offices 2:30 p.m. to 4 p.m. at least one day a week. Thirty-two participants signed up; eight participants never placed an order, one participant left the organization, and two participants in the Personalization condition dropped out after two deliveries due to the inconvenient delivery schedule. We ended up with 21 participants, nine in the Personalization condition and 12 in the No Personalization condition. There were eight women ranging in age from 23–49 and 13 men ranging in age from 22–51. The participants included eleven graduate students, eight staff members, one post-doc, and one faculty member. All were members of a computer science school. Only one participant had prior exposure to the robot. Knowledge of programming did not statistically differ in the two conditions. Knowledge of robotics was a little higher in the No Personalization condition, but not statistically significant.

7.3.4 Procedure

The robot delivered snacks from 2:30–4pm Mondays, Wednesdays, and Fridays. We provided snacks for free to compensate users' participation in surveys and interviews. Participants could place an order anytime before noon on the day of snack delivery. If participants were not in their offices, their snack was placed in a paper bag and hung on their office door. Because we could not deliver snacks to all 21 participants in a day, those who joined the service early were retired from the study after two months of usage.

7.3.5 Data Sources

Interaction logs

The robot's camera and microphone recorded all interactions between the robot and participants. Except for one day when the robot's recording was turned off accidentally, and a few other cases when the camera was turned away from participants, 175 interactions were audio recorded and 161 interactions were video recorded when participants were in their offices.

Surveys

Participants completed a background survey after registering for the study, robot and service evaluation surveys at the end of Periods 1 and 2, and an exit survey. The background survey included questions about participants' demographic information, their snacking routines, and their orientations toward services, adapted from (Lee et al., 2010). The evaluation survey included self-report measures of rapport development adopted from (Aaker et al., 2004). The exit survey measured participants' overall satisfaction with the service, and checks on the manipulation of personalization.

Interviews

The first author conducted 30–60 minute semi-structured interviews with the 21 participants at the end of the study. The interview began with questions about participants' experiences with the robot and the service. Then, we asked participants how they felt their experiences with the robot changed over time, whether they saw other participants interacting with the robot, how other people around them behaved, what types of breakdowns they experienced and how they reacted to them, what they liked and disliked about the service, whether they had any concerns about the service. To avoid biasing the interview, the protocol did not include explicit questions about personalization. All but one participant consented to audio recording of the interview.

7.3.6 Measures

Participants' Service Orientation

Our previous work showed that people's service orientations influenced their reactions to and satisfaction with a robotic service (Lee et al., 2010). Therefore we used items (7-point Likert scales) from (Lee et al., 2010) to assess participants' food service orientation—relational vs. utilitarian. Using principle component analysis, we constructed a social orientation scale with three items (*Cronbach's* $\alpha = .78$), and a utilitarian scale with six items (*Cronbach's* $\alpha = .52$). Participants in the No Personalization condition ($M = 5.31$, $SE = 0.34$) had a higher social orientation than those in the Personalization condition ($M = 4.07$, $SE = 0.40$), $F(1, 20) = 5.55$, $p < .05$, so we included the social orientation scale as a control variable in our statistical analysis model.

Rapport

We measured rapport strength by using the constructs liking, closeness, and self-connection, suggested by the literature on relationship with brands (Aaker et al., 2004) and politeness (Brown & Levinson, 1987). Subjective measures were included in the surveys, and behavioral measures were taken from participants' behavior during snack delivery. We first read all interaction transcripts, identifying behaviors that show participants liked the robot and felt close to it. We do not discuss behaviors equally exhibited in both conditions (e.g., greetings). Two coders coded for the following three behaviors.

Flattery and gift giving. These behaviors convey that people are cooperators, specifically, that the speaker wants to satisfy the hearer's wants (Brown & Levinson, 1987). We coded instances when participants complemented the robot (e.g., "you are inspirational to me," "I'm glad you came.") or gave a gift to the robot; *Cohen's Kappa* = .78).

Self-disclosure. Self disclosure indicates that two people feel close to each other (Brown & Levinson, 1987). We coded instances where participants shared

information about themselves that was not solicited or goes beyond the typical response given to the robot (e.g., Snackbot: “Get ready for a new week.”

Participant B: “That’s right. We’ll see. We have a big presentation tomorrow. Hopefully we’ll be okay.”; *Cohen’s Kappa* = .70).

Greeting using the robot’s name. We coded instances when the participants greeted the robot using its name.

Closeness. The evaluation survey included two 7-point Likert items adopted from Aaker et al., 2004 (I have a personal relationship with the robot, I feel close to the robot; *Cronbach’s α* = .76).

Self-connection. The evaluation survey included two 7-point Likert items adopted from Aaker et al., 2004 (Snackbot represents the personal service that I would want, The service fits my current lifestyle; *Cronbach’s α* = .60).

Cooperation

Cooperation measures consisted of participants’ responses to three requests the robot made in three visits towards the end of each participant’s service experience. We invented cooperation tasks to meet the following criteria: participants would have to listen to the robot, comply with a request for a favor entailing a new behavior by the participant that would not happen without the request, and would be different in each case, to avoid learning or habit effects. We standardized the measures by transforming scores so that each distribution has a mean of zero and a standard deviation of 1.

Help request. The robot explained to participants that it needed to give visitors a tour of the building, and asked whether they could suggest good locations to add to the tour. We counted the number of locations that participants suggested.

Neck stretch. The robot explained to participants that taking a break has been shown to boost people’s productivity. The robot said it knew how to do a neck

stretching exercise that helps release the tension around a person’s neck and shoulders. The robot asked participants whether they would like to try the exercise. We coded whether the participants completed the exercise with the robot or not (yes = 1, no = 0).

Mystery snack. The robot explained to participants that it was carrying a special “fresh and good” mystery snack (Figure 24). The robot asked whether participants would like to try the mystery snack instead of the snack that they ordered. The mystery snacks were baked goods such as a lemon bar or cupcake that had not been part of the service. We coded whether the participants took the mystery snack or not (yes = 1, no = 0).



Figure 24. Mystery snack placed on the robot’s tray.

Engagement

To measure engagement during service encounters, we coded participants’ postures and facial expressions, which can indicate people’s engagement in social interaction (Knapp & Hall, 2009). We do not discuss measures that did not differ between the two conditions (e.g., gaze, head nodding). We did not code proxemics because we could not reliably measure the distance between participant and robot from the recorded videos.

Facial expression. We coded for instances of smiling, laughter and general facial expression (positive, neutral, negative).

Standing posture. We coded whether participants were upright, leaning against the door, or leaning forward. The frequency of leaning forward did not vary by condition. Compared to leaning against the door, standing upright is a less relaxed behavior, and indicates a positive attitude, more attention to an addressee (Knapp & Hall, 2009), and is exhibited when the addressee is of a higher status (Mehrabian, 1970).

Service Satisfaction

The exit survey included questions on participants' overall service satisfaction, their willingness to continue the service on a 7-point Likert scale, and how much they would be willing to pay per month to continue to use the service.

Analysis of Interviews

We transcribed the interviews and interaction logs and did thematic coding, using the NVivo 8 software. We followed an inductive process that involved reading through the interview and interaction scripts and investigating emerging categories and relationships (Strauss & Corbin, 1998). We started by open coding a small sample of scripts, adjusted and added categories, and then proceeded to open coding of all the data. In the phase 2, we grouped the lower-level codes into thematic clusters and drew connections among them to tell a story about how participants made sense of the robot, and how the robot changed and evoked social behaviors that created interesting ripple effects. We do not report themes that concern functional and aesthetic qualities of the robot, and ideas for new features; they were practical suggestions unique to our service platform. In this process, we compared what we were learning with existing concepts such as sensemaking and structuration. We also counted how much participants spoke and relational behaviors from the interaction logs.

Analysis of Quantitative Data

We used a multi-level regression model to analyze the codes from the interaction log, comparing responses during Period 1 vs. Period 2. For the evaluation surveys,

we used ordinary least squared regression analysis to measure rapport after Period 2, controlling for initial rapport after Period 1. For the exit survey, we used ordinary least squares ANOVA. We included the social orientation scale as a control variable in all the models, because, as noted above, social orientation differed between conditions.

7.4 Results: Effects of Personalization

Our results provide substantial evidence that personalization of the robot improved participants' service experience.

7.4.1 Overall Service Usage

There were 261 orders, on average 6 orders per day ($SD = 4.53$). On average, each participant ordered 12 snacks ($SD = 3.96$) throughout the study. The participants could order only one snack at a time. Excluding the times participants were not in their offices, they interacted with the robot 9 times on average ($SD = 3.07$). Each interaction averaged one minute and six seconds long ($SD = 37$ seconds), included 7 turns ($SD = 2.28$) by the participant and 8 turns ($SD = 2.27$) by the robot. The average number of words in participants' dialogues was 35.13 ($SD = 23.08$). The difference between conditions in interaction duration and number of turns was not statistically significant.

In the interviews, participants expressed that their initial excitement wore off after two to three interactions with the robot. Snackbot's visit became routine, as participants knew how the interaction generally unfolded, and got used to seeing the robot (Feldman & Pentland, 2003). Some people wanted to interact with the robot socially, and were disappointed if they missed it. They made an effort to be in their offices if they did not have afternoon appointments:

Participant U: Yeah, it was definitely something we added to my Monday, Wednesday and Friday routine and I was always sad if I missed it.

Participant O: I was having a conversation with a coworker about whatever it was that I was going to do that afternoon, and I realized, I heard myself say, “Well, it doesn’t matter, ‘cause I’m not missing my Snackbot visit now.”

Participant M: Oh yeah, [my boss’s] office is down the hall from mine, and I was in a meeting with him and then I heard Snackbot coming down the hall towards my office, and so I ran out of the meeting to go to my office and wait for Snackbot...

These participants expressed that social conversation with the robot created a nice break from work, adding extra pleasure and value to receiving the robot’s deliveries:

Participant J: It gives you more to talk about and it’s funny. And it’s more entertaining. Like, it’s not hard to walk to a vending machine.

7.4.2 Manipulation Check

In the exit survey, we asked participants if the robot remembered their previous snack choices (Personalization $M = 6.70$ ($SE = 0.56$), No Personalization $M = 4.31$ ($SE = 0.48$), $F(2,19) = 9.38$, $p < .01$), other customers’ snack choices (Personalization $M = 6.63$ ($SE = 0.63$), No Personalization $M = 4.33$ ($SE = 0.50$), $F(2,19) = 7.18$, $p = .02$), and how personal the service felt (Personalization $M = 6.13$ ($SE = 0.44$), No Personalization $M = 4.90$ ($SE = 0.38$), $F(2,19) = 4.01$, $p = .06$). These results show that the personalization manipulation was effective.

7.4.3 Rapport

As predicted in Hypothesis 1, recorded interactions show that participants exhibited social behaviors more frequently when the robot personalized its dialogues (see Figure 25).

Flattery and gifts. Participants in the Personalization condition were more likely to flatter the robot or to give it a gift during Period 2 ($M = 0.22$, $SE = .05$) than during Period 1 ($M = 0.07$, $SE = .05$), $F(1, 163.1) = 5.84$, $p < .05$, and more than those in the No Personalization condition ($M = 0.03$, $SE = .04$), $F(1, 34.7) = 9.16$,

$p < .01$; period x condition interaction, $F(1, 163.3) = 2.61, p = .1$). Here is one example:

Participant E: (starts laughing). I have a snack for you.

Snackbot: Please take your orange.

Participant E: I have a snack for you Snackbot. It's a battery.

Snackbot: Thanks, [participant name]. Enjoy your snack.

Participant E: Bye Snackbot.

Snackbot: I hope you have a wonderful day. Goodbye.

Participant E: You too, enjoy your snack.

Self-disclosure. Participants also disclosed more about themselves in the Personalization condition during Period 2 ($M = 0.68, SE = 0.10$) than during Period 1 ($M = 0.26, SE = 0.11$), $F(1, 162.4) = 14, p = .001$, and those in the No Personalization condition ($M = 0.25, SE = 0.09$), $F(1, 25.84) = 9.11, p < .01$; period x condition interaction, $F(1, 159.5) = 4.92, p = 0.03$).

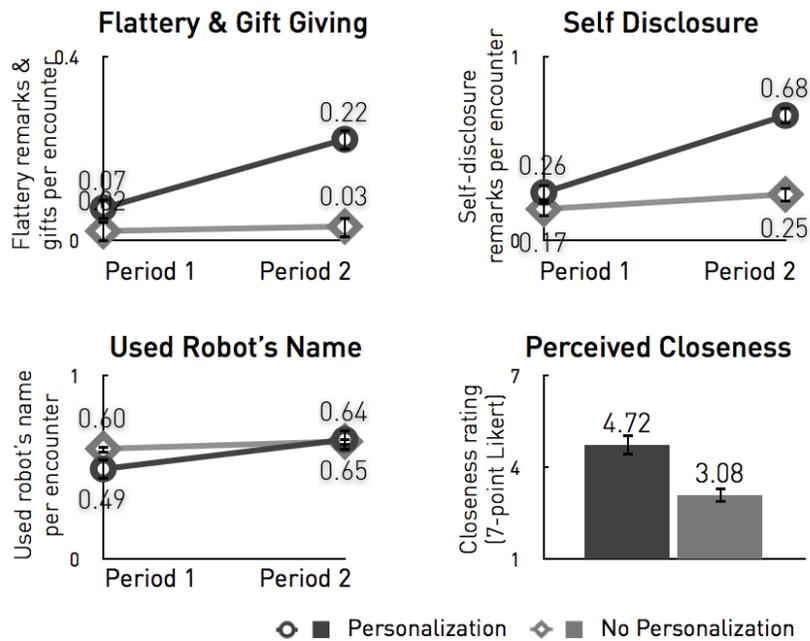


Figure 25. Measure of rapport.

Using the robot’s name. Participants in the Personalization condition greeted the robot with the robot’s name (i.e., “Hi, Snackbot”) more frequently ($M = 0.65$, $SE = 0.13$) during Period 2 than Period 1 ($M = 0.49$, $SE = 0.13$), $F(1, 143.7) = 5.23$, $p < .05$. This result suggests a potential ceiling effect, but we could not think of reasons why there would be a ceiling in this rate.

Perceived closeness. Participants in the Personalization condition felt closer to the robot ($M = 4.72$, $SE = 0.71$) than those in the No Personalization condition ($M = 3.08$, $SE = 0.52$; $F(3,16) = 3.05$, $p = .1$) but the difference was only marginally significant. Perceived self-connection did not differ by condition.

7.4.4 Cooperation

Personalization increased participants’ cooperation, as predicted in Hypothesis 2. We derived a summary measure of cooperation for each participant by standardizing scores on all three measures (see Figure 26) and calculating a mean for each person.

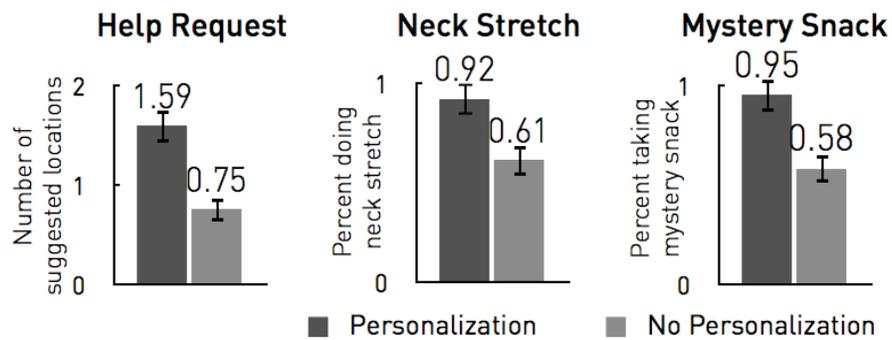


Figure 26. Measures of cooperation.

The results showed people’s willingness to cooperate with the robot was greater in the Personalization condition ($M = 0.49$, $SE = .28$) than in the No Personalization condition ($M = -0.45$, $SE = .22$), $F(2,18) = 3.48$, $p = 0.02$. We provide an example below.

Snackbot: I need to give a tour of [building] for visitors, I am still new to this building and I am not sure where to bring them. Could you suggest some interesting places in [building]?

Participant F (No Personalization condition): Snackbot, let's not be ridiculous, can I take my snack? Can I have my snack?

Participant L (Personalization condition): Let's see. You could visit the [exhibit name] on the first floor or the third floor. The second floor has a lot of cool other robotic stuff that you could check out or show people, I think they would like that [...].

7.4.5 Engagement

Participants' engagement with the robot appeared to be more sustained when the robot personalized its remarks (see Figure 27).

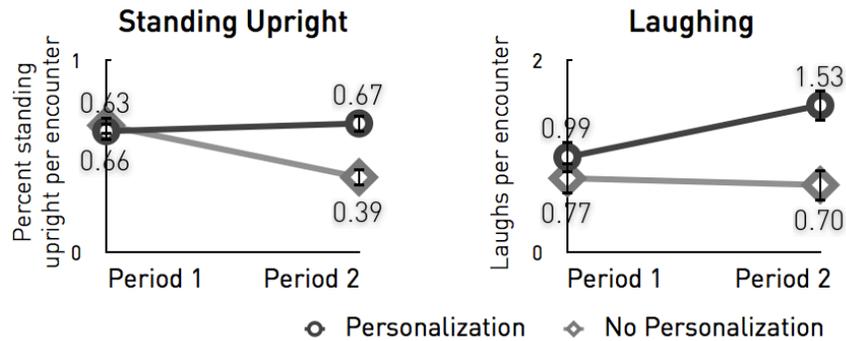


Figure 27. Measures of engagement.

Laughing. Participants laughed more during personalized interactions during Period 2 ($M = 1.53$, $SE = 0.36$) than during Period 1 ($M = 0.99$, $SE = 0.36$), $F(1,146.1) = 4.94$, $p < .05$ and more than those in the No Participation condition ($M = 0.70$, $SE = 0.32$), $F(1, 27.91) = 2.75$, $p = .10$; period x condition interaction: $F(1, 145.3) = 3.27$, $p = .07$).

Standing posture. The percentage of the participants who sustained their upright standing posture did not change over time in the Personalization condition. In the No Personalization condition, the percentage of the participants who stood upright when interacting with the robot decreased from Period 1 ($M = 0.66$, $SE = 0.1$) to Period 2 ($M = 0.39$, $SE = 0.1$), $F(1,140.2) = 11.25$, $p = .001$. More

participants in the No Personalization condition leaned against their office doors while interacting with the robot in Period 2, signaling higher status and/or less attention.

7.4.6 Service Satisfaction

The ratings of service satisfaction did not statistically differ by condition.

Participants in both conditions were highly satisfied with the service (Personalization $M = 6.05$ ($SE = 0.24$); No Personalization $M = 6.22$ ($SE = 0.21$)), and were willing to continue the service (Personalization $M = 6.40$ ($SE = 0.41$); No Personalization $M = 6.53$ ($SE = 0.35$)). Participants in the Personalization condition said they would pay more to continue to use the service ($M = \$16.19$, $SE = 4.09$) than those in the No Personalization condition ($M = \$12.4$, $SE = 3.48$), but the difference was not statistically significant.

7.4.7 Discussion

Our analyses suggest that personalizing the interactions with the robot reinforced participants' rapport, cooperation, and engagement. Our post-study interview results helped us understand how participants interpreted the personalization strategy. As noted above, in the interviews, we did not mention personalization, so the answers we received were unsolicited.

Receiving Personal Attention

Consistent with the literature on personalization, participants seemed to like personal attention from the robot. We designed Snackbot's personalization to build on real experiences between the robot and the person, creating an interaction that was unique to each participant. When the robot remembered even a small detail about a participant, for example, their favorite snack, it seemed to elicit feelings of closeness. For example, Participant N said:

Surprisingly Snackbot knows that he never dies on me. (Interviewer: How did you feel about it?) So I feel good. I feel special that I communicate with Snackbot with no problem.

By contrast, in the No Personalization condition, most participants expressed a desire to have more tailored interactions with the robot, as Participant U said:

But I felt like over time [...] if he shows up every week, Monday, Wednesday and Friday, you would hopefully learn [my] name or that the conversation would get to the point where it could be a little bit more personal.

The rapport created through personalization may have played a role in influencing people's willingness to cooperate or help the robot. Participant I in the No Personalization condition said during the interview that the robot's tour help question was one of his negative experiences with the robot:

I think it was mostly that you don't have enough of a rapport with it to answer that question. So if it was like someone—if it was like Justin or someone who works with me, I could be like "Oh we should show them the thing down in that lab where you work."

Sustaining Interest

According to the interviews, participants in the Personalization condition were more engaged with the robot over the course of the study. We surmise that the robot's interactions became more meaningful over time. For those in the No Personalization condition, interaction with the robot became less meaningful as participants realized that their conversation with the robot did not have any bearing on the robot's future behavior. This caused people to lose interest in conversing with the robot. By contrast, in the Personalization condition, the robot made comments based on its past performance or the participant's use of the service, building common ground and shared history. The robot's telling stories related to the participant each time caused excitement and expectation, as participants waited for new stories.

Participant L (Personalization condition): We even commented to each other a couple times; What do you think he's going to say today or do you think he's going to say something about carnival?

Disadvantages of Personalization

As in human interaction, personal conversation can create discomfort because people feel invested in the relationship. Some of Snackbot's personalized dialogues evoked negative responses, especially when participants felt uneasy about the behaviors that were the topics of conversation. The most sensitive topics pertained to participants' not being present when the robot arrived, and to their choice of snacks.

Participant M (Personalization condition): But then my most negative [feeling] was one time he said, "I notice that you always order Reese's Cups. You must really like Reese's Cups," and that was kind of awkward for me because it's like, "Oh, I'm the one ordering all the junk food, and eating junk food every day, and now he's pointing it out."

We were initially concerned that participants would have privacy concerns or feel more pressure to be social with the robot when the robot personalized its interactions. Participants mentioned that they did not have privacy concerns with the topics or events that the robot used to personalize. In both conditions, participants seemed to feel some pressure to be social and polite with the robot as the interactions took place in a social setting, the workplace, and others might hear these interactions.

7.5 Results: Changes in Service Orientation

The previous section explained the design of two adaptive interaction strategies and show that the personalized interaction was effective in improving rapport, engagement and cooperation. The interview results show how participants' service orientations changed over time and was influenced by social context. The participants formed either relational or utilitarian orientations with Snackbot. There were many factors that influenced this formation and changes of orientations. Interestingly, the interaction between individuals and the robot created a ripple effect in the workplace, resulting in new social dynamics within the organization and its influence on individual orientations.

7.5.1 Service Orientation with Snackbot

Over the course of the snack service trial, two notable phenomena emerged. Participants began to attach a workplace role to the robot, and incorporated the service into their daily routine. Two service orientations emerged – relational and utilitarian orientations. The interviews suggest that potential factors that influenced the formation of the service orientation, including the robot’s breakdowns, the “role” or “functionality” that the robot plays in participants lives and contexts, and the rapport and relationship built from the robot, their initial belief about the machine.

Initial Orientation with Snackbot

Before the introduction of the service, all participants seemed to have utilitarian orientation with the Snackbot service. The results of the first survey suggested that participants’ initial expectations of the service focused on its utilitarian benefits such as offering good quality snacks, and on efficient service such as getting snacks on time. Further, retrospectively, in the final interview, all participants said they had not expected to interact with a robot and recalled that they had expected it would be like a delivery cart that just left snacks. The fact that the service was quite different changed their minds. With the exception of one participant who did not want to continue the service at the end of the study, participants reported that they liked the service, talked about the service with their friends and families, and felt positively about it.

Formation of Relational vs. Utilitarian Orientations with Snackbot

In the interviews at the end of the field experiment, 16 participants reported having relational orientations with Snackbot, and five participants reported having utilitarian orientation with Snackbot.

Relational Orientation

Research shows that new or unfamiliar situations, new technology, or new services trigger a process of sensemaking, whereby people attach particular

meaning to events (Orlikowski, 2000; Weick, 1979; Weick, 1995). Our introduction of a social robot into the workplace resulted in participants attaching a social role to the robot beyond a snack deliverer (as we intended), forming extended relationships with the robot:

Participant M: Snackbot is non-judgmental, yet you can kind of feel like you have some sort of some kind of relationship. I mean, whether it'd be a deep relationship, probably not, but just that constancy.

For example, this participant's explanation shows that one factor that influenced her orientation was the role that the robot played in her life and the meaning that she attached over time. She knew, cognitively, that the robot was a machine, but still she related to this robot relationally as it reminded her of her previous coworker.

Participant O: [the robot] reminded me of a coworker that I used to have that used to stop by and, like, make sure that you got a break during the day. And so it was, kind of, interesting. Because I was, like, wow. This is just a machine that comes to visit me. But it actually makes me feel better and reminds me of people that aren't around me anymore. Which is, I think, kind of, important to me.

Relational orientations seemed to help people be more tolerant, patient, understanding of the robot's breakdowns. Breakdowns that happened during the robot's interaction interrupted the flow of conversation. These breakdowns could cause confusion and frustration, but participants who felt a connection to the robot sometimes made jokes or filled in the blanks to cope with the awkward pauses in conversation with the robot:

Participant J: Are we having a staring contest? I think you will win.

Snackbot, after 18 second delay: Please take your Snickers.

For other participants, the robot's breakdowns were entertaining, robot-like qualities that they desired in a delivery robot. For example, Participant J said:

If he had just come and, you know, had a nice little conversation and given me the snack, I actually don't think I would've liked it as much as I did. ... But if it's just, sort of, cutely robotic in a way where it's not able to accomplish what a human could. Then, it's, like, better than if it was just really, really good at what it did, I think. Because ultimately, you know you're interacting with a computer. You're not going to be tricked into thinking it's a person.

Utilitarian Orientation

Five participants liked the service but reacted negatively to having a social conversation with the robot. Three of the participants said social interaction with the robot was meaningless because the robot was not a creature or a person, and that social interaction was not something they desired from a delivery person in any case.

For example, Participant C, who described the robot as “an ATM that dispenses snacks,” said:

Yes, I know the robot's not a person that's going to miss me so it's like somebody has programmed it to say "I'm going to miss you," and it's just like funny in a way, it is, but it's not meaningful.

For this participant with a utilitarian orientation, his cognitive belief that the robot is a programmed machine played the biggest role.

For other participants, their existing orientation about delivery service seemed to influence their orientation with the robotic delivery service. For example, Participant A said:

Participant A: “Do you want a service robot to be very conversational?” ...I'm a little reluctant with these human analogies in general, but in the sense that if you're in a hotel room and somebody knocks and says, “Room service,” you don't start chatting with them.

The robot's inability to carry out natural conversation also contributed to these reactions. All participants expressed that the robot should know if they were hurried or busy with work before it started social conversation:

Participant D: I like the snack delivery thing. Sometimes I would actually come to campus just because I ordered a Snackbot snack, and I liked to be here when he showed up. Other times, I was kind of cranky and didn't feel like talking to him and sort of wanted to just grab the snack and walk away. But I felt bad, so I didn't do that.

Changes in Orientation

A few participants self-reported their changes in orientation that happened over time. These changes were often triggered by events that disconfirmed their initial orientation such as breakdowns (Weick, 1995). As with any technology used in real world settings, robot breakdowns were not an uncommon event.

Breakdowns were occasions for people to change their conceptions of the robot and to reevaluate their connections to it. For those who had a connection to the robot, breakdowns shattered the illusion of the robot having social intelligence. For example, Participant M reported the transition from relational to utilitarian orientation was caused by incidents that where a robot was talking to a wall.

Interviewer: Any suggestions for the next version of the service?

Participant M: Not to talk to a door. . . I thought it was sad. Talking to a door, you know it's undignified. . . you know just in general, don't embarrass yourself, you're supposed to be a human here. You know, don't ruin the illusion.

On the other hand, one transition from utilitarian to relational orientation was caused by the robot's apology. Participant E told us that she did not care about the robot for the first few weeks; but once the robot apologized to her, she started to feel different about the the robot:

Participant E: But the one thing that really shocked me was the day, it was a few weeks ago, when he came to the office and said that he was embarrassed because he broke down the first few times in front of my office. And I was, I felt bad for the robot. And suddenly, I noticed I was suddenly thinking it was a person, or reacting to him like a person.

Participant E put a flashlight battery in the robot's tray, as a gift during her last visit, in case the robot would run out of battery life as had happened during a prior delivery run.

7.5.2 Ripple Effects: Influence of Social Context on Orientation

The interviews and recordings of the interactions suggest that one-to-one interactions between the robot and individuals created a ripple effect within the workplace, resulting in new social practices among employees. Participants' individual service orientations were also influenced by these social practices. The convergence of social norms into a polite model and group's social interaction encouraged (influenced) participants to have relational orientation with the robot.

Formation of Norms through Interpersonal Influence

All robot deliveries happened at employees' office doors, and conversations with the robot could be overheard by officemates or passers-by. Both participants and non- participants eavesdropped and observed others' interactions with the robot. Surprising to us, these behaviors continued throughout the service deployment. Recorded interaction logs showed that often one or more people were watching when someone interacted with the robot. When something out of the ordinary happened, for example, if the robot made a funny comment, observers laughed or remarked about the incident. On one occasion, the robot came to a participant's office door while the participant remained at her desk and yelled at the robot. At this point the robot said, "Please stand in front of me." Everyone in the office laughed.

Some participants said they felt self-conscious or awkward when others overheard their interactions:

Participant J: If people were in the hallway or across in their offices, and you're just, sort of, the spotlight's... on you a little bit when he comes to your door.

In overhearing and observing other participants' interactions, employees developed a consensus on how a typical interaction should unfold and the types of inputs that the robot could understand:

Participant B: I think definitely seeing maybe what worked when people interacted with him and what didn't kind of like primed you like how or things you should kind of say or could say to Snackbot in order for him to understand you.

Participants learned to be polite to the robot. For example, they waited until the robot was finished speaking, took snacks only after the robot invited them to do so, and did not make impolite remarks:

Participant R: I think I was a little bit meaner to Snackbot before I saw [Participant O] talking to him. I was like, "Oh, she's actually really nice and she says bye properly and "Have a good day," whereas I'm just like, "Bye Snackbot." After I saw her, I was like, "Oh, I should really be nicer to Snackbot.

The analysis of interaction logs of the participant above shows that, in her earlier interactions, she took a snack before the robot was finished talking, and used more directive language (e.g., "Snackbot, go away."). In her later interactions, she was more conversational and polite.

Robot as Member of the Workplace

After a few weeks, some participants in the workplace began treating the robot like a member of the workplace, and it became the norm to protect the robot from criticism. For example, the interaction log of Participant N shows when he complained that the robot was slow, his officemate made excuses for it:

Participant J: Hey, it's Monday.

Another participant talked about this phenomenon as follows:

Participant F: Snackbot doesn't have feelings but I wouldn't want to just take the snack and shut the door in its face. Or one time I told Snackbot--I think Snackbot asked me if there was maybe a tour of [building] or something, which room should Snackbot take me up to, and I just told Snackbot that probably someone would program it. It's a robot. It's probably not going to make those choices. And then my office mate was like, "Oh. Now you've gone and made Snackbot feel bad." So I think part of it is about how my relationship with Snackbot is not just about Snackbot but about other people who are around and kind of see us.

In the subsequent visit after the incident reported above, Participant F apologized to the robot.

The behavior logs show participants exhibited more relational and in-group member interactions over time. On average, more percentage of participants made meta-relational comments during their interaction (e.g., using “us” or referring to the robot as “friend”) in Period 2 ($M = 0.25$, $SE = 0.05$) than Period 1 ($M = 0.13$, $SE = 0.05$), $F(1, 161.3) = 4.56$, $p = .03$. Significantly fewer percentage of participants took snacks before the robot gave permission to do so in Period 2 ($M = 0.18$, $SE = 0.05$) than Period 1 ($M = 0.05$, $SE = 0.04$), $F(1, 136.1) = 7.73$, $p < .01$. Finally, they smiled more frequently during the interaction in Period 2 ($M = 1.51$, $SE = 0.26$) than in Period 1 ($M = 2.10$, $SE = 0.25$), $F(1, 136.9) = 8.44$, $p < .01$.

One issue was that our interaction design did not allow for an easy way for people to interrupt and end a conversation. Participants may have felt some social pressure to be polite, even when they wanted to end the dialogue because they were busy, or the robot was experiencing a delay:

Participant N: It's kind of awkward because when [the robot] crashes you don't know what to do because sometimes it turns away and you're trying to take the cookie or something and then people will be like why are you stealing from Snackbot? Snackbot didn't ask you to take the cookie yet.

In all social groups, people develop feelings around fairness and the distribution of resources (Diekmann, Samuels, Ross, & Bazerman, 1997). In one hallway where five participant's offices were located, perceptions of the robot as a workplace member developed to the point that participants seemed to think that the robot had personal preferences for some workers and felt slightly envious when the robot seemed to prefer others. For example, a purely mechanical decision, such as the order of office visits, was interpreted as evidence for the robot's preference.

Participant L: I don't know if it was numeric or just alphabetic or whatever it was and we thought "Oh, why he always goes to her first because he likes her best."

Participant E: I think he's flirting with her. I wonder if he likes her. Because he seemed to talk to her longer than anyone else.

The analysis of the interaction logs showed that the robot spoke the same amount of words to Participant J as to any other participant.

When the robot made the mistake of calling Participant J's name at a different participant's office, participants interpreted the mistake as additional evidence for the robot's "crush" on Participant J.

Participant L: We were kind of all at our doors here looking this way and [then] he then went over to Participant M's [office] and asked for Participant J again . . .and we all said "I knew it! I knew he has a crush on [Participant J] because he keeps looking for her." I think it was because we thought he was talking to her more than he was talking to the rest of us. That's what made us first think. We said "Oh, gosh. He says so many different things to her."

Being the first to receive a personalized interaction from the robot also made some participants feel special:

Participant L: Like when he had the mystery snack for me and he hadn't given it to anybody else.

Personalization strategies also contributed to social comparison:

Participant J: Yeah, I think that the robot complimented one girl, [Participant E], one lady, on always being in her office. And how she must be a hard worker. How he would miss her and things like that. And then, I felt a little jealous.

Deliveries as an Occasion to Socialize

The initial survey and the exit interview included questions about participants' snacking routines before the study. Participants ate snacks during long afternoons, usually at their desks; many made individual trips to vending machines without socializing. This practice may reflect US workplace culture that values efficiency. Snack consumption increased when the service was used to get

snacks regularly (to have healthy snacks or curb hunger), and did not change when the service substituted snacks that participants used to bring from home. Participants' experiences did not differ by these snacking patterns.

In one hallway where many participants' offices were near each other, participants began to routinely socialize when Snackbot visited, calling the days the robot made deliveries "Snackbot days":

Participant N: I really liked, enjoyed the Snackbot. And it has been like in the hallway of like the [building] [room number], everyone is looking forward to Monday, Wednesday and Friday. They call it Snackbot day. Sometimes I go into the office and people will be yelling today is Snackbot day.

Participant J: I'm just finishing up my first year over here. And people, kind of, mostly keep to themselves. And a lot of times, people aren't even in their office. And I think people might've even been showing up more to get the snacks. So it's usually pretty, like, quiet in my hall. You know, even if people are in, they might close their door or something. But I think people are more likely to be around and laughing and feeling sociable when the robot was there.

Participants' responses suggest that the robot became a common boundary object that participants could easily relate to, creating a topic of conversation and an occasion to socialize, in the way that dogs do in a public park (Robins, Sanders, & Cahill, 1991).

In another hallway, a few participants who shared a lab space started impersonating each other during the Snackbot visit when the participants who ordered a snack were not in their office. While doing this, they usually mimicked personal characteristics such as tone of voice and accent to entertain themselves and other passersby.

Participant B: Let's see, who was I? I was Participant S who wasn't in the lab. And my other friend, [Participant U], I think he was [Participant I]. Participant I is Australian, so he tried to do an Australian accent. But Snackbot didn't seem to like that. <laughing>. [...] I guess I tried to impersonate his mannerisms and the way he interacted with Snackbot. I mean, it

really wasn't really for entertainment purposes with the robot. It was more for the other people that were in the office.

When probing further, the participants who impersonated each other could not explain why they started; they said someone started and it seemed fun. It became a pattern to imitate anyone who was not in the office when the robot came to make a delivery.

7.5.3 Discussion

Our findings show employees attached different social roles to the robot beyond a delivery person as they incorporated the robot's visit into their workplace routines. Beyond one- on-one interaction, the robot created a ripple effect in the workplace, triggering new positive and negative behaviors among employees, including politeness, protection of the robot, mimicry, social comparison, and even jealousy.

The ripple effects were quite unanticipated, and they lasted and grew richer over time. This was not our design intention. The initial purpose of this study had been to evaluate the feasibility and usefulness of a social robot to perform delivery services in a workplace, and to examine how the robot's interactions could be designed to support repeated interactions with customers. Yet, we gathered a great deal of evidence to support the fact that social dynamics around the robot and service evolved. In the following sections, we discuss different aspects of the ripple effect, and how the design of the robot and the workplace culture contributed to these effects. We believe that this result is partly due to the interaction and service design of the robot, with its repeated travels and conversations through the workplace.

The robot interacted through conversations that could be overheard, causing people to pay attention and to observe what was going on. The robot's mobile form also made it easy to be noticed, as compared to a screen agent on a kiosk or

a computer. Perhaps another influence was the afternoon delivery time, possibly more conducive to socializing than mornings.

To make the robot more sociable and interesting, we designed the dialogues to change over time, using different topics, and (for half of the participants) building on prior events to spark more personalized interactions. We think if the robot had enacted the same dialogues for four months, interest would probably have flagged.

Our results also suggest that the decision of anthropomorphic vs. non-anthropomorphic systems has tradeoffs, and social qualities should be employed adaptively depending on individual preferences and situational contexts (e.g., busyness). The literature suggests that services can be successfully transactional or social depending on the situation (e.g., a postman who delivers mail to a large city apartment vs. a postman who delivers mail to a small community over time, and is treated as a community member) (Gwinner, Gremier, & Bitner, 1998). We explored a social interaction model appropriate for our workplace context – the same robot visiting people’s offices repeatedly, unfolding social interaction over time. To our surprise, 75% of the participants appreciated these interactions over time. However, for 25%, social interactions were a reason to devalue the service as it incurred interruption or did not match their conception of the robot as machine.

7.6 Implications for Design

In this section, we briefly address how to design successful personalization for repeated interactions. We then discuss how we could harness interpersonal interactions in the workplace and social dynamics that unfold around technology.

7.6.1 Using Personalization Strategies

When To Use Personalization

Human-robot interaction may benefit from personalized behaviors when it is important for the service to track and be aware of past service events. Customers know the business has a record of interactions and may expect a social robot to reflect these past interactions. For example, a snack delivery robot in a nursing home could be aware of what time meals were last served. Personalized behaviors may also be useful when the robot needs to be persuasive, for example, in choosing a healthy snack over an unhealthy snack, or when the robot needs help or input from customers (Rosenthal, Veloso, & Dey, 2012). Personalized behavior will be also useful in situations where the robot is assisting people doing boring and repetitive tasks since personalized behaviors over time can create surprise, joy and more engagement.

How To Use Personalization

We suggest personalization is best used to define a meaningful relationship between a robot and a person. As we learned in our study, the events that are selected to make meaning must be chosen carefully. For example, comments about liking a particular kind of candy were embarrassing rather than meaningful. Like human interaction, not all facts bear repeating. Consideration must be given to what critical moments in an interaction are and how they can be detected. For example, an assistive robot in a care facility might call out moments of independence and ability to complete activities of daily living rather than breakdowns or calls for assistance.

Challenges and Opportunities

Individuals differ in their receptivity to personalization. It will be important to develop mechanisms to detect responses to specific strategies and ways for the robot to recover from mistakes. Personalization also offers new opportunities in services. One avenue for research will be to investigate personalization unique to

robots; for examples, unlike humans, the robot has a perfect record of past interactions. In a setting where a human could not easily employ personalization techniques (e.g., a vendor in a big store), robots can personalize their interactions and change the dynamics of encounters. Another interesting avenue is self-aware robotic services. Compared to systems personalized to users' tracked behaviors, our attempt to use the robot's own tracked behaviors to personalize its interaction is relatively new. Our study suggests that it can be a promising area for the design of repeated interactions.

7.6.2 Leveraging the Ripple Effect

Much discussion of social agents has concerned their immediate effect on individuals and tasks (e.g., Bickmore & Picard, 2005). We believe positive ripple effects instigated through group interaction can be anticipated and leveraged to help members of an organization to adopt, and adapt (Orlikowski, 2000) new technology in the workplace. Here we present a few factors that are important in promoting positive ripple effects.

Making Interaction Visible

Our findings show several benefits of having the interaction between social agents and people in a place where other people can overhear or join in. This visibility of interaction helps people learn how to interact with a novel system by providing examples, and developing usage norms based on a group consensus. Increasing visibility of interaction can create a passive form of socialization, for example, as happens often in online communities when newcomers watch how old timers interact. Many online community members derive entertainment and learning benefits from watching other people's conversations online. Watching prepares them to join the interaction and socialize with others later. Features like embodiment, interaction location, and timing can be used to increase the visibility of agent- group interaction.

Harnessing In-Group Effects

As workers in our study began to think of the robot as a part of the organization, a desire to protect it emerged. Previous research shows that the influence of a person in the group gets stronger as group members like each other (Fulk, 1993). Having social agents perceived as a group member can encourage the development of norms that are more favorable and generous to a social agent. In our study, repeated, consistent exposure to employees, social interaction around organizational topics, and the robot's persona (not pretending to be more capable than it actually was) seemed to contribute to people's acceptance of it as a member of the organization.

Encouraging and Discouraging Social Comparison

A few participants compared the robot's treatment of them with how the robot treated others. They attributed preferences to the robot, even when it was a result of a purely mechanical decision. People's tendency to anthropomorphize an agent could be used to encourage more frequent interaction with the agent. For example, an agent in a rehabilitation center could publicly encourage a patient who followed its orders well to promote social comparisons. In other cases where such attribution is not desirable, designers should make it clear that the social agent does not have such biases or preferences.

Promoting Socializing

Our field study suggests social agents can be used to promote social activities and even celebrations among people. In our study, the robot's visit created an occasion to socialize. It offered topics of conversation and an excuse to take a break. Engaging in topics that are of interest to a group will be one way for an agent to facilitate socializing.

Starting and Ending Interactions

Research has shown that interaction has a natural opening, middle section, and closing (Goffman, 1963). Limitations in our dialogue design meant that these

rules were often violated as people waited uneasily for the robot to finish its script. To encourage better adaptation to a busy workplace, a social agent must be able to start and end an interaction fluidly at any moment. Social interaction with an agent may be too demanding at a given time, therefore, agents need to offer a graceful way for people turn down an interaction.

Social Context Awareness

For social agents to instigate or encourage group interaction, they need to be aware of the possibilities for social interaction to unfold. To improve this capability, they should be aware of who might be near the focal person or persons, and be able to adaptively deliver personalized messages aimed at the group. Additionally, recognizing who is busy and who is free to socialize or interact will be important.

7.7 Limitations

Conducting a field experiment using a realistic service increased the ecological validity of our results, but also entailed three notable constraints. First, we randomized conditions across the hallways to avoid contamination. Nonetheless, participants in the same hallway sometimes socialized during the Snackbot visit, and the existing culture of the hallway may have influenced the results reported in the paper. Second, we used different styles of personalization to elicit surprise and enjoyment. For this reason, we cannot distinguish among the effects of specific personalization tactics. We do not know whether our strategy would be as effective if only one of the personalization topics were used. Third, the robot took one or two more speaking turns in the Personalization condition than in the No Personalization condition. It could have been more effective simply because it spoke more.

Our study was also limited due to technical constraints. It was conducted on one floor of a computer science building, where the robot could operate reliably, with

access to engineering help if it broke down. (Studying an organization's prototype within that organization is not uncommon for this reason (Venolia et al., 2010).) None of our participants were part of the Snackbot development but some would have had a bias to like robots. Also, our study used a Wizard of Oz technique for the selection of nodes in the dialog script. When we asked participants if they believed the robot was autonomous, they wondered how much the robot was autonomous, but no one believed that they were communicating with an operator through the robot.

The snack service was operated as compensation for participating in the trial for at least two months. Free snacks may have contributed to high service satisfaction in both conditions. Also, we recorded all the interactions with participants' consent. Recording may have influenced participants' behaviors, as well. Finally, the robot was anthropomorphic. Generalizing the results to different service domains and robots will require further investigation.

7.8 Summary

Through a longitudinal study, we provide evidence that personalization with memory reinforces people's rapport, cooperation, and engagement with a robot. The personalization strategies also influenced social dynamics around the users, facilitating the adoption of the robot in the workplace. The results suggest, however, that the efficacy of these strategies may depend on whether people want to interact with the service in a relational vs. functional manner.

8

Understanding Personalization Process of Experienced Personal Service Providers

This chapter explores ways for technology services to personalize service offerings. I interview experienced personal service providers to learn how they personalized their service offerings. The findings show that providers worked with their clients to define goals, tailor the service and style of interaction, and get feedback. I discuss these findings in relation to the current state of personalized technology services and offer design opportunities that will allow such services to help users better define their needs and to adapt their services over time.

8.1 The Practices of Human Service Providers as a Model for Technology-Based Services

Advances in computing technology have resulted in a proliferation of computer-based services, ranging from information services to online health services to social networking services and more. Technology that tracks and models people's preferences and behavior creates a new opportunity to offer more personalized service experiences (Mitchell, 2009). Personalization can enable service delivery to satisfy each person's unique needs and characteristics.

The success of personalized technology services depends on how well the technology understands people. Most current personalized systems use a user-driven or system-driven approach to achieving this understanding. User-driven systems learn about people by explicitly asking them questions or allowing them to personalize settings directly. These systems function most effectively when

people have a clear sense of their goals and preferences. People, however, do not always know their preferences or may not be able to articulate them (Simonson, 2005). For instance, a person may look for advice on dieting, but his hidden goal is to look more attractive. Further, what people say they want may not be an optimal solution for their problems or needs. They may lack domain knowledge, such as sufficient medical expertise, to weigh the value of the advice they seek, or they may have decision biases such as overconfidence (Kahneman, 2003).

System-driven approaches make inferences about people from their behaviors. These systems function best if people's behaviors remain consistent over time and context. However, people's preferences and the contexts for using technology services do vary and change. For example, a person might order a book for his nephew but the context of "gift-giving for children" differs from that of "ordering professional books for myself." Uncertainty surrounding user preferences and contexts, especially over time, poses challenges in creating personalized technology services.

We studied personal service providers as a model for computer-based services to mitigate these issues. Personalization of services existed long before computers were invented. Service roles such as aide-de-camp, court jester, tutor, and butler existed centuries ago for highly placed or wealthy people. Today, personal service providers perform a myriad of personalized services at many levels of society. In this paper, we ask how personal service providers have navigated the problems of determining and adapting to people's preferences and their changes of preference or context over time. Prior work suggests that practitioners use principled and systemic approaches to understanding unique problem situations and creating solutions optimized for each problem (Schon, 1982).

The goal of this research was to explore the strategies and practices of human service providers to inform the design of personalized technology services. We interviewed personal service providers to understand how they determined their

clients' goals and preferences and how they accommodated their clients' changing preferences and experiences over time and context. We use providers' approaches to reflect on current practices of computer-based services, and offer design and research opportunities to help users define their needs and to adapt personalized services over time.

8.2 Method

We interviewed professionals who provided personal services. We defined personal services as activities that provide skillful, individualized care. Personal service is a rapidly expanding sector of the American labor market (Bureau of Labor Statistics, 2012). We studied personal service providers because they had experience and expertise in tailoring their services to meet each client's needs. We sought to interview a wide range of personal service providers, instead of focusing on one type of service to extract common strategies that providers used.

8.2.1 Sample

We interviewed 17 personal service providers (13 women, 4 men) who offered 13 different personal services in Pittsburgh. We recruited using referrals from clients and snowball sampling. We also directly contacted professionals whose business contacts were publicly available. We aimed to recruit providers with experience, as experts are reported to provide better care than novices (Jensen, Gwyer, Shepard, & Hack, 2000). Interviewees' professions ranged widely (Table 19). Providers had practiced for 11.7 years ($SD = 8.7$) on average.

The interviewees offered three types of interpersonal services (Bitner, 1992). The first group performed support tasks in the home or office that clients delegated to them. For example, a personal chef did cooking tasks that her clients could not or did not want to do by themselves. The second group helped clients acquire or transform their possessions or appearances. For example, a professional organizer removed, rearranged, and created spaces for clients' belongings so that their

| <i>Service Provider</i> | <i>Experience</i> | <i>Service Description</i> |
|---|-------------------|--|
| Perform support tasks in home or office | | |
| Personal chef | 7 years | Prepares meals at clients' kitchen once a week to provide a week worth of meal |
| Personal secretary | 7 years | Assists office work by arranging budget, travel, meetings, email |
| Personal secretary | 4 years | Assists office work by arranging budget, travel, meetings, email |
| Caregiver for a quadriplegic | 1 year | Executes daily tasks such as driving, preparing food, bathing, giving medication |
| Caregiver for a quadriplegic | 5 months | Executes daily tasks such as driving, preparing food, bathing, giving medication |
| Acquire or improve possessions or appearance | | |
| Home contractor | 20 years | Remodels clients' houses |
| Real estate agent | 5 years | Assists clients purchasing a house |
| Professional organizer | 15 years | Assists clients organizing their houses or offices |
| Hair stylist | 15 years | Provides hair cuts and hair color styling |
| Improve physical or mental wellbeing | | |
| Massage therapist | 20 years | Gives therapeutic and relaxation massages |
| Chiropractor | 25 years | Administers treatments for physical ailments |
| Personal trainer | 5 years | Creates exercise and diet programs and coaches work-out sessions |
| Personal trainer | 2 years | Creates exercise programs and coaches work-out sessions |
| Physical therapist | 20 years | Creates a rehabilitation program and coaches clients in their residence |
| Physical therapist | 28 years | Creates a rehabilitation program and coaches outpatient clients |
| Counselor | 15 years | Counsels clients in recovery from alcohol and drug addiction |
| Personal tutor | 10 years | Tutors students in chemistry and mathematics for college admission tests |

Table 19. Interviewee description.

home would be better organized and easier to navigate. The third group helped clients improve their wellbeing. For example, a personal trainer created personalized exercise programs and coached workout sessions for his clients. This kind of service usually involved close proximity and physical contact between the provider and the client.

All of our interviewees had steady clients that they interacted with over time. Some services, such as physical therapist and personal tutor, required long-term interactions with clients to obtain desirable service outcomes. Others had voluntary repeat clients (e.g., hair stylist) or were based on a long-term contract (e.g., personal chef).

8.2.2 Interview Protocol

The interview began with questions about the professionals' work context and service. Then we asked about their typical relationships with their clients, and whether any of their work was tailored to different individuals. Further questions probed what information service providers needed to do a good job with their clients, how they listened to feedback and observed cues in their clients' behavior, and whether and how their interactions with repeat clients changed over time. We also asked about breakdowns or negative experiences of interactions with their clients. The interviews took one hour; all interviewees consented to confidential audio recording of the interview.

8.2.3 Analysis

The interview recordings were transcribed for analysis. The analysis involved two parallel processes. First, we extracted personalization instances using three stages in the process of personalizing services (Tuzhilin, 2009) as a coding category: (1) how service providers collected information about clients, (2) how they tailored their service, and (3) how they evaluated their service and used feedback. This analysis allowed us to systematically analyze what was personalized. While we

extracted personalization instances, we added codes that described how and why providers personalized their service. We grouped these lower-level codes into thematic clusters that revealed themes of personalization across the three stages of the process. As we continued to develop categories, we also compared what we were learning with existing theories and frameworks of personalization, and used the themes to reflect on current practices of technology personalization.

8.3 Findings

All the service providers that we interviewed referred to their customers as “clients.” In doing so, providers assumed a personal service role rather than an economic goal of selling goods or services (Remen, 1991). During the interviews, all providers were articulate about their detailed planning and the amount of thought they put into personalizing services for each of their clients. All providers emphasized that they were not offering preferential treatment to certain clients by personalizing their services; they personalized their service to create service solutions that worked best for each client, while ensuring an equal level of effectiveness that would not be achievable by offering the same service for all different clients.

8.3.1 Collecting Client Information

Providers sought to understand their clients’ goals, motivations, needs, expectations, constraints, preferences and personalities in order to personalize the services that they provided and their interaction style. To collect this information, providers asked their clients questions and observed them, making inferences about their behavior. Using both methods allowed them to understand clients’ subjective opinions and desires, and what clients did not realize or did not express but was still critical to the service and its delivery.

Building Rapport with Clients

All providers described the importance of building rapport with clients in their initial meeting as explained in (Price & Arnould, 1999). Clients needed to share personal information with providers. Providers worked closely with them, sometimes in their homes. To facilitate this process, providers sought to make clients feel at ease and to establish trust with them. Providers used small talk to establish rapport with their clients, showing a personal interest by asking clients what they did for a living or talking about common interests.

All providers also emphasized the importance of expressing genuine care toward their clients. They said they showed interest by maintaining eye contact and listening to their clients attentively. Over time, providers remembered details about clients and paid them individualized attention, remembering and mentioning mutual experiences, events, or concerns that clients had discussed with them. Doing so helped providers and clients maintain good rapport over repeated service encounters.

Asking Questions about Clients

As rapport was established, providers asked clients questions to gather information that they needed to create personalized service offerings. All professionals inquired about clients' demographic information (e.g., age, jobs), goals and constraints (e.g., medical conditions, budget, likes and dislikes) and current ways of dealing with problems. Some providers also used a checklist or a form to make sure they collected all critical information.

Providers did not assume that solutions or services that clients requested were a good solution for them. They asked clients what they expected to get out of the service, and gathered information that allowed them to assess whether what clients wanted was an optimal solution for them. All providers mentioned that they asked questions such as “What works and does not work with your current [service]?” or “What do you like or not like about your current [possessions]?”

These questions allowed clients to express their opinions grounded in concrete examples. By establishing common ground, providers could understand what clients' problems were and where their real difficulties or needs arose.

Asking Probe Questions

All providers used a prototypical list of questions in their initial meetings with clients. Clients' answers served the role of helping service providers to infer clients' personalities. For example, if clients gave a long answer to a question such as "how are you doing," it could be an indicator of how talkative a client was and if he or she desired a relational interaction with the service provider.

During his first visit to a client's home, the contractor we interviewed asked clients, "May I place my [heavy, old] briefcase on your [dining room] table?" Based on his client's answer, the contractor inferred whether the client was treating him as a professional or as a servant. He also inferred their expectations about cleanliness, and their boundaries between private versus semi-private places. Using this information, the contractor tailored how he interacted with clients and how he behaved in the house. His question was a natural part of his work process. He needed to place his briefcase somewhere during his visit, so clients never guessed that he was testing them.

Observing Clients and their Possessions

While providers conversed with their clients, they also paid attention to clients' interaction styles, behaviors, appearance, possessions and the context where the service would take place. For example, most providers explained that they paid attention to how talkative their clients were in order to infer whether their clients were extroverts or introverts. Then, they tailored the way that they built rapport with the clients, and gave feedback or instructions. For example, a physical therapist explained that she would use a more directive sentence with outgoing, talkative clients ("Now hop on the bike"), whereas she would use a suggestive instruction ("How about doing biking now?") with a reserved client. She

explained that from her experience, she learned that using direct instructions with introverted clients would make them uncomfortable, whereas using suggestive instructions with outgoing clients could invite further conversation that was unrelated to therapy.

Some providers observed their clients' appearance and possessions in order to infer their preferences about service outcomes. For example, a hairstylist observed her clients' outfits and accessories to understand their idea of beauty. She used this information to decide whether to recommend a hairstyle that was high versus low maintenance, trendy versus more conservative, and warm versus cool color. The professional organizer observed whether her clients' homes were decorated with a focus on visual or textural materials. The contractor observed how his client's home had been previously maintained or improved, looking at details like hardware on the windows to understand whether his client prioritized precision or aesthetics. This distinction was important to him because he might have to sacrifice one or the other to finish the work on time and within budget.

8.3.2 Co-Creating Personalized Services

Building on the information that providers collected through conversation, observation, and probing questions, providers and clients collectively discussed the scope of the service and the client's goals. This process is known as service co-creation (Vargo, Maglio, & Akaka, 2008). Our interview results revealed three ways that personalized services were co-created depending on their clients' knowledge of service domains, and their clarity of goals and preferences: co-creation where the client led, co-creation where the service provider led, and equally led co-creation.

Client-Led Co-Creation

When clients had clear preferences and goals for the service, service providers' input was minimal. In our study, service providers such as the personal chef, the caregiver for the quadriplegic, and the personal secretary let their clients take the

lead in personalizing the service. On these cases, clients were very specific about their service preferences. Thus, in their initial meeting, providers focused on acquiring information about clients' goals and preferences with as much detail as possible. For example, the caregiver for the quadriplegic described how each of her clients wanted their activities of daily living such as bathing and dressing executed differently. The personal secretary needed to learn her clients' travel preferences and apply them each time she arranged a trip. These fine-grained differences had to be learned to offer personalized service.

When clients led the co-creation process, providers refined their personalized services iteratively over time. Clients periodically requested changes and modifications. For example, the personal chef explained that once she collected information about her clients, she tailored her menu based on this feedback, and provided a week's worth of meals. Then, her clients provided feedback on the menus and requested modifications such as less salt or larger portions. One of her clients wanted her salad prepared in a very specific way, with the cucumber peeled and cubed instead of sliced.

Provider-Led Co-Creation

When clients lacked knowledge and expertise about the service domain, providers played active roles in personalizing the service. In our study, providers who offered expert services to improve clients' physical and mental well being, such as personal tutor, personal trainer, and physical therapist, exhibited this pattern of personalization. Providers had knowledge and expertise in creating services that could help clients reach their goals, whereas clients typically lacked this knowledge. In their initial meeting with clients, providers sought to understand clients' goals and motivations, but clients' initial preferences played a comparatively little role in determining the nature of the personalized service. For example, a physical therapist created a simple storefront mock-up for her client who wanted to return to her cashier job in a bakery after her injury. Together, they practiced sales interactions and monitored the client's progress in managing

the cognitive and physical demands of the job. Although the client had provided information about her goals and motivations, it was the therapist's role to devise a rehabilitation program with specific tasks.

The delivery of these services mandated clients' presence and active participation. This requirement differs from the situation described in the previous section, in which clients specified the service and delegated tasks. One of the goals of personalization in provider-led services was eliciting and sustaining participation from clients. Providers explained that understanding clients' goals was important for maintaining engagement over time. Sometimes clients, however, sometimes did not think about or initially share their deeper motivations with service providers. In these cases, providers guided clients to think more deeply about their goals and motivations, so that they could personalize their interactions and sustain their motivation over time.

A personal trainer said he tried to understand the deeper motivations of his clients, even when they initially appeared to have well-defined goals:

(my clients) may say that their goal is to lose ten pounds, or lower their BMI, blood pressure or to be more fit or healthy; these are very generic goals, and are not what truly motivates them. For example. . . they have known that they may die early if they do not change their lifestyles, but this did not prompt or motivate them to change and come here. There is always something more emotional and motivating for different individuals, even though it may not be related to their biggest problem (threats to their health). So I keep asking 'why' to truly understand that motivation. At the end, they say it is to fit into their jeans that they used to wear when they were young. . . or look better to their husband, or get over the break-up. They want to fulfill their emptiness.

What is noteworthy here is that the service provider's role is to repeatedly ask questions to get at hidden motivations. Providers used this information to tailor service offerings and interaction styles. For example, the personal trainer used this information to personalize his exercise programs to support the goals that clients cared most about. He also tailored his feedback to his clients. A personal

tutor used his understanding of a student's goal of going to pharmacy school in tailoring examples that he used to explain math or chemistry concepts. These shared experiences made the client-provider relationship more collaborative and caring.

Although providers initially set out a personalized service plan for each client, they adjusted each session and the overall plan depending on clients' progress. For example, the physical therapists paid attention to objective measures (strength, balance) that showed whether patients were making progress in improving their physical condition. The chiropractor used a paper form on which patients could specify what areas of the body were in pain and how bad the pain was. Based on changes in their pain ratings, she gauged how frequently her clients needed to visit her practice. The personal trainer asked his clients to fill out a journal and follow the personalized diet and exercise program at home. He adjusted the programs depending on his clients' progress:

If they're trying and it's just not working, then I'll just say, "You know what? We'll try a different route, because that's my fault." [...] So maybe we just give them one exercise to do at home and one nutrition habit. So now they're succeeding. After the first month they've done a fantastic job. They have done one stretch in reality, it only took them like two minutes to do a day, but they had success and so then we could keep building on success.

Personalized feedback played an important role in keeping clients engaged. Service providers emphasized the importance of talking about concrete measures of improvement, because this conveyed authentic care and encouragement. For example, one physical therapist gave feedback to her clients by saying "You did five steps last time; now you are doing ten steps. Good job!" rather than saying "You are improving. Good job." A personal trainer sent out personalized messages to his clients, wrapping up the previous session and setting expectations for future sessions. He sent these reminders twice a week, "Hey, great workout yesterday. I look forward to Tuesday." During the weekend, he would remind

clients, “Hey, we had a great week. Have a great weekend and get ready for the next week.”

Equally Shared Co-Creation

When clients knew their preferences, but lacked knowledge about the service domain, the creation of the service was more or less equally led by providers and clients. Service providers who helped clients acquire or improve their appearance or personal possessions such as hair stylist, contractor, professional organizer, and real estate agent, exhibited this pattern of personalization. These providers would describe available service options and offer their recommendations in order to help clients make the best decisions about their personalized service.

A hair stylist choosing a hairstyle for a client that satisfied both her and her client’s idea of beauty illustrates this point. She explained that many of her clients visited her salon with a desire for a new hairstyle, but only a vague idea of what they wanted. Some clients brought pictures that showed desirable hairstyles. In some cases, their preferences were influenced by the models’ faces rather than the hairstyles themselves, and the chosen hairstyles would not complement the client’s face. The provider’s goal in this situation was to create a style personalized to the client that satisfied both of their preferences. To do this, the hair stylist sought to understand her client’s idea of beauty and interest in beauty by observing their clothes and jewelry, and discussing what clients wanted to do with their hair. She would then suggest a style that complemented the client best. If the client did not like her suggestion, or felt it was too drastic, they would negotiate a satisfactory middle ground.

A professional organizer took a different approach to helping clients concretize their preferences. Rather than showing examples or providing recommendations, she asked a series of questions and rephrased the answers so that clients could narrow down what they wanted. One anecdote illustrates this point.

I start the first session by asking, 'Take me through your space you want to work on. Tell me what works and what doesn't work about the space.' The client may start by saying, 'This room is a total waste; this is my home office.' I may ask, 'What do you do in your home office?' The client may say, 'I pay my bills.' I may say, 'It sounds like it's working for paying bills.' 'I guess so.' I may ask, 'Why do you think it is not working as an office?' The client may think for a little bit and answer, 'I guess I originally wanted to pay bills and make art.' 'Ah, so you want a studio and an office. It is working as an office but not as a studio. So you need a studio.'

8.3.3 Evaluating and Adapting Personalized Services

Providers in our study did not assume that their clients' goals and preferences were static. They used multiple feedback channels to constantly monitor their clients and re-evaluate their services in relation to the goals and scope they had created together. Based on this feedback, providers drew from their repertoire of strategies until they found effective ones. Even when providers settled on successful strategies, they continued to personalize based on situation and context.

Monitoring Clients' Experience and Context

Providers checked whether their personalization strategy was working as intended in various ways. One method they used was to observe clients' non-verbal behavior and reactions to their service in real time. For example, one personal trainer said he adjusted the weights according to his client's posture. The massage therapist monitored the lines of expression around her client's eyes to monitor their comfort level.

Another way to monitor the service was through conversation and small talk. Providers asked clients directly what they thought or felt about the service. Questioning could be done periodically, or at the end of a service encounter. For example, the professional organizer would always ask "what excited you most?" at the end of each session with her clients. Providers also periodically followed up by sending email or calling to ask whether their clients were still satisfied with the service. If providers had administered an intake survey, they would administer

the form again to check whether assumed conditions for service still remained the same.

Clients also proactively communicated their opinions about the service outcome. For example, the personal secretary's client told her that she had done a good job managing email and travel. At that time, they also communicated about changes in the client's goals and preferences. Sometimes, changes in clients' personal circumstances required changes in personalized service. For example, a personal chef talked about how her client wanted to increase the protein content of their meals when their son started playing football. Other times, the service evoked positive change in client outcomes, and subsequently, client goals had to be re-examined. The physical therapists and caregivers described how their clients often saw improvement in their physical condition, leading to a need to change the service. For example, a quadriplegic client identified new goals of going to college and entering the work force. A physical therapist's client initial goal was to become pain free. As she progressed in her program, she discovered a new joy in exercise, and set a goal of participating in a 5K marathon. Small talk played an important role in this resetting of goals. It created opportunities for providers and clients to communicate these changes naturally and spontaneously.

Another way to monitor service outcomes was specific measurement. For example, the physical therapist and personal tutor administered tests to quantitatively measure improvement. Even without tests, over time, service providers could distinguish routine and non-routine responses from their clients. These could be used as a further opportunity to monitor their service and receive feedback. When clients exhibited atypical behavior, especially behavior that seemed to reflect a negative attitude toward the service or a change in personal circumstance, providers followed up to find out whether there was new information they needed to know. For example, the personal trainer described how clients who are usually talkative are quiet, he asks if everything is okay with

his clients, to find out whether there was a problem with the training session or if a client simply had a bad day. The personal chef became concerned if she saw her meals piling up in her clients' fridge, and she would proactively follow up. The counselor, chiropractor, and physical therapist described how they would ask if everything was ok if clients began to pay less attention to their clothes or their hygiene, to check for possible depression.

Using these feedback channels, service providers assessed whether their interaction and service offerings had the intended effect on their clients, and whether clients changed in response.

Experimenting with Personalization Strategies

If a personalization strategy was not effective, providers experimented with other personalization strategies. To do so, they relied on a set of heuristics that they had developed over time. They could then experiment with other strategies until the personalization was working well. For example, one physical therapist described how she tried different rapport building strategies until she discovered what worked best with her clients. She might first tell a joke; if this was unsuccessful, she would move on to asking about the client's job or profession. If small talk was unsuccessful, she would try working with two clients at once so that they could socialize with each other. Sometimes this strategy worked wonders:

I had one gentleman who was in a boating accident and he had a compression fracture in his spine. So he had to wear a brace. And he was miserable. He was angry. Poor him. He was just absolutely cranky, miserable. While he was in the clinic, a woman came in who had been paralyzed from the waist down. Her boyfriend had shot her. And that man, that day he saw that, made sure that that woman had everything she needed. [...] She had flowers in her room. He had balloons sent to her room. He had animals. He had every piece of equipment she needed sent to that room, and all of a sudden his problem wasn't so big. And it made a huge difference. So getting people to connect like that helps them and helps me.

This example suggests that the ability to try different personalization strategies to find one that works is just as important as getting the personalization perfect in

the first place. Many providers described an iterative process of finding a strategy that felt right. For example, the personal tutor tried several ways to explain abstract mathematical concepts. The professional organizer tried different strategies to get clients to articulate their organization goals. Service providers who focused on health and wellness tried different intervention strategies to find the best way to co-create the service.

The Role of Prediction in Personalized Service

Experienced service providers developed a general schema to evaluate and predict the outcomes of their clients. While realizing that every client's experience was different, they paid attention to cues that could signal that client experience was changing in an expected way. For example, a personal tutor knew students' progress would slow down after one month of tutoring. When iterating on or personalizing the service over time, providers took these typical changes in user experiences into account, and proactively changed their service. Providers expected people's mood and emotions to change with the context. The real estate agent knew anxiety often overwhelmed clients who were searching for a home, repeatedly visiting houses and not finding a desirable one. Anxiety also affected clients when they had to make a purchase decision. One of her roles was to provide assuring feedback at these times. The personal trainers knew that his clients would be bored after repeating the same program for a few times, so he proactively changed their programs periodically:

So I purchased [...] gloves and mitts so she could box. So now, again, it adds to the whole motivation. I've been showing them a new exercise that's only for them. I went out and bought the equipment for her and it's specifically for her. Her name's on it and it was a gift because I knew she couldn't do this style [of] training. But there's always other options. So it's keeping them interested.

A few providers such as the personal tutor or the professional organizer emphasized the importance of communicating negative experiences that might arise over time with their clients. For example, the professional organizer

forewarned her clients about potential emotional distress from discarding their possessions in their initial meeting:

“Some people hit it [the wall] the first day. Some people hit it the 85th day. When you hit the wall, I’m going to point it out to you... It doesn’t mean that we have to stop [organizing the house] and it also doesn’t mean you have to do anything horrible. It just means that I have to help you through whatever that is.”

8.4 Design Opportunities

Our interview findings reveal rich and nuanced ways that service providers personalized their services for each client. Providers and clients collaboratively defined why clients wanted to use the service, and worked together to find personalized solutions to satisfy client goals and preferences. Providers also personalized how they interacted with their clients and expressed care to facilitate this co-creation process. Through repeated interaction, providers and clients communicated whether there were any changes in clients’ motivation and goals in order to adjust their personalization strategies. We use these findings to inform the design of computer-based services that help users define their needs and adapt personalized services over time.

8.4.1 Defining the Goals and Scope of Personalization

Providers and clients co-created the goals and scope of personalization. Providers first sought to understand *why* clients wanted the service, rather than *what* service they wanted. Asking clients about their reasons for using the service, providers understood the clients’ needs and underlying motivation. This allowed providers to verify whether the goals and scope of service that the clients wanted satisfied their needs. At times, clients clearly knew their preferences and goals and took the lead in refining their preferences over time. Nevertheless, all providers began personalizing asking for their reasons and motivation. If the clients were

unsure how to address their needs, they worked with the clients to define the goals and scope of personalization.

This co-creation process is in stark contrast to the personalization process of most computer-based services. Many of these services, such as e-commerce and news websites, allow users to personalize the content and structure of the service, assuming that users have a deep understanding on why they wanted the service, have well-defined goals that satisfy their needs and know how to achieve their goals through the service. Very few systems start the personalization process by working with people to verify or confirm their needs and goals, or check if personalization could provide better ways of reaching those goals. This is an interesting challenge for future design and research of personalized systems. How can these systems assist people in articulating their goals that satisfy their needs and refining the scope of the service? Providers' strategies revealed in the study suggest three potential approaches.

Promoting Reflection on Why People Use Services

One approach to a more collaborative personalization system might be to use questions that prompt people to reflect on their reasons for using the service. Providers in our studies used questions to prompt clients to think about why they wanted to use the service. In doing so, providers helped clients define their own goals. Asking a why question is a very simple activity, but providers' examples show that it can be effective in guiding people to think more deeply about their reasons for using a service. Providers who provided a service that improved clients' well being said that having clients reveal their underlying goals was important in sustaining their engagement. An interactive education or health management site could incorporate such a component into the system. For example, activity-tracking software could prompt a question asking people to reflect on the reasons *why* they exercise, helping them realize their underlying motivation, such as to be energetic to achieve more things in their lives, before they set a concrete and specific goal for exercise, such as burning 500 kcal per

week. The software could use this information to tailor its feedback and show how different service options could help people achieve their underlying goals. Expert crowd-sourced applications could help create the links between personal reasons and motivations with appropriate solutions and feedback (Aleahmad, Alevan, & Kraut, 2009).

Prompting people to reflect on their reasons for using a service can be useful for services where people might benefit from orienting themselves to long-term, deeper goals. Previous research in decision making shows that asking people the reasons for doing an activity leads them to focus on higher-level goals, long-term consequences of their choices, and decisions that are more in line with their ideal selves (Carver, & Scheier, 2000; Fujita, & Han, 2009; Trope & Liberman, 2010). For example, online news websites can present questions to users, right before setting the personalization features, which ask them to list reasons why they read news. People may set their personalization settings to be more in line with their ideal selves (e.g., read world and technology news), rather than only what they usually do (e.g., read celebrity gossip). Personalization of news content can play a big role in influencing people's exposure to a variety of news (Sunstein, 2009), and initial settings of more self-beneficial choices might encourage this diverse exposure.

Translating Misconceived Goals to Optimal Ones

A second approach would be to design systems that can help people diagnose their current problems and needs in order to help them to find service solutions. All the providers we interviewed asked, "what works and does not work with your current service?" This question became an important anchor point for their assessment of clients' needs and problems. Computer-based services can also use this technique to understand users' current problems and suggest optimal solutions. An interesting example is suggested by a new online bra shopping site (True & Co.). This site offers staged questions that customers answer about how well their current bras fit. Based on this assessment, the service suggests a bra size

and style for customers and displays appropriate options. What is unique about this website is that it implements an algorithm that links common misconceptions and problems with optimal solutions. This approach could be used for other services to help people devise more optimal, personalized solutions.

Offering Goal-Centered Decision Guidelines

When clients knew their goals, but had vague ideas about how to reach them, providers helped clients concretize their ideas by asking simple questions that guided them to think about important decision criteria. Many computer-based services offer very detailed search options or keywords to help users find and customize the options they like most. But these sites often lack structured assistance to help people make good decisions. For example, an online shoe shopping site allows people personalize their search by color, style, or size, but they do not show a check list of factors that they need to consider buying a right shoe such as the right fit or tradeoffs between design and durability.

Incorporating decision criteria for articulating the goals and scope of a service may help people become more aware of their decisions. For example, a web plug-in or an app could show checklists of criteria related to people's goals. People could then check off criteria as they go through different options. Automated analyses of service reviews could be used to automatically create these guidelines.

Understanding Users through Implicit Probes

Asking questions to users about their goals and preferences can be obtrusive for some services. A strategy that the contractor used could be adopted by a computer-based service to implicitly understand user preferences. The contractor asked a question (whether he could place his briefcase on a living room table) that set up a situation that required people to make a decision. The answer to this question revealed important individual differences that he used to personalize his service. Similarly, computer-based systems could start their service asking users "how are you doing?" and provide the user with options to answer this greeting. Users' answers to this question could reveal whether they want to have social

interaction with computer-based services (Lee, Kiesler, & Forlizzi, 2010) or whether users are extroverts or introverts.

Balancing Autonomy of Users and Systems

People vary widely in their level of self-understanding and desire for personalized services; this understanding varies across domains and within a service. For example, some people may want personalized service for choosing food or exercise, but not for improving their relationships. Personalized technology systems should assess and respect these differences and flexibly provide different levels of assistance and guidance. For example, design features that prompt more reflective behaviors should respect the right boundary between being too easy to ignore and being too intrusive for users.

8.4.2 Adapting Personalization over Time and Context

Service providers in our study used multiple feedback channels to constantly evaluate their personalized services in relation to their clients' needs over time. This allowed providers to detect changes in users and the context where services were co-created. Except for content recommendation systems that allow people to rate recommended offerings (e.g., Pandora), very few personalized systems evaluate their personalized services in-situ, or prompt people to readjust or reevaluate their personalization settings when their goals and context change (Adomavicius & Tuzhilin, 2005; Tuzhilin, 2009).

Building Feedback Channels and Personalization Repertoire

In our study, providers and clients shared the task of tracking changes in their goals or context. Providers followed up when they saw atypical behaviors of a client, such as a client being unusually quiet, or changes in context such as food accumulating in the refrigerator. When clients had new goals or changes in circumstances, they communicated about these changes with providers, so that both could collectively adjust service plans.

Personalized computer-based services could use this communication pattern to sustain feedback loops over time. To do so, it would be important to add system capabilities to detect non-routine events and proactively follow up with clients. These systems could also provide incentives for people to communicate changes in their goals and/or context. For example, an activity tracking application could prompt a question when users' activities decrease to find out whether they had been traveling or whether the exercise program was too demanding or hard.

With a repertoire of personalization strategies, service providers could flexibly experiment with different strategies that would work with changing situations. One of the tradeoffs often discussed in the personalization literature is the risk of providing an inaccurate personalization (Ashman, Brailsford, & Brusilovsky, 2009). Using a feedback loop and multiple strategies, personalized systems might recover from a suboptimal personalization strategy. The system would need a repertoire of strategies for a particular domain. For example, an online education site could model how an experienced human tutor provided feedback to students and how the tutor recovered from a less effective one to find an alternative.

Using a Temporal Model of User Experiences

Providers in our study had a general understanding of temporal changes they expected in clients' experience with their service, but their focus changed over time from determining initial goals and service to adapting it to new goals or changes in the client. They both focused on clients' emotional and behavioral experience over time. Personalized technology services could also benefit from developing a temporal model of people's experience in order to guide how they should personalize over time. Doing so would allow systems to personalize their services based partly on prediction along with past behavior and explicit inputs. The trend to analyze huge datasets and to measure client experience could eventually lead to a temporal model of service experience.

8.5 Design Space for Co-Creation of Personalized Technology-Based Services

In this section, we explore ways to design co-created, personalized technology-based services, and discuss which design dimensions should be considered in this process. The first step of personalization is configuration of the services; improper configuration can create a cascading effect in the later phases of a service experience. Thus, this section focuses on the initial configuration process. This process consists of (1) information collection for personalization and (2) configuration of service features and goals. We explain how this process is currently done and what design strategies could be used to mitigate the shortcomings of the existing methods.

8.5.1 Current Personalization Methods for Technology-Based Services

Current personalization methods could be categorized by the varying degrees of system and user input/labor required to make decisions (Parasuraman & Sheridan, 2000) (Table 20).

| <i>Information Collection (Input)</i> | <i>Creation / Configuration of Services (Output)</i> |
|--|--|
| (Information in users' heads without being externalized) | 1. Users personalize service features on their own |
| Information entered by user | 2. System suggests personalized services and users can modify them |
| | 3. System suggests options and users choose one |
| Information implicitly collected by system | 4. System suggests personalized services |
| | No personalization (Same service for everyone) |

Table 20. Current methods of personalization listed along the spectrum of user-driven and system-driven personalization.

In user-driven personalization, users personalize service features without externalizing their thought processes, going over needs, requirements, and preferences internally. In system-driven personalization, systems either ask people questions in order to understand them, or infer their characteristics based on implicitly observed behaviors. The system then uses the information it has gathered to offer personalized suggestions or services.

We postulate that when users are given more opportunities to participate in the process, they enjoy fuller freedom and flexibility in personalizing service features. Thus their perceived ownership of and motivation to use the service increases. The downside is that users need to put time and effort into personalizing the service. In addition, when users do not have enough knowledge or information about the service domains, their personalization choices may not reflect the best solution for them. Increasing system involvement decreases the amount of effort users need to put into personalizing services, which is more convenient for users. When systems have expertise that users do not have, system-driven personalization could also lead to more effective solutions for users. However, if user participation is minimal, users may feel less ownership over the systems and thus be less motivated to use them. How can we utilize system expertise on personalizing solutions while involving users in the personalization process, empowering and motivating them to use the service? The following sections describe different ways in which users and systems can co-create personalized services.

8.5.2 Empowering Users When Collecting Information for Personalization

The process of collecting information for personalization could be enhanced in two ways.

Helping Users Realize and Focus on Their Underlying Motivations

In current personalized services, systems generally use the information they collect about users to classify users into pre-defined categories. For example, health services such as Fitbit and Highmark wellbeing management services ask people to answer questions aimed at understanding users' current dietary and exercise lifestyles, then use this information to suggest certain exercise or diet regimes. However, these questions usually only ask about people's behaviors or opinions on their lifestyles; they do not address tendencies to stick to their surface goals. One way that this personalization process could be enhanced is for systems to ask reflective questions in the beginning of the personalization process, as some experienced personal service providers do. These questions might nudge people to think more deeply about important aspects of services.

Conceptualizing Users as People Rather than Numbers

To collect data from users, most current personalization features in computational applications ask questions that can be readily processed by computational systems. Most of these questions collect information in objective data forms such as numbers, scales, or categories. This process implies that systems are conceptualizing users as collections of numbers and facts. Depending on the service type and domain, however, people may desire to be treated and understood as a person, rather than one point on a data plot. One way to make users feel that they are being treated as a person and that the service is truly personalized for them is to ask users subjective questions about themselves. Asking subjective questions, whether systems actually use the answers for personalization or not, may make users feel that they are regarded as a person rather than as a data point on a graph.

8.5.3 Empowering Users When Creating Personalized Services

Informing Users' Personalization Choices

One of the downsides of user-driven personalization is that people may not have the right information or level of expertise to create services that satisfy their own needs. Many system-personalized services currently offer convenient solutions (such as personalized news). In domains such as healthcare and education, people may be more invested in creating personalized services (as compared to news). One way to give users control over personalized systems is to provide them with the information that systems use to personalize – either raw data or logic-level information – give them the opportunity to personalize on their own. We postulate that providing users with helpful information (instead of prescribing a solution) will be just as effective as providing system recommendations; and people may be more receptive to and more motivated to use the service than they are when receiving system recommendations. In addition, allowing users to access the information and make their own choices may have a long-term effect on their abilities to personalize services, helping them obtain basic, beneficial knowledge about using services (on their own) later.

Giving Users an Opportunity to Personalize on Their Own

Along the spectrum of personalization methods, from user-driven to system-driven, systems can use scaffolding ways to decide what kinds of personalization methods to use for a particular individual. For example, systems can show recommendations first, and then give users the option to modify the recommendations; or systems can allow users to customize options first, and then present them with a link or button to access personalized recommendations. Depending on the type of service, this seemingly small design choice could be used to strategically influence how people personalize service. Viewing system recommendations will strongly influence users' opinions on service solutions by creating “anchoring effects” (Tversky & Kahneman, 1974). When users do not have knowledge about or experience with the service, they are often heavily

influenced by the provided examples and rarely change the defaults. In certain domains, examples can help, but in other domains, showing the recommendations first could destroy an opportunity for users to create their own services. For these domains, systems should ask users to personalize services first, then provide feedback or show comparable examples.

8.5.4 Psychological Ways of Increasing Perceived Control

Timing of Participation

When people's explicit participation is called for in the personalization process could also influence people's perceptions of control over the personalization process. For example, previous research suggests that people tend to remember how the experience ended better than how the experience in the middle unfolded (Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993). Thus, when the participation is requested at the final stage (such as to confirm personalized service features) rather than in the middle of the experience, people may be more likely to remember having control over the process.

Multiple Choice vs. Free Response

The method of asking for participation from users may also influence people's perceptions of the service. Two common ways of collecting user information are giving people lists of choices or giving them direct control over the values of different features. People may feel greater control and feel that services are personalized to a greater degree when they are, at least once, asked to adjust some values directly, rather than choosing from suggested options every time.

8.6 Limitation and Future Work

A limitation of this study is that we interviewed a small sample of service providers with unknown selection bias. Interviewees who participated in the study might have felt more confident in their practices than other service providers. We did not interview clients. In our future work, we will investigate

clients along with their service providers, and empirically test the efficacy of the design strategies. We also acknowledge that another way to implement a personalized system is to connect experienced human service providers with users. We chose to elaborate on improving computer-based services, as these services have the potential to provide scalable services to the masses with affordable costs. Investigating computer-systems that can help providers do a better job in personalizing their service could be an important future research area.

As with any technology, there will be risks in personalizing computer systems further. Personalization can threaten privacy and security if the system is insecure or allows third parties to use people's data. Some people may not want this level of information available to a system whereas others may share their personal information with a system as long as it benefits them. We need to learn more about how people feel about systems that learn about them over time.

Conclusion

This thesis contributes new principles and knowledge/systems for creating personalized technology-based services, and in the process, contributes to service design research.

9.1 Review of Thesis Contributions

This thesis offers several contributions in the fields of HCI, HRI, interaction design and service research.

9.1.1 Contributions to Personalization, HCI and HRI Research

a. This thesis broadens the concept of personalization discussed in HCI and HRI. This thesis argues that designers of personalized systems should consider situations where people do not clearly know what they want or have preferences that are sub-optimal for their needs or change over time. I present a set of design considerations for future personalization technology, based on observing the practices of experienced personal service providers.

b. This thesis draws attention to the importance of recognizing people's orientations to interactive technology – whether they treat the robot as a relational being or a utilitarian tool. The social actor paradigm highlights people's tendency to interact with interactive systems as they do with other people. Literature on social agents and robots focuses on using social strategies with people to build rapport. This thesis investigated two robotic systems deployed in the real world for an extended period of time. The results suggest that not everyone treats interactive systems as social

actors even when the systems are anthropomorphic. The same individual can have different orientations across different contexts. Future interactive systems need to be sensitive to this difference across individuals and one individual's changes in orientation in different contexts.

c. This thesis designs and evaluates two types of adaptive interaction for technology systems: a social interaction strategy that relies on knowledge common to an organization, and a personalized interaction strategy that learns about people over time and adapts interaction accordingly. These strategies have been evaluated in field studies, and they appear to be effective in improving engagement, rapport, and cooperation.

d. This thesis contributes to the research on understanding rapport and supporting long-term interaction between technology systems and people. Previous studies used social strategies such as small talk, demonstrating empathy, and mimicking gestures, but they did not investigate how these strategies might adapt over time. The result of the longitudinal field study shows that personalization strategies, when combined with social strategies, are more effective than social strategies alone.

e. This work contributes to understanding how people interact with technology over time. Most studies of screen-based and robotic social agents have been conducted in isolated settings. This thesis, however, shows that interpersonal interactions in the workplace influence the social dynamics that unfold around technology. Future research should take into account the social and organizational contexts of technology.

9.1.2 Contributions to Interaction Design Research and Service Research

- a. I documented the process of designing and implementing a social robot and snack delivery service, which is one of the first research of a human-centered design process applied to a social robot and its service.
- b. Most HRI research centers around the robot – its physical form and interaction, yet less attention has been given to the design of the service it provides. I illustrated how to employ a service design approach when designing both a product and its service, taking into account multiple stakeholders in the system.

9.2 Ethical Discussion

The personal and social implications of the technology services studied as part of this dissertation, and of the broader category of autonomous, intelligent technology services in general, deserve careful consideration.

9.2.1 Social Technology

People attribute agency to technology-based services consciously or unconsciously, regardless of whether designers intended to evoke such effects. As technology becomes more intelligent, researchers have begun to investigate ways to make technology-based services more socially beneficial. For example, the study in **Chapter 7** suggests that robotic service providers may be able to build rapport with people and promote social interaction among users. This study adds to the growing evidence that social technology can promote social or psychological wellbeing. In addition, the right (minimum) level of social interaction could facilitate daily service interactions (the way that social, friendly strangers or service employees can cheer people up.). Studies in **Chapters 4** and **5** suggest that adding a social component to technological interactions may also

make service experiences more enjoyable (as Apple's Siri does with its sometimes whimsical social behaviors), or help users tolerate service breakdowns without losing their sense of control.

How can we weigh the upsides and downsides of this type of social (or sociable) technology? Some researchers have cautioned that this kind of social technology could isolate people from other human beings (Turkle, 2012). People may prefer to interact with technology rather than with human beings because social technology can provide comfort and benefits without the baggage that can come with human social relationships, such as reciprocity or fear of rejection.

At this point, the extent to which people's relationships with social technology could develop is unclear. As suggested in **Chapter 7** and in studies with a robotic dog AIBO, people know that technology is mechanical, not human; although they may apply some social expectations and norms to machines, they do not apply all of those that they would with other humans. It is probable that a new kind of relationship type is emerging as social technology develops: something in between an interpersonal relationship and an interaction with an inanimate object.

However, like virtual games that can be used either as healthy entertainment or as a replacement for real life, social, sociable technology may have dual effects. How can we design technology to provide social benefits while mitigating potential negative impacts? Should technology detect signals that people's relationships with and mental models of technology are becoming unhealthy? When should we promote sociality or social agency and when should we not? How can we diminish people's tendencies to give agency to and anthropomorphize technology? Is it possible to understand the long-term effects of such technology before it goes on the market for wider adoption?

9.2.2 Personalization Technology

Personalizing services that have not been traditionally personalized could also result in unintended consequences. For examples, online services that used to deal with public domains, such as maps, are now increasingly becoming personalized. For example, Google announced that instead of showing the same map to everyone, they could potentially personalize information on a Google map to match people's interests and needs. What would happen if this traditionally public service were personalized? Would this influence people's perceptions of "public space"? Could this personalization disadvantage some individuals and benefit others? What if personalization unintentionally promotes comparison and discrimination, as even a robot's seemingly small personalized social gestures did among a social group (Chapter 7)? Careful thought needs to be given to whether certain tasks or services should never be personalized, and whether some personalized services that benefit individuals may not be beneficial to society in general.

9.2.3 Autonomous Technology and Human Labor

The advancement of autonomous, intelligent technology may improve people's quality of life by doing dangerous or laborious tasks and expanding services that human service providers usually offer to a broader spectrum of the population. For example, assistive telepresence and mobile robots can help older adults live independently in their homes by allowing human care givers to remotely check-in and monitor, and assisting with tasks that older adults cannot easily do by themselves. With this kind of technology service, people who could not afford to hire human in-home caregivers could still benefit from such services.

The development of these intelligent technology services needs to be accompanied by further research on how they impact our society. As previous technological developments in history have changed the types of jobs and skills in demand, this technology will also ultimately transform how we work and how we

educate people. For example, Brynjolfsson and McAfee (2011, 2013) conjecture that the development of intelligent machines will let people have more time and resources to focus on tasks that need human creativity and innovation; however, they argue that these machines may take away jobs from people, especially the “blue-collar” workers such as truck drivers and those who work in warehouse facilities. They argue that we need to start preparing people for a world where many jobs will be replaced by machines and focus on improving people’s creativity.

Robotics researchers argue that it will be a considerable amount of time – at least a decade to a half-century – until these machines will be completely able to execute tasks that people must currently do (Brook, 2013); in the mean time, people will be needed to be manage and program these intelligent technologies. Designers and developers of such technology will need to ensure that people do not feel threatened, and that humans and technology can share tasks and responsibilities in ways that respect human dignity.

References

- Aaker, J., Fournier, S., & Brasel, S. (2004). When good brands do bad. *Journal of Consumer Research*, 31(1), 1-16.
- Abelson, R. P. (1981). Psychological status of the script concept. *American Psychologist*, 36(7), 715-729.
- Adomavicius, G., & Tuzhilin, A. (1999). User profiling in personalization applications through rule discovery and validation. In *Conference on Knowledge Discovery in Data: Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*, 15(18), 377-381.
- Adomavicius, G., & Tuzhilin, A. (2005). Personalization technologies: A process-oriented perspective. *Communications of the ACM*, 48(10), 83-90.
- Aggarwal, P. (2004). The effects of brand relationship norms on consumer attitudes and behavior. *Journal of Consumer Research*, 31(1), 87-101.
- Akerlof, G., & Kranton, R. (2010). *Identity economics: How our identities shape our work, wages and well-being*. Princeton University Press.
- Aleahmad, T., Alevan, V., & Kraut, R. (2009). Creating a corpus of targeted learning resources with a web-based open authoring tool, *IEEE Transactions on Learning Technologies*, 2(1), 3-9.
- Amazon mTurk, <https://www.mturk.com>.
- Anand, S. S., & Mobasher, B. (2003). Intelligent techniques for web personalization. In *Proceedings of the 2003 international conference on Intelligent Techniques for Web Personalization*, 1-36.
- Anderson, B., Hagen, C., Reifel, J., & Stettler, E. (2006). Complexity: Customization's evil twin. *Strategy & Leadership*, 34(5), 19-27.

- Ansari, A., & Mela, C. F. (2003). E-customization. *Journal of Marketing research*, 40(2), 131-145.
- Arbanowski, S., Ballon, P., David, K., Droegehorn, O., Eertink, H., Kellerer, W., Van Kranenburg, H., Raatikainen, K., & Popescu-Zeletin, R. (2004). I-centric communications: Personalization, ambient awareness, and adaptability for future mobile services. *Communications Magazine, IEEE*, 42(9), 63-69.
- Ashman, H., Brailsford, T., & Brusilovsky, P. (2009). Personal services: Debating the wisdom of personalization. In *Proceedings of ICWL 2009*, 1-11.
- Awad, N. F., & Krishnan, M. S. (2006). The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS Quarterly*, 30(1), 13-28.
- Balabanovic, M., & Shoham, Y. (1997). Fab: content-based, collaborative recommendation. *Communication of ACM*, 40, 66-72.
- Ball, D., Coelho, P. S., & Vilares, M. J. (2006). Service personalization and loyalty. *Journal of Services Marketing*, 20(6), 391-403.
- Barley, S. R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly*, 31(1), 78-108.
- Battarbee, K., & Koskinen, I. (2005). Co-experience: User experience as interaction. *CoDesign*, 1(1), 5-18.
- Belk, R. W. (1988). Possessions and the extended self. *Journal of Consumer Research*, 15(2), 139-168.
- Bell, A. (1984). Language style as audience design. *Language in Society*, 13, 145-204.

- Bellisle, F., Dalix, A.M., Mennen, L., Galan, P., Hercberg, S., de Castro, J.M., & Gausseres, N. (2003). Contribution of snacks and meals in the diet of French adults: A diet-diary study. *Physiological Behavior*, 79, 183-189.
- Berkovsky, S., Freyne, J., & Oinas-Kukkonen, H. (2012). Influencing individually: Fusing persuasion and personalization. *ACM TiiS*, 2(2), 9-16.
- Berkovsky, S., Kuflik, T., & Ricci, F. (2008). Mediation of user models for enhanced personalization in recommender systems. *User Modeling and User-Adapted Interaction*, 18(3), 245-286.
- Bethel, C. L., & Murphy, R. (2006). Affective expression in appearance constrained robots. In *Proceedings of HRI'06*, 327-328.
- Bickmore, T., & Cassell, J. (2001). Relational agents: A model and implementation of building user trust. In *Proceedings of CHI'01*, 396-403.
- Bickmore, T., Pfeifer, L., & Jack, B. (2009). Taking the time to care: Empowering low health literacy hospital patients with virtual nurse agents. In *Proceedings of CHI'09*, 1265-1274.
- Bickmore, T., Pfeifer, L., Schulman, D., Perera, S., Senanayake, C., & Nazmi, I. (2008). Public displays of affect: Deploying relational agents in public spaces. In *Proceedings of CHI'08, Work-In-Progress*, 3297-3302.
- Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction*, 12, 293-327.
- Bitner, M. J. (1992). Servicescapes: The impact of physical surroundings on customers and employees. *Journal of Marketing*, 56, 57-71.
- Bitner, M.J., Booms, B.H., & Tetreault, M.S. (1990). The service encounter: Diagnosing favorable and unfavorable incidents. *The Journal of Marketing*, 71-84.

- Bitner, M.J., & Brown, S.W. (2006). The evolution and discovery of services science in business schools. *Communications of the ACM*, 49, 7, 72-78.
- Bitner, M.J., Brown, S. W., & Meuter, M. L. (2000). Technology infusion in service encounters. *Journal of the Academy of Marketing Science*, 28(1), 138-149.
- Bitner, M. J., Faranda, W. T., Hubbert, A. R., & Zeithaml, V. A. (1997). Customer contributions and roles in service delivery. *International Journal of Service Industry Management*, 8(3), 193-205.
- Bitner, M. J., Ostrom, A. L., & Meuter, M. L. (2002). Implementing successful self-service technologies. *Academy of Management Executive*, 16(4), 96-109.
- Bitner, M.J., Ostrom, A.L., & Morgan, F.N. (2007) Service blueprinting: A practical tool for service. *California Management Review*, 50(3), 66
- Blom, J. (2000). Personalization: A taxonomy. In *CHI'00 extended abstracts on Human factors in computing systems*, 313-314.
- Blom, J., & Monk, A. F. (2003). Theory of personalization of appearance: Why users personalize their PCs and mobile phones. *Human-Computer Interaction*, 18(3), 193-228.
- Bloom, B. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4-16.
- Blow, M., Dautenhaun, K., Appleby, A., Nehaniv, C.L., & Lee, D. (2006). The art of designing robot faces: Dimensions for human-robot interaction. In *Proceedings of HRI'06*, 331-332.
- Brave, S., Nass, C., & Hutchinson, L. (2005). Computers that care: Investigating the effects of orientation of emotion exhibited by an embodied computer agent. *International Journal of Human-Computer Studies*, 62, 161-178.

- Breazeal, C. (2003). Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies*, 59, 119- 155.
- British Design Council, service design case study:
<http://www.designcouncil.org.uk/en/Case-Studies>
- Brown, P., & Levinson, S. (1987). *Politeness: Some universals in language usage*. Cambridge University Press.
- Bruce, A., Nourbakhsh, I., & Simmons, R. (2002). The role of expressiveness and attention in human-robot interaction. *In Proceedings of ICRA'02*, 4138-4142.
- Brynjolfsson, E., & McAfee, A. (2011). *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving productivity, and Irreversibly Transforming Employment and the Economy*. Lexington, MA: Digital Frontier Press.
- Brook, R. (2013). Oral presentation at TED 2013
- Bureau of Labor Statistics, U.S. Dept. of Labor. Occupational outlook handbook. <http://www.bls.gov/ooh/About/Projections-Overview.htm> (accessed September, 2012).
- Burgard, W., Cremers, A.B., Fox, D., Hähnel, D., Lakemeyer, G., Schulz, D., Steiner, W., & Thrun, S. (1999). Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114, 3-55.
- Cai, Y., & Sullivan, K. J. (2005). A value-oriented theory of modularity in design. *ACM Sigsoft Software Engineering Notes*, 30(4), 1-4.
- Carver, C. S., & Scheier, M. F. (2000). Autonomy and self-regulation. *Psychological Inquiry*, 11, 284–290.
- Center for Behavioral Decision Research, <http://www.cbdr.cmu.edu/>

- Cesario, J., Grant, H., & Higgins, E. T. (2004). Regulatory fit and persuasion: Transfer from “feeling right.” *Journal of Personality and Social Psychology*, 86, 388-404.
- Chase, R. B., & Apte, U. M. (2007). A history of research in service operations: What's the *big idea*?. *Journal of Operations Management*, 25(2), 375-386.
- Chesbrough, H., & Spohrer, J. (2006). A research manifesto for services science. *Communications of the ACM*, 49(7), 35-40.
- Cingil, I., Dogac, A., & Azgin, A. (2000). A broader approach to personalization. *Communications of the ACM*, 43(8), 136-141.
- Clark, H., & Brennan, S. (1991). Grounding in communication. In *Perspectives on Socially Shared Cognition*, 127-149.
- Claypool, H. M., Mackie, D. M., Garcia-Marques, T., McIntosh, A., & Udall, A. (2004). The effects of personal relevance and repetition on persuasive processing. *Social Cognition*, 22(3), 310-335.
- Cordova, D. I., & Lepper, M. R. (1996). Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of Educational Psychology*, 88(4), 715-730.
- Craigslist, <http://pittsburgh.craigslist.org/>.
- Cramer, H., Evers, V., Ramlal, S., Van Someren, M., Rutledge, L., Stash, N., Aroyo, L., & Wielinga, B. (2008). The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction*, 18(5), 455-496.
- Dabholkar, P. A., & Bagozzi, R. P. (2002). An attitudinal model of technology-based self-service: Moderating effects of consumer traits and situational factors. *Journal of the Academy of Marketing Science*, 30(3), 184-201.

- Dabholkar, P. A., & Sheng, X. (2012). Consumer participation in using online recommendation agents: Effects on satisfaction, trust, and purchase intentions. *The Service Industries Journal*, 32(9), 1433-1449.
- Das, A. S., Datar, M., Garg, A., & Rajaram, S. (2007). Google news personalization: Scalable online collaborative filtering. In *Proceedings of the 16th international conference on World Wide Web*, 271-280.
- Dautenhahn, K. (2004). Robots we like to live with: A developmental perspective on a personalized, life-long robot companion. In *Proceedings of ROMAN'04*, 17-22.
- Dautenhaun, K., Walters, M., Woods, S., Koay, K.L., Nehaniv, C.L., Sisbot, A., Alami, A., & Simeon, T. (2007). How May I Serve You?: A robot companion approaching a seated person in a helping context. In *Proceedings of HRI'07*, 172-179.
- DeSanctis, G. & Poole, M.S. (1994). Capturing the complexity in advanced technology use: Adaptive structuration theory. *Organization Science*, 5(2), 121-147.
- Diekmann, K., Samuels, S., Ross, L., & Bazerman, M. (1997). Self-interest and fairness in problems of resource allocation: Allocators vs. recipients. *Journal of Personality and Social Psychology*, 72, 1061-1074.
- Dillenbourgh, P., Jermann, P., Schneider, D., Traum, D., & Buiou, C. (1997). The design of MOO agents: Implications from an empirical CSCW study. *Artificial Intelligence in Education*. IOS Press.
- DiSalvo, C., Gemperle, F., Forlizzi, J., & Kiesler, S. (2002). All robots are not created equal: The design and perception of humanoid robot heads. In *Proceedings of DIS'02*, 321-326.

- Dourish, P., & Bly, S. (1992). Portholes: Supporting awareness in a distributed work group. In *Proceedings CHI'92*, 541-547.
- Ducheneaut, N., Wen, M. H., Yee, N., & Wadley, G. (2009). Body and mind: A study of avatar personalization in three virtual worlds. In *Proceedings of CHI'09*, 1151-1160.
- Duck, S. (1991). *Understanding Relationships*. Guilford Press.
- Dweck, C. (2006). *Mindset: The New Psychology of Success*. Random House Digital, Inc.
- Eirinaki, M., & Vazirgiannis, M. (2003). *Web mining for web personalization*. *ACM Transactions on Internet Technology (TOIT)*, 3(1), 1-27.
- Enembreck, F., & Barthes, J. P. (2002). Personal assistant to improve CSCW. *Proceedings of CSCW in Design*, 329-335.
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4), 864- 886.
- Evers, V., Maldonado, H., Brodecki, T., & Hinds, P. J. (2008). Relational vs. group self-construal: Untangling the role of national culture in HRI. In *Proceedings of HRI'08*, 255-262.
- Fan, H., & Poole, M. S. (2006). What is personalization? Perspectives on the design and implementation of personalization in information systems. *Journal of Org. Computing and Electronic Commerce*, 16(3-4), 179-202.
- Feil-Seifer, D., & Mataric, M. J. (2005). Defining socially assistive robotics. In *Proceedings of ICORR'05*, 465-468.
- Feldman, M., & Pentland, B. (2003). Reconceptualizing organizational routines as a source of flexibility and change. *Administrative Science Quarterly*, 48, 94-118.

- Fink, J., & Kobsa, A. (2000). A review and analysis of commercial user modeling servers for personalization on the world wide web. *User Modeling and User-Adapted Interaction*, 10, 209-249.
- Fischer, K. (2007). The Role of Users' Concepts of the Robot in Human-Robot Spatial Instruction. *Lecture Notes In Computer Science*, 4387, 76-89.
- Fischer, K. (2006). The Role of Users' Preconceptions in Talking to Computers and Robots. In *Proceedings of the Workshop on How People Talk to Computers, Robots, and other Artificial Communication Partners*, 112-130.
- Flanagan, J. C. (1954). The critical incident technique. *Psychological Bulletin*, 51(4), 327-358.
- Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., & Yang, J. (2005). Predicting human interruptibility with sensors. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 12, 119-146.
- Fong, T., Nourbakhsh, I., & Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42, 143-166.
- Forlizzi, J. (2007). How robotic products become social products: An ethnographic study of cleaning in the home. In *Proceedings of HRI'07*, 129-136.
- Forlizzi, J., & Battarbee, K. (2004). Understanding experience in interactive systems. In *Proceedings of DIS'04*, 261-268.
- Forlizzk, J., & DiSalvo, C. (2006). Service robots in the domestic environment: A study of the Roomba vacuum in the home. In *Proceedings of HRI'06*, 258-265.
- Forlizzi, J., & Ford, S. (2000). The building blocks of experience: An early framework for interaction designers. In *Proceedings of DIS'00*, 419-423.
- Fournier, S. (1998). Consumers and their brands: Developing relationship theory in consumer research. *Journal of Consumer Research*, 24(4), 343-353.

- Friedman, B., Kahn, P.J., & Hagman, J. (2003). Hardware companions?: What online AIBO discussion forums reveal about the human-robotic relationship. In *Proceedings of CHI'03*, 273-280.
- Fujita, K., & Han, A. (2009). Moving beyond deliberative control of impulses: The effect of construal levels on evaluative associations in self-control conflicts. *Psychological Science*, 20, 799-804.
- Fulk, J. (1993). Social construction of communication technology. *Academy of Management Journal*, 921-950.
- Gajos, K., Czerwinski, M., Tan, D. S., & Weld, D. S. (2006). Exploring the design space for adaptive graphical user interfaces. In *Proceedings of AVI'06*, 201-208.
- Gallant, L., Irizarry, C., & Kreps, G. L. (2006). User-centric hospital web sites: A case for trust and personalization. *E-Service Journal*, 5(2), 5-26.
- George, M. (2008). Interactions in expert service work: Demonstrating professionalism in personal training. *Journal of Contemporary Ethnography*, 37, 108-131.
- Glushko, R. J. (2010). Seven contexts for service system design. *Handbook of Service Science*, 219-249.
- Glushko, R. J., & Nomorosa, K. J. (2013). Substituting information for interaction: A framework for personalization in service encounters and service systems. *Journal of Service Research*, 16(1), 21-38.
- Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A.C., & Wang, J. (2005). Designing robots for long-term social interaction. In *Proceedings of IROS'05*, 2199-2204.
- Gockley, R., Simmons, R., & Forlizzi, J. (2006). Interactions with a moody robot. In *Proceedings of HRI'06*, 186-193.

- Gockely, R., Simmons, R., & Forlizzi, J. (2007). Natural person-following behavior for social robots. In *Proceedings of HRI'07*, 17-24.
- Goffman, E. (1959). *The Presentation of Self in Everyday Life*. Anchor Books, New York, NY.
- Goffman, E. (1963). *Behavior in public places*. The Free Press.
- Goetz, J., Kiesler, S., & Powers, A. (2003). Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *Proceedings of ROMAN'03*, 55-60.
- Grasso, F., & Paris, C. (2011). Preface to the special issue on personalization for e-health. *User Modeling and User-Adapted Interaction*, 21(4), 333-340.
- Griffith, T. L. (1999). Technology features as triggers for sensemaking. *Academy of Management review*, 24(3), 472-488.
- Gronroos, C. (1990a). Chapter 2. The nature of services and service quality. In *Service Management and Marketing: Managing the Moments of Truth in Service Competition*. 25-48.
- Gronroos, C. (1990b). Chapter 3. Managing customer-perceived service quality. In *Service Management and Marketing: Managing the Moments of Truth in Service Competition*. 49-70.
- Gronroos, C. (2008). Service logic revisited: Who creates value? And who co-creates? *European Business Review*, 20(4), 298-314.
- Guo, L., Arnould, E. J., Gruen, T. W., & Tang, C. (2013). Socializing to co-produce: Pathways to consumers' financial well-being. *Journal of Service Research*, online first publication.

- Gwinner, K.P., Gremier, D.D., & Bitner, M.J. (1998). Relational benefits in services industries: The customer's perspective. *Journal of Academy of Marketing Science*, 26(2), 101-114.
- Hada, Y., Gakuhari, H., Takase, K., & Hemeldan, E.I. (2004). Delivery service robot using distributed acquisition, actuators and intelligence. In *Proceedings of IROS'04*, 2997-3002.
- Hallowell, R. (1996). The relationships of customer satisfaction, customer loyalty, and profitability: An empirical study. *International Journal of Service Industry Management*, 7(4), 27-42.
- Harris, S. G. (1994). Organizational culture and individual sensemaking: A schema-based perspective. *Organization Science*, 5, 309-321.
- Hart, C.W.L., Heskett, J.L., & Sasser Jr, W.E. (1990). The profitable art of service recovery. *Harvard Business Review*, 148-156.
- Hayashi, K., Sakamoto, D., Kanda, T., Shiomi, M., Koizumi, S., Ishiguro, H., Ogasawara, T., & Hagita, N. (2007). Humanoid robots as a passive-social medium: A field experiment at a train station. In *Proceedings of HRI'07*, 137-144.
- Heckhausen, J., & Schultz, R. (1995). A life-span theory of control. *Psychological Review*, 102, 284-304.
- Hinde, R. A. (1976). Interactions, relationships and social structure. *Man*, 1-17.
- Hinde, R. A. (1995). A suggested structure for a science of relationships. *Personal Relationships*, 2(1), 1-15.
- Hinds, P., Roberts, T. L., & Jones, H. (2004). Whose job is it anyway? A study of human-robot interaction in a collaborative task. *Human-Computer Interaction*, 19, 151-181.

- Ho, S. Y., Davern, M. J., & Tam, K. Y. (2008). Personalization and choice behavior: the role of personality traits. *ACM SIGMIS Database*, 39(4), 31-47.
- Houston, M. J., & Jefferson, R. W. (1975). The negative effects of personalization on response patterns in mail surveys. *Journal of Marketing Research*, 12(1), 114-117.
- Ishizaki, S. (2010). Service design and technical/professional communication. *Professional Communication Conference*, 1-5.
- Jensen, G. M., Gwyer, J., Shepard, K. F., & Hack, L. M. (2000). Expert practice in physical therapy. *Physical Therapy*, 80, 1, 28-43.
- Johnson, D., Gardner, J., & Wiles, J. (2004). Experience as a moderator of media equation: The impact of flattery and praise. *International Journal of Human-Computer Studies*, 61(3), 237-258.
- Jones, H., & Hinds, P. (2002). Extreme work teams: Using SWAT teams as a model for coordinating distributed robots. In *Proceedings of CSCW'02*, 372-381.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics, *The American Economic Review*, 93(5), 1449-1475.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Kanda, T., Hirano, T., Eaton, D. & Ishiguro, H. (2004). Interactive robots as social partners and peer tutors for children: A field trial. *Human Computer Interaction*, 19(1), 61-84.
- Kaptein, M., Ruyter, B. D., Markopoulos, P., & Aarts, E. (2012). Adaptive persuasive systems: A study of tailored persuasive text messages to reduce snacking. *ACM TiiS*, 2(2), 10-34.
- Karapanos, E., Hassenzahl, M., & Martens, J.B., User experience over time. In *Proceedings of CHI'08*, 419- 432

- Karapanos, E., Zimmerman, J., & Forlizzi, J. (2009). User experience over time: an initial framework. In *Proceedings of CHI'09*, 729-738.
- Kay, J., & Kummerfeld, B. (2012). Creating personalized systems that people can scrutinize and control: Drivers, principles and experience. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 2(4), 42 pages.
- Keaveney, S. (1995). Customer switching behavior in service industries: An exploratory study. *The Journal of Marketing*, 59(2), 71-82.
- Kelley, S.W., Hoffman, D.H., & Davis, M.A. (1993). A typology of retail failures and recoveries. *Journal of Retailing*, 69, 429-429.
- Kendon, A. (1990). A description of some human greetings. In *Conducting interaction: Patterns of behavior in focused encounters*. CUP Archive.
- Kidd, C. D., & Breazeal, C. (2008). Robots at home: Understanding long-term human-robot interaction. In *Proceedings of IROS'08*, 3230-3235.
- Kiesler, S., Powers, A., Fussell, S. R., & Torrey, C. (2008). Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition*, 26(2), 169-181.
- Kim, H. R., & Chan, P. K. (2003). Learning implicit user interest hierarchy for context in personalization. In *Proceedings of the 8th international conference on Intelligent user interfaces*, 101-108.
- Kim, T., & Hinds, P. (2006). Who should I blame? Effects of autonomy and transparency on attributions in human-robot interaction. In *Proceedings of ROMAN'06*, 80-85.
- Kittur, A., Chi, E.H., & Suh, B. (2008). Crowdsourcing user studies with Mechanical Turk. In *Proceedings of CHI'08*, 453-456.
- Knapp, M. L., & Hall, J. A. (2009). *Nonverbal Communication in Human Interaction*. Cengage Learning.

- Kobsa A. (2001). Generic user modeling systems. *User Modeling and User-Adapted Interaction*, 11, 49-63.
- Kobsa, A. (2007). Privacy-enhanced web personalization. In Brusilovsky, P., Kobsa, A., & Nejdl, W. (Eds.), *The Adaptive Web* (pp. 628-670).
- Komiak, S., & Benbasat, I. (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS Quarterly*, 30(4), 941-961.
- Kramer, J., Noronha, S., & Vergo, J. (2000). A user-centered design approach to personalization. *Communications of the ACM*, 43(8), 44-48.
- Kramer, T., Suri, S-W., & Thakkar, M. (2007). The effect of cultural orientation on consumer responses to personalization. *Marketing Science*, 26(2), 246-258.
- Kross, E., Ayduk, O., & Mischel, W. (2005). When asking “why” does not hurt: Distinguishing rumination from reflective processing of negative emotions. *Psychological Science*, 16, 709-715.
- Kulesza, T., Stumpf, S., Burnett, M., & Kwan, I. (2012). Tell me more?: The effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of CHI'12*, 1-10.
- Kunda, Z. (1999). *Social Cognition: Making Sense of People*. Cambridge, MA:MIT Press.
- Kuno, Y., Sadazuka, K., Kawashima, M., Yamazaki, K., Yamazaki, A., & Hideaki Kuzuoka. (2007). Museum guide robot based on sociological interaction analysis. In *Proceedings of CHI'07*, 1191-1194.
- Kuzuoka, H., Pitsch, K., Suzuki, Y., Kawaguchi, I., Yamazaki, K., Yamazaki, A., Kuno, Y., Luff, P., & Heath, C. (2008). Effect of restarts and pauses on achieving a state of mutual orientation between a human and a robot. In *Proceedings of CSCW '08*, 201-204.

- Lai, J., Levas, A., Chou, P., Pinhanez, C., & Viveros, M. (2002). BlueSpace: Personalizing workspace through awareness and adaptability. *International Journal of Human-Computer Studies*, 57(5), 415-428.
- Lawrence, R. D., Almasi, G. S., Kotlyar, V., Viveros, M., & Duri, S. S. (2001). Personalization of supermarket product recommendations. *Applications of Data Mining to Electronic Commerce*, 11-32.
- Lee, E. J., Nass, C., & Brave, S. (2000). Can computer-generated speech have gender?: An experimental test of gender stereotype. In *Proceedings of CHI'00, Extended Abstract*, 289-290.
- Lee, M.K., & Forlizzi, J. (2009). Designing adaptive robotic services. In *Proceedings of IASDR'09*.
- Lee, M. K., Forlizzi, J., Kiesler, S., Rybski, P., Antanitis, J., & Savetsila, S. (2012). Personalization in HRI: A longitudinal field experiment. In *Proceedings of HRI'12*, 319-326.
- Lee, M.K., Forlizzi, J., Rybski, P.E., Crabbe, F., Chung, W., Finkle, J., Glaser, E., & Kiesler, S. (2009). The Snackbot: Documenting the design of a robot for long-term human-robot interaction. In *Proceedings of HRI'09*, 7- 14.
- Lee, M.K., Kiesler, S., & Forlizzi, J. (2008). How do people snack? Understanding the context of a mobile robot snack service. Unpublished ms. Carnegie Mellon Univ., Pittsburgh, PA. 15213.
- Lee, M.K., Kiesler, S., & Forlizzi, J. (2010). Receptionist or information kiosk? How do people talk with a robot? In *Proceedings of CSCW'10*, 31-40.
- Lee, M. K., Kiesler, S., & Forlizzi, J. (2011). Mining behavioral economics to design persuasive technology for healthy choices. In *Proceedings of CHI'11*, 325-334.

Lee, M. K., Kiesler, S., Forlizzi, J., & Rybski, (2012). Ripple effects of an embodied social agent: A field study of a social robot in the workplace. In *Proceedings of CHI'12*, 695-704.

Lee, M.K., Kiesler, S., Forlizzi, J., Srinivasa, S., & Rybski, P. (2010). Gracefully mitigating breakdowns in robotic services. In *Proceedings of HRI'10*, 203-210.

Lee, M. K., & Makatchev, M. (2009). How do people talk with a robot?: An analysis of human-robot dialogues in the real world. In *Proceedings of CHI'09, Extended Abstract*, 3769-3774.

Lee, S., Lau, I.Y., Kiesler, S., & Chiu, C. (2005). Human mental models of humanoid robots. In *Proceedings of ICRA'05*, 2767-2772.

Lenzo, K.A., & Black, A.W., Theta, Cepstral, <http://www.cepstral.com>.

Leshed, G., Hancock, J. T., Cosley, D., McLeod, P. L., & Gay, G. (2007). Feedback for guiding reflection on teamwork practices. In *Proceedings of Group'07*, 217-220.

Liang, T-P., Lai, H-J., & Ku, Y-C. (2006). Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *Journal of Management of Information Systems*, 23(3), 45-70.

Liechty, J., Ramaswamy, V., & Cohen, S. (2001). Choice menu for mass customization: An experimental approach for analyzing customer demand with an application to a web-based information service. *Journal of Marketing Research*, 38(2), 183-196.

Lim, B. Y., Dey, A. K., & Avrahami, D. (2009). Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of CHI'09*, 2119-2128.

McAfee, A. (2013). What happens after the droids take our jobs? Oral presentation at TED 2013.

- Mager, B., & Evenson, S. (2008). The Art of Service: Drawing on the Arts to Inform Service Design and Specification. In Eds. Bill Hefley and Wendy Murphy, *Service Science: Research and Innovations in the Service Economy*. London, UK: Springer (pp. 75-76).
- Makatchev, M., Lee, M. K., & Simmons, R. (2009). Relating initial turns of human-robot dialogues to discourse. In *Proceedings of HRI'09*, 321-322.
- Malheiros, M., Jennett, C., Patel, S., Brostoff, S., & Sasse, A. (2012). Too close for comfort: A study of the effectiveness and acceptability of rich-media personalized advertising. In *Proceedings of CHI'12*, 579-588.
- Maslow, A. H. (1943). A theory of human motivation. *Psychological review*, 50, 370-396.
- Mason, M., & Lopes, M. (2011). Robot self-initiative and personalization by learning through repeated interactions In *Proceedings of HRI'11*, 433-440.
- Mayer, J. D., Faber, M. A., & Xu, X. (2007). Seventy-five years of motivation measures (1930–2005): A descriptive analysis. *Motivation and Emotion*, 31(2), 83-103.
- Mayor-Schonberger, V., & Cukier, K. (2013). *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Eamon Dolan / Houghton Mifflin.
- McCarthy, J., & Wright, P. (2004). *Technology As Experience*. MIT Press.
- McCracken, G. (1986). Culture and consumption: A theoretical account of the structure and movement of the cultural meaning of consumer goods. *Journal of Consumer Research*, 71-84.
- McDonough, E. F., Kahn, K. B. & Barczaka, G. (2001). An investigation of the use of global, virtual, and collocated new product development teams. *Journal of Product Innovation Management* 18(2), 110-120.

- Mehrabian, A. (1970). A semantic space for nonverbal behavior. *Journal of Consulting and Clinical Psychology*, 35(2), 248-257.
- Meuter, M. L., Bitner, M. J., Ostrom, A. L., & Brown, S. W. (2005). Choosing among alternative service delivery modes: An investigation of customer trial of self-service technologies. *Journal of Marketing*, 69, 61-83.
- Meuter, M. L., Ostrom, A. L., Roundtree, R. I., & Bitner, M. J. (2000). Self-service technologies: Understanding customer satisfaction with technology-based service encounters. *Journal of Marketing*, 64(3), 50-64.
- Michalowski, M. P., Sabanovic, S., & Simmons, R. (2006). A spatial model of engagement for a social robot. In *Advanced Motion Control '06*, 762-767.
- Milkman, K. L., Chugh, D., & Bazerman, M. H. (2009). How can decision making be improved? *Perspectives on Psychological Science*, 4(4), 379-383.
- Mitchell, T. (2009). Mining our reality. *Science*, 1644-45.
- Mittal, B., & Lassar, W. M. (1996). The role of personalization in service encounters. *Journal of Retailing*, 72(1), 95-109.
- Moeller, S., Ciuchita, R., Mahr, D., Odekerken-Schorder, G., & Fassnacht, M. (2013). Uncovering collaborative value creation patterns and establishing corresponding customer roles. *Journal of Service Research*, online first publication.
- Moon, Y., & Nass, C. (1996). Adaptive agents and personality change: Complementarity versus similarity as forms of adaptation. In *Proceedings of CHI'96*, 287-288.
- Morelli, N. (2002). Designing product/service systems. A methodological exploration. *Design Issues*, 18, 3-17.

Morkes, J., Kernal, H., & Nass, C. (1999). Effects of humor in task-oriented human-computer interaction and computer-mediated communication: A direct test of SRCT theory. *Human-Computer Interaction, 14*, 395- 435.

Mutlu, B., & Forlizzi, J. (2008). Robots in organizations: The role of workflow, social, and environmental factors in human-robot interaction. In *Proceedings of HRI'08*, 287-294.

Mutlu, B., Osman, S., Forlizzi, J., Hodgins, J., & Kiesler, S. (2006). Perceptions of ASIMO: An exploration of co- operation and competition with humans and humanoid robots. In *Proceedings of HRI'06*, 351-352.

Nabe, S., Kanda, T., Hiraki, K., Ishiguro, H., Kogure, K., & Hagita, N. (2007). Analysis of human behavior to a communication robot in an open field. In *Proceedings of HRI'07*, 234-241.

Nakagawa, K., Shiomi, M., Shinozawa, K., Matsumura, R., Ishiguro, H., & Hagita, N. (2011). Effect of robot's active touch on people's motivation. In *Proceedings of HRI'11*, 465-472.

Nakauchi, Y., & Simmons, R. (2002). A social robot that stands in line. *Autonomous Robots, 12*, 313-324.

Nansel, T. R., Iannotti, R. J., Simons-Morton, B. G., Cox, C., Plotnick, L., Clark, L. M., & Zeitzoff, L. (2007). Diabetes personal trainer outcomes: Short-term and 1-year outcomes of a diabetes personal trainer intervention among youth with type 1 diabetes. *Diabetes Care, 30*(10), 2471-2477.

Nass, C., & Lee, K. M. (2000). Does computer- generated speech manifest personality? An experimental test of similarity-attraction. In *Proceedings of CHI'00*, 329-336.

- Nass, C., & Lee, K. M. (2001). Does computer-synthesized speech manifest personality? Experimental tests of recognition, similarity-attraction, and consistency attraction. *Journal of Experimental Psychology*, 171-181.
- Nelson, H.G., & Stolterman, E. (2003). *The Design Way*. Englewood Cliffs, NJ: Educational Technology Publications.
- Nisbett, R., Caputo, C., Legant, P., & Marecek, J. (1973). Behavior as seen by the actor and as seen by the observer. *Journal of Personality and Social Psychology*, 27(2), 154-164.
- Nisbett, R., & Wilson, T. (1977). "Telling more than we can know": Verbal reports on mental processes. *Psychological Review*, 84(3), 231-259.
- Nomura, T., Kanda, T., Suzuki, T., & Kato, K. (2008). Prediction of Human Behavior in Human--Robot Interaction Using Psychological Scales for Anxiety and Negative Attitudes Toward Robots. *IEEE Transactions on Robotics*, 24(2), 442-451.
- Nowak, M., & Nass, C. (2012). Effects of behavior monitoring and perceived system benefit in online recommender systems. In *Proceedings of CHI'12*, 2243-2246.
- Oliver, A. (1980). Cognitive model of the antecedents and consequences of satisfaction decisions. *Journal of Marketing Research*, 460-469.
- Orlikowski, W. J. (2000). Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations. *Organization Science*, 11(4), 404-428.
- Ostrom, A. L., Bitner, M. J., Brown, S. W., Burkhard, K. A., Goul, M., Smith-Daniels, V., Demirkan, H., & Rabinovich, E. (2010). Moving forward and making a difference: Research priorities for the science of service. *Journal of Service Research*, 13(1), 4-36.

Ovaskainen, M.L., Reinivuo, H., Tapanainen, H., Hannila, M-L., Korhonen, T., & Pakkala, H. (2006). Snacks as an element of energy intake and food consumption. *European Journal of Clinical Nutrition*, 60, 494- 501.

Padmanabhan, B., Zheng, Z., & Kimbrough, S. O. (2001). Personalization from incomplete data: What you don't know can hurt. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge Discovery and Data Mining*, 154-163.

Parasuraman, A., Zeithaml, V. A., & Berry, L. L. (1985). A conceptual model of service quality and its implications for future research. *Journal of Marketing*, 49, 41-50.

Parise, S., Kiesler, S., Sproull, L., & Water, K. (1996) My partner is a real dog: Cooperation with social agents. In *Proceedings of CSCW'96*, 399-408.

Pariser, E. (2011). *The filter bubble: How the new personalized web is changing what we read and how we think*. NY, USA: Penguin Group.

Paulhus, D. (1983). Sphere-specific measures of perceived control. *Journal of Personality and Social Psychology*, 44(6), 1253-1265.

Pfeifer, L. M. & Bickmore, T. (2011). Is the media equation a flash in the pan? The durability and longevity of social responses to computers. In *Proceedings of CHI'10*, 777-780.

Pine, B. J., & Gilmore, J. H. (1998). Welcome to the experience economy. *Harvard Business Review*, 76, 97-105.

Pinhanez, C. (2009). Service as customer-intensive systems. *Design Issues*, 25(2), 3-13.

Pinhanez, C. S. (2009) A service science perspective on human-computer interface issues of online service applications. *International Journal of Information Systems and Services Sector*, 1(2),17–35.

- Pinhanez, C. S. (2011). Human aspects of internet services: Considering the needs of users and providers. *Journal of Internet Services and Application*, 1(3), 155-164.
- Polaine, A., Lovlie, L., & Reason, B. (2013). *Service Design*. Rosenfeld Media.
- Pollack, M.E., Engberg, S., Matthews, J. T., Thrun, S., Brown, L., Colbry, D., Orosz, C., Peintner, B., Ramakrishnan, S., Dunbar-Jacob, J., McCarthy, C., Montemerlo, M., Pineau, J., & Roy, N. (2002). Pearl: A mobile robotic assistant for the elderly. In *Proceedings of AAAI Workshop on Automation as Eldercare*, AAAI, 2002.
- Powers, A., & Kiesler, S. (2006). The advisor robot: Tracing people's mental model from a robot's physical attributes. In *Proceedings of HRI'06*, 218-225.
- Powers, A., Kiesler, S., Fussell, S., & Torrey, C. (2007). Comparing a computer agent with a humanoid robot. In *Proceedings of HRI'07*, 145-152.
- Prahalad, C. K., & Ramaswamy, V. (2004). Co-creation experiences: The next practice in value creation. *Journal of interactive marketing*, 18(3), 5-14.
- Price, L. L., & Arnould, E. J. (1999). Commercial friendships: Service provider-client relationships in context. *Journal of Marketing*, 63, 38-56.
- Raju, P. S. (1980). Optimum stimulation level: Its relationship to personality, demographics, and exploratory behavior. *Journal of Consumer Research*, 7(3), 272-282.
- Rao, S., Hurlbutt, T., Nass, C., & JanakiRam, N. (2009). My Dating Site Thinks I'm a Loser: Effects of personal photos and presentation intervals on perceptions of recommender systems. In *Proceedings of CHI'09*, 221-224.
- Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. New York: Cambridge University Press.

- Reinders, M. J., Dabholkar, P. A., & Frambach, R. T. (2008). Consequences of forcing consumers to use technology-based self-service. *Journal of Service Research, 11*(2), 107-123.
- Remen, R. N. (1991). In the Service of Life. In *Consciousness & Healing: Integral Approaches to Mind-Body Medicine*.
- Riecken, D. (2000). Personalized views of personalization. *Communications of the ACM, 43*(8), 27-28.
- Ringber, T., Odekerken-Schroder, G., & Christensen, G.L. (2007). A cultural models approach to service recovery. *Journal of Marketing, 71*, 194-214.
- Riquelme, H. (2001). Do consumers know what they want? *Journal of Consumer Marketing, 18*(5), 437-448.
- Robins, D., Sanders, C., & Cahill, S. (1991). Dogs and their people: Pet-facilitated interaction in a public setting. *Journal of Contemporary Ethnography, 20*, 3-25.
- Rosenthal, S., Veloso, M., & Dey, A. K. Is someone in this office available to help me? Proactively seeking help from spatially-situated humans. *Journal of Intelligent & Robotic Systems, 66*(1-2), 205-221.
- Ross, L., Greene, D., & House, P. (1977). The false consensus phenomenon: An attributional bias in self-perception and social perception processes. *Journal of Experimental Social Psychology, 13*, 279-301.
- Rybski, P. E., Yoon, K., Stolarz, J., & Veloso, M. M. (2007). Interactive robot task training through dialog and demonstration. In *Proceedings of HRI'07*, 49-56.
- Sabelli, A. M., Kanda, T., & Hagita, N. (2011). A conversational robot in an elderly care center: An ethnographic study. In *Proceedings of HRI'11*, 37-44.

- Saari, T., & Turpeinen, M. (2004). Towards psychological customization of information for individuals and social groups. In Kara., C., Blom, J., & Karat, J. (Eds.), *Designing Personalized User Experiences in eCommerce* (pp. 19-37).
- Sampson, S. E., & Froehle, C. M. (2006). Foundations and implications of a proposed unified services theory. *Production and Operations Management*, 15(2), 329-343.
- Schon, D. A. (1982). *The Reflective Practitioner*. Basic Books.
- Schubert, P., & Ginsburg, M. (2000). Virtual communities of transaction: The role of personalization in electronic commerce. *Electronic Markets*, 10(1), 45-55.
- Schubert, P., Uwe, L., & Risch, D. (2006). Personalization beyond recommender systems. In *Project E-Society: Building Bricks* (pp. 126-139). Springer US.
- Schulte, J., Rosenberg, C., & Thrun, S. (1999). Spontaneous, short-term interaction with mobile robots. In *Proceedings of Robotics and Automation*, 658-663.
- Severinson-Eklundh, K., Green, A., & Huttenrauch, H. (2003). Social and collaborative aspects of interaction with a service robot. *Robotics and Autonomous Systems*, 42, 223- 234.
- Sheridan, T.B. (2002). *Humans and automations: System design and research issues*. Wiley-Interscience.
- Shostack. G.L. (1982). How to design a service. *European Journal of Marketing*, 16(1), 49-63.
- Shiomi, M., Kanda, T., Ishiguro, H., & Hagita, N. (2006). Interactive humanoid robots for a science museum. In *Proceedings of HRI'06*, 305- 312.
- Sidner, C.L., Lee, C., Morency, L.-P., & Forlines, C. (2007). The effect of head-nod recognition in human-robot conversation. In *Proceedings of HRI'07*, 290-296.

- Siino, R., &Hinds, P. (2004). Making sense of new technology as a lead-in to structuring: The case of an autonomous mobile robot. *Academy of Management Best Paper Proceedings*, New Orleans.
- Siino, R., & Hinds, P. (2005). Robots, gender & sensemaking: Sex segregation's impact on workers making sense of a mobile autonomous robot. In *Proceedings of ICRA'05*, 2773 – 2778.
- Simmons, R., Goldberg, D., Goode, A., Montemerlo, M., Roy, N., Sellner, B., Urmson, C., Schultz, A., Abramson, M., Adams, W., Atrash, A., Bugajska, M., Coblenz, M., MacMahon, M., Perzanowski, D., Horswill, I., Zubek, R., Kortenkamp, D., Wolfe, B., Milam, T., & Maxwell, B. (2003). Grace: An autonomous robot for the AAAI robot challenge. *AAAI Magazine*, 24, 51–72.
- Simonson, I. (2005). Determinants of customers' responses to customized offers: Conceptual framework and research propositions. *Journal of Marketing*, 32-45.
- Smith, A. Bolton, R., & Wagner, J. (1999). A model of customer satisfaction with service encounters involving failure and recovery. *Journal of Marketing Research*, 36, 356-372.
- Spraragen, S., & Chan, C. Service Blueprinting: When Customer Satisfaction Numbers are not enough. In *Proceedings of International DMI Education Conference*.
- Spreng, R, Harrell, G., & Mackoy, R. (1995). Service recovery: Impact on satisfaction and intentions. *Journal of Services Marketing*, 9(1), 15- 23.
- Srinivasa, S. S., Ferguson, D., Helfrich, C. J., Berenson, D., Collet, A., Diankov, R., Gallagher, G., Hollinger, G., Kuffner, J, & Weghe, M. V. (2010). HERB: A home exploring robotic butler. *Autonomous Robots*, 28(1), 5-20.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences*, 23, 645-726.

- Stickdorn, M., & Schneider, J. (2010). *This Is Service Design Thinking: Basics, Tools, Cases*. BIS Publishers.
- Strauss, A. L., & Corbin, J. (1998). *Basics of qualitative research: Techniques and procedures for developing grounded theory* (2nd ed.). Thousand Oaks: Sage.
- Sung, J.Y., Christensen, H. I., & Grinter, R.E. (2009). Robots in the wild: understanding long-term use HRI. In *Proceedings of HRI'09*, 45-52.
- Sung, J.Y., Grinter, R.E., & Christensen, H. I. (2009). "Pimp My Roomba": Designing for personalization. In *Proceedings of CHI'09*, 193-196.
- Sung, J., Guo, L., Grinter, R., & Christensen, H. (2007). "My Roomba is Rambo": Intimate home appliances. In *Proceedings of Ubicomp'07*, 145-162.
- Sunikka, A., & Bragge, J. (2012). Applying text-mining to personalization and customization research literature – Who, what and where?. *Expert Systems with Applications*, 39(11), 10049-10058.
- Sunstein, C. R. (2007). *Republic.com 2.0*. NJ, USA: Princeton University Press.
- Surprenant, C.F., & Solomon, M.R. (1987). Predictability and personalization in the service encounter. *The Journal of Marketing*, 86- 96.
- Tam, K. Y. & Ho, S. Y. (2005). Web personalization as a persuasion strategy: An elaboration likelihood model perspective. *Information Systems Research*, 16(3), 271-291.
- Tam, K. Y., & Ho, S. Y. (2006). Understanding the impact of web personalization on user information processing and decision outcomes. *MIS Quarterly*, 30(4), 865-890.
- Tapus, A. & Mataric, M. (2008). User personality matching with a hands-off robot for post-stroke rehabilitation therapy. *Experimental Robotics*, 165-175.

- Tax, S. Brown, S., & Chandrashekar, M. (1998). Customer evaluations of service complaint experiences: Implications for relationship marketing. *The Journal of Marketing*, 62(2), 60-76.
- Teevan, J., Dumais, S., & Horvitz, E. (2010). Potential for personalization. *ACM TOCHI*, 17(1), 1-31.
- Thomke, S (2003). R&D Comes to Services: Bank of America's Pathbreaking Experiments. *Harvard Business Review*, 81(4), 70-79.
- Thrun, S. (2004). Toward a framework for human-robot interaction. *Human-Computer Interaction*, 19(1-2), 9-24.
- Thrun, S., Schulte, J., & Rosenberg, C. (2000). Interaction with mobile robots in public places. *IEEE Intelligent Systems*, 7-11.
- Torrey, C. (2009). *How robots can help: Communication strategies that improve social outcomes*. Ph.D. Dissertation. Carnegie Mellon University.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 440-463.
- True & Co. <https://trueandco.com>
- Turkle, S. (2005). Relational artifacts/children/elders: The complexities of cyber companions. In *Proceedings of CogSci Workshop on Android Science*, 62-73.
- Turkle, S. (2012). *Alone Together: Why We Expect More from Technology and Less from Each Other*. Basic Books.
- Tuzhilin, A. (2009). Personalization: The state of the art and future directions. In Admavicius, G. & Gupta, A. (Eds.) *Business Computing* (pp. 3-43).
- Vaajakallio, K., Mattelmaki, T., Lehtinen, V., Kantola, V., & Kuikkaniemi, K. (2009). Literature review on service design. *eXtreme Design-project*, University of Art and Design Helsinki & Helsinki University of Technology. Online document.

- Valenzuela, A., Dhar, R., & Zettelmeyer, F. (2009). Contingent response to self-customization procedures: Implications for decision satisfaction and choice. *Journal of Marketing Research*, 46(6), 754-763.
- Vargo, S. L. & Lusch, R. F. (2004). Evolving to a new dominant logic for marketing. *Journal of Marketing*, 68(1), 1-17.
- Vargo, S. L., & Lusch, R. F. (2008). Service-dominant logic: Continuing the evolution. *Journal of the Academy of Marketing Science*, 36, 1-10.
- Vargo, S. L., Maglio, P. P., & Akaka, M. A. (2008). On value and value co-creation: A service systems and service logic perspective. *European Management Journal*, 145-152.
- Venolia, G., Tang, J., Cervantes, R., Bly, S., Robertson, G. G., Lee, B., & Inkpen, K. (2010). Embodied Social Proxy. In *Proceedings of CHI'10*, 1049-1058.
- Verhagen, T., Van Nes, J., Feldberg, F., & Van Dolen, W. (2011). Virtual customer service agents: Using social presence and personalization to shape online service encounters. *Research Memorandum*, 10.
- Verplanken, B. (2006) Beyond frequency: Habit as mental construct. *British Journal of Social Psychology*, 45, 639-656.
- Wada, K. Shibata, T. Saito, T., & Tanie, K. (2002). Robot assisted activity for elderly people and nurses at a day service center. In *Proceedings of ICRA'02*, 1416-1421.
- Walters, M.L., Syrdal, D.S., Dautenhahn, K., Boekhorst, R., & Koay, K.L. (2008). Avoiding the uncanny valley: Robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. *Autonomous Robots*, 24, 159-178.

- Waytz, A., Epley, N., & Cacioppo, J. T. (2010). Social Cognition Unbound: Insights Into Anthropomorphism and Dehumanization. *Current Directions in Psychological Science*, 19(1), 58-62.
- Weick, K. E. (1979). *The Social Psychology of Organizing* (2nd ed.). Addison-Wesley.
- Weick, K.E. (1995). *Sensemaking in Organizations*. Sage, Thousand Oaks, CA.
- Weld, D. S., Adar, E., Chilton, L., Hoffmann, R., & Horvitz, E. (2012, July). Personalized Online Education—A Crowdsourcing Challenge. In *Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence*.
- Weld, D. S., Anderson, C., Domingos, P., Etzioni, O., Gajos, K., Lau, T., & Wolfman, S. (2003, August). Automatically personalizing user interfaces. In *Proceedings of International Joint Conference on Artificial Intelligence*, 18, 1613-1619.
- White, R. W., & Kelly, D. (2006). A study on the effects of personalization and task information on implicit feedback performance. In *Proceedings of the 15th ACM international conference on Information and knowledge management*, 297-306.
- Woods, S. N., Walter, M. L., Koay, K.L., & Dautenhahn, K. (2006). Comparing human robot interaction scenarios using live and video based methods: towards a novel methodological approach. In *Proceedings of Workshop on Advanced Motion Control*, 750-755.
- Wunderlich, N. V., Wangenheim, F. V., & Bitner, M. J. (2013). High tech and high touch: A framework for understanding user attitudes and behaviors related to smart interactive services. *Journal of Service Research*, 16(1), 3-20.
- Yamaoka, F., Kanda, T., Ishiguro, H., & Hagita, N. (2006). How contingent should a communication robot be? In *Proceedings of HRI'06*, 313- 320.

Yamazaki, A., Yamazaki, K., Kuno, Y., Burdelski, M., Kawashima, M., & Kuzuoka, H. (2008). Precision timing in human-robot interaction: Coordination of head movement and utterance. In *Proceedings of CHI'08*, 131- 140.

Yee, N., Bailenson, J., & Ducheneaut, N. (2009). The proteus effect: Implications of transformed digital self-representation on online and offline behavior. *Communication Research*, 36(2), 285-312.

Yee, N., Bailenson, J., & Rickertsen, K. (2007). A meta-analysis of the impact of the inclusion and realism of human-like faces on user experiences in interfaces. In *Proceedings of CHI'07*, 1-10.

Zomerdijk, L. G., & Voss, C. A. (2010). Service design for experience-centric services. *Journal of Service Research*, 13(1), 67-82.