Designing Interpersonal Intelligence and Ownership Models for Social Agents

Michal Luria

A thesis presented for the degree of
Doctor of Philosophy

Thesis Committee:
Jodi Forlizzi (Co-Chair), Carnegie Mellon University
John Zimmerman (Co-Chair), Carnegie Mellon University
Geoff Kaufman, Carnegie Mellon University
Malte F. Jung, Cornell University

Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, USA
February 2022
CMU-HCII-21-108

This work was supported by the National Science Foundation Grant SES-1734456
and NASA Grant 80NSSC19K1133

Abstract

This doctoral thesis investigates how intelligent personal agents and social robots should be designed to behave and interact in social environments. Fifty million Americans now own smart speakers, and over 40% use chatbots regularly. These agents are gaining access to people's personal information, and they need increasingly sophisticated rules on how to behave and how to both share and protect personal information. Yet, at the moment, they are designed as one-on-one devices (one agent and one user), whereas in reality, they exist in socially complex spaces. This body of work uses design research approaches to examine how designers might break through current underlying assumptions of agent and robot design, map a broader design space for future personal agents and robots, and suggest considerations and guidelines for more sophisticated, transparent, and trustworthy social agents. One aspect of agent design that was revealed in this work was that of ownership. A sense of ownership over artifacts provides individuals with a sense of control, trust, and comfort. It is not clear, in current designs, who an agent belongs to and whether and how agents create a sense of ownership for their users. Do agents belong to one individual or to a group? Do they belong to the person who uses them, or to the company that provides them? The second part of this thesis examines design opportunities within this space, and suggests how different ownership models might impact agent perceptions and interaction with them.

Acknowledgments

I could never have imagined so many people would take an active part in my path towards successfully completing a Ph.D. First and foremost, I would like to thank my advisors, Jodi Forlizzi and John Zimmerman, for their ongoing support and dedication throughout this time, and for always reassuring me that I am on the right track, even when at times it did not seem so. Jodi has supported me from my very first day at CMU and was always available to share a word of wisdom, question a research direction, or provide a shoulder to cry on. John has been an infinite source of knowledge, always bringing a fascinating and original view on the world of designers, researchers, and HCI as a whole. He also taught me how to stand up for my ideas and to push back and argue when needed.

I would also like to thank my committee members, Malte F. Jung and Geoff Kaufman. Over the years, I had the wonderful opportunity to engage with them in conversation, in class and outside of it. I thank them for supporting my work and providing valuable feedback, as well as a fresh and interesting perspective on my thesis.

I was lucky enough to collaborate with many more individuals during my Ph.D. and to learn something from each and every one of them. Thank you to Dan Lockton, Stuart Candy, Molly Wright Steenson, Oren Zuckerman, Aaron Steinfeld, Jessica Hammer, Jessica Hodgins, Liz Carter, Amit Zoran, Henny Admoni, and Golan Levin for providing unique perspectives and for being mentors along the way.

This time would not have been the same without the support of my peers, who also became my friends. Yasmine Kotturi has been my person at HCII since the open house event, where we both decided to go to CMU and spend our first year as roommates. Thank you, Yasmine, for being my first friend in Pittsburgh and for your support and friendship throughout. Thank you to Sam Reig for being such a great collaborator, HRI-peer, friend, and just someone who I could always rely on. Thank you to Nur Yildirim for getting the inner desires and worries of a designer, for your sincere empathy, and for hosting such a wonderful troast. Thanks to Joselyn McDonald for being the person who always got me, research-wise, interests-wise, and just for sprinkling fun into my life. Thanks to Jonathan Dino for many evenings and afternoons at Pittsburgh's best location—The Row House. Thanks to Rushil Khurana, Sai Swaminathan, and Kristin Williams for being my office roommates and venting partners. There are so many others that I had the chance to work with or to spend time

together. I thank you all for preserving my sanity through social and professional encounters throughout my Ph.D: Joseph Seering, Stephanie Valencia-Valencia, Julia Cambre, Zhi Tan, Reuben Aronson, Rachel Gita Karp, Judeth Oden Choi, Qian Yang, Steven Dang, Tianying Chen, Judith Uchidiuno, Mary-Beth Kery, Ken Holstein, Anna Kasunic, David Gerritsen, Fannie Liu, Haojian Jin, Zheng Yao, Kirabo Byabashaija, Cori Faklaris Lea Albaugh and many others.

Thank you to Queenie Kravitz for many joyful hallway conversations, and for always providing support at difficult times. Thanks to Marian D'Amico for the very much needed administrative support for doing research and completing a Ph.D. at CMU.

Most importantly, I would like to thank my partner, Guy Hoffman, for supporting me every single day. Countless tears and laughs, thoughts, fears, random ideas, and lots of stress that Guy helped me navigate in an effort to move forward with renewed strength. I truly could not have done it without you. Thank you to my son, Geffen Luria Hoffman, who only came along in the last ten months, but feels as if he was always part of my journey. I would also like to thank my friends back in Israel, as well as my parents, siblings, and extended family, for their ongoing support regardless of the distance.

Contents

1 INTRODUCTION

This body of work aims to investigate how intelligent personal agents and social robots can and should behave; how they can have more social sophistication to effectively navigate interpersonal relationships and function in socially complex contexts. Today, more than 50 million Americans own smart speakers, and over 40% have recently used chatbots [106]. Robots are appearing in workplaces, and companies are attempting to design robots that would remotely work with people in their homes. Overall, more and more people are interacting with agents that appear social, or that use social cues to interact [103]. However, these agents are very limited in understanding or responding to social nuance.

As agents gain access to people's personal information through a range of digital activities and applications, they need increasingly sophisticated rules on how to behave, how to both share and protect personal information, and how to infer aspects of the social context. Today, these devices mostly take the form of a single agent with access to a single user's information. But what happens when an agent or even a set of agents gain access to a family's full set of personal data? What should an agent do when a mother-in-law asks for their daughter-in-law's location or when a teen asks an agent to lie and tell their parents that they have been studying? How can an agent bring both value and harmony within a group of people? Is the agent a product that belongs to the user, or is it a service that is owned by the service provider?

While the definition of agents and robots strongly varies, I focus on interfaces that interact with people using *social* cues—whether it be voice, text, or nonverbal behavior. In the field of human-robot interaction (HRI) and research on social agents more broadly, the implicit convention is to design robots and agents as human-like. Researchers are refining robotic arms to understand and act in the physical world as humans do, chatbots attempt to address inquiries just like a human agent, and companies like Amazon aim towards conversational agents that fully understand and

respond to verbal communication. One downside of this convention is that it unintentionally frames agents and robots as "less than human" interfaces. They attempt to act as close to humans as possible but fail to do so entirely. An alternative approach is to view agents as *different-abled* interfaces that can give value to their users by engaging in *non-human* capabilities and social behavior. For example, agents can move their "social presence", or "entity" from one body to another; they can be in many places at the same time; they can focus on multiple tasks simultaneously without sacrificing their more-than-human precision; they can develop nonverbal behavior that is unique to robots and does not strive to be human-like.

With that in mind, my work examines what it might mean for AI agents to be socially sophisticated and aware. However, this opportunity space is too broad and cannot be easily or systematically covered. In contrast to other improvements in technology, which have a clear goal of what "better" should be (faster, cheaper, smaller), in the case of socially-aware agents, it is unclear what would constitute a "better" interface. Should agents fully use every piece of information to understand social context? Should they attempt to be a "moral compass" and call out users who misbehave? Or should they primarily be obedient?

Given these and many other unanswered questions, it is too soon to take an engineering research approach. Design research methods are more suitable for exploring this opportunity space to understand the goals for socially complex agents and how they can provide the most value to users. Therefore, this body of work uses design research to learn about the under-explored design space of complex interaction instances with agents in social environments. The findings of my work can play a critical role in identifying key agent capabilities for the development of socially sophisticated agents, and provide an understanding of what "better" social agents should do.

In the first part of my thesis, I conducted three research studies that used exploratory design methods to examine socially complex personal [87], interpersonal [116] and group [88] interactions. Findings indicated that people expected agents to be "better" social communicators by nuanced use of their data to support them in the acts of daily living. For example, participants wanted agents to use their sensing and learning capabilities to improve their understanding of social roles, generate continuous interactions with users over time and to provide highly personalized services [87, 116, 88], but also to have transparency about which data is being

collected, and to avoid agents that can "listen", "watch" or "record" [88]. Participants indicated that a single agent "entity" that can move from body to body and access their personal information would be preferable than many agents across many services [87, 116].

The finding of preferring a single entity extended beyond merely personalized preferences. Participants expressed a desire to know who *the owner of an agent is* and, as a result, who it is accountable to. They indicated that this knowledge would allow them to better anticipate AI behaviors in complex social situations, for instance, in a case of conflicting requests from different users, or when one user asks to gain access to another's data [88]. The desire to know more about who "owns the agent" raises new research questions regarding the influence and importance of ownership in interaction with agents, how agents account for ownership and affiliation in their design, and whether they should be conceptualized as products or as services.

Thus, the second part of my thesis focuses solely on the topic of agent ownership and attachment. I attempt to form an understanding of the overall concept of agent ownership and potential value, as well as to lay out the design space of possible future social agents that take ownership into consideration. This final set of design recommendations can contribute to a sense of attachment, trust, and an overall improved experience with socially sophisticated agents.

Research Overview

The research in this thesis consists of several overarching questions. In the first part I look into how agents should behave and make use of data as they are integrated in socially complex contexts. In the second part, the research focuses on when and how should agents communicate ownership to users, and how might service providers address the perceived affiliation of agents with them while maintaining users' sense of ownership.

PART 1: AGENTS IN SOCIALLY COMPLEX CONTEXTS

Study 1: Personal interactions

How should agents be designed to address a *single user* in a range of social contexts? Should different agents serve across services in homes and public spaces at different touchpoints? Or can and should agents utilize

their non-humanlike qualities of being able to move an agent's "entity" or "social presence" from one location to another to facilitate sequential and personal interactions across space?

This work used Speed Dating with User Enactments [34] in which participants interacted with agents that had one social presence per body, that could *re-embody* (move their social presence from body to body), and that could *co-embody* (move their social presence into a body that already contains another). Reactions showed that participants felt comfortable with familiar and re-embodying agents, who created more seamless and efficient experiences. Situations that required expertise or concentration on the task raised concerns about a single "can do it all" agent. We report on our insights regarding collaboration and coordination with several agents in multi-step interactions [87].

Study 2: Interpersonal interactions

How should agents be designed to address *interpersonal* interactions in public contexts? Should there be a single agent that addresses several users, should services provide personal agents that are affiliated with the service, or should each user have their personal agent that "follows" them from one location to another as suggested in Study 1?

This study made use of a more structured format of Speed Dating with User Enactments, as the research questions were more specific and structured and could be more rigorously tested. Participants interacted in situations where agents either belonged to the service (like current standard service agents), agents belonged to the service but generated a unique agent for a specific user, or agents that belonged to *the user*, and served them in public and in their personal space while leveraging personal information. Findings suggested that people prefer the latter and felt comforted by a single entity that "knew them" and is capable of using their personal data in a range of contexts without having to transfer it across systems [116].

Study 3: Interpersonal interactions with a shared agent

Studies 1 and 2 began to outline how agents should behave in public scenarios. But when moving agents into personal spaces new challenges surface. Current "home agents", such as Google Home and Amazon Echo, assume their devices are shared between household members. Yet their design does not address the interpersonal challenges that agents

might need to address as they become more socially sophisticated. For example, what should an agent do when a mother-in-law in the home asks for her daughter-in-law's location? Should the agent share this personal information? Should it prevaricate, stall, or redirect the subject? Should it snarkily refuse? What should an agent do if a teen asks it to lie and tell parents the teen has been studying? Should it keep secrets? Should it actively deceive? Or should it tattle?

The findings of this study were the first in which participants explicitly stated that they would like to know who the agent *belongs to*. Participants argued that this kind of knowledge, knowing who the agent belongs to and is accountable to, would assist in setting their expectations about how an agent is likely to handle a range of interpersonal situations. Furthermore, findings suggested that the *social roles* and *presence of individuals* are critical pieces of information for the agent to interact and respond in a more nuanced and socially appropriate way.

Part 1 Summary

The three studies combined successfully began to map the initial design space of socially sophisticated agents. One of the important outcomes of this work was the identification of *agent ownership and affiliation* as an important aspect to consider when designing agents, a topic that has not been previously explored in the context of social interfaces. While findings in the last study explicitly expressed the importance of user ownership and service affiliation as part of agent design, support for its importance can also be found in the first two studies, as I will further discuss in the next thesis chapters.

PART 2: AGENT OWNERSHIP AND AFFILIATION

Study 4: Ownership and affiliation of agent (and) services

In order to tackle the complex topic of ownership in the context of agents, first, it was critical to consider the concept of *service* ownership as a whole. By definition, ownership and service contradict each other. Yet some aspects of ownership might still be evident in services, as my prior work has shown, particularly in the case of agent services.

On the one hand, in the current structure of social agent services (e.g., Amazon Echo, Google Home) and their users, service providers are the ones who *own* agents to some extent, as the service providers control the

content and interaction with the agent, and instead provide users with a *service*. On the other hand, people purchase devices that "include" these services and therefore perceive these devices as *theirs*.

While we know little about the tension between service and product specifically with social agents, this co-existence is not new—recent years have shown many examples of "servitization" of products and "productization" of services [7]. Therefore, in this research effort, I asked people to reflect and compare social agent services that they use to range of to other products and services, technological and non-technological. Some of the research questions this effort addresses include: (1) To what extent do current users feel ownership over their agents, in comparison to other material possessions and technological devices? (2) What role does the service model play in people's sense of ownership over their agent? and (3) What are some alternative models for service agent design? This study used a combination of interviews, card sorting, and sketching activities better to understand current stances on agent service ownership and affiliation. I conclude with opportunities for changes in service design that can address users' concerns and potentially improve their sense of ownership and trust in personal agents.

Study 4: Exploration of community-owned agents

This part of the thesis focuses on the possibility of having *community-owned* agents. It investigated social agents in the form of a chatbot designed to be shared within an online community to improve moderation and encourage new and more positive interactions between community members. As part of this goal, the project attempted to design a sense of collective ownership within that group. We created a concept "baby" chatbot that learned and "grew" within a community and implemented it in a game-streaming online community on Twitch.

Findings from a three-week field study showed that the implemented chatbot was successful in creating a sense of collective ownership over it. We discuss insights about three design choices for the bot that contributed to its perceived ownership: (1) enabling users to creatively shape it, (2) promoting a sense of individual responsibility, and (3) creating opportunities for personal interactions.

Study 5: Exploration of personally-owned agents

Finally, the last exploration in my thesis looks into *personally-owned* agents. For this goal, the study used co-design with theatre professionals to devise an immersive performance, "Robotic Futures," as a form of knowing-through-doing. The performance allowed learning about the felt-experience of interacting with personally-owned agents in the ecosystem of the home.

Examining the nuance of interactions that unfolded throughout the devising process raised guidelines and questions about how personally-owned agents should behave and when they might be designed to be shared as opposed to personally-owned. Findings suggested three types of agents in the home: "social" agents, "expert" agents and "self" agents. The latter two were found to be more suitable as personally-owned, as opposed to shared between family members. Furthermore, the study identified a tension between who the primary *user* is, and whether they are also the *owner* of the agent. The ability to *use* an agent but not *own* it, and vice-versa, can create a more complicated and nuanced interaction with agents, and perhaps more nuanced behavior on the agents' side.

Thesis Contribution

The contribution of this body of work is in two resolutions: from a broader perspective, my work aims to begin to map out the design space of personalized agents, that understand social context and roles and that are better socially integrated; agents that make use of their super-human capabilities, rather than striving to be human-like; agents that collect and use information ethically and that align with their users' boundaries and concerns. Identifying this design space allows to fully examine interaction challenges and concerns before the technology is ripe. Furthermore, it reveals important aspects that developers and designers should aim for as agents transition towards becoming more socially sophisticated.

The second part of my work delves deeper into one area within the design space that has emerged in the first part: ownership and affiliation of social agents. The knowledge generated in this second part include a better understanding of where agents fit within the overview of "things" people own, as well as guidelines for how agents can be designed as personally- or community-owned agents: when is each appropriate for the social context at hand, and how are they likely to be perceived by users.

2 RELATED WORK AND THEORY

This chapter draws on three bodies of research and their intersection: (1) Research about social agents in personal environments and social contexts; (2) Theory and prior work about ownership of material and technological possessions; (3) Service theory related to product versus service mindsets and technological applications. The intersection of the three, ownership and possession of service agents, is a topic that has been under-explored in research, and is the focus of this thesis.

Agents and Robots in Social Contexts

Agents in the Home

When designing technology for the home, designers need to address a range of unique challenges, as homes are complex, emotionally-oriented social spaces shared by multiple people with many different roles, genders, and ages [30]. To better understand this complex space, previous literature has attempted to define the challenges that should be addressed to encourage acceptance of technology within the home [45]. Research has also examined what makes smart homes and social agents in the home desirable and how they should be designed to be accepted as part of the domestic space [35, 151, 129]. "Placemail", a task management software based on location, looked into who in the household should be able to receive reminders and information [84]. Brown *et al.*'s work on the Whereabouts Clock explored how publicly displayed information in the home could remain private through the use of ambiguity, where family members use knowledge of each other's routines to extract hidden and pertinent information [18]. Pina and colleagues examined opportunities to design family-centered health-tracking devices as opposed to ones designed for isolated individuals [112]. Pierce used a design-led inquiry to explore the boundaries of what may be perceived as "creepy" for future smart

home technology [111]. These examples indicate a complex landscape of privacy and personal information when using technology in a social environment. However, interactions that consider the complexity of social aspects embedded in an environment have been under-explored in the context of social agents.

The topic of interaction with technology in socially-complex situations is critical to consider with social agents, as these introduce another layer of socialness beyond those of other technological devices. In contrast to their simple services, agents' use of speech causes people to perceive them as social [103], and to interact with them similarly to how people interact with other people [102, 114]. This gets reinforced by design choices that make agents' performance come across as humorous and intelligent, which also drives increased agent personification [85]. Situating interactions in social spaces, such as people's homes, also increases personification [113]; people who interacted with Alexa together with other family members were more likely to refer to Alexa in a personified way, specifically in the context of a household or in the presence of children.

According to industry reports, as agents develop and become even more widespread in people's homes, they will also become more personal and social [1, 118]. Research supports this prediction, as it shows that matching agent personality [97], preferences [11], humor type and interaction style [13] can all benefit interaction and engagement, and shape a positive attitude towards social agents. Agents can also increase their perceived socialness by gaining awareness of their environments, recognizing emotions, being trusted by users, and demonstrating unpredictability [37]. Rogers and colleagues have examined people's reactions to the notion of future agents as personal companions from a more critical perspective [122]. Such speculative scenarios may not be very distant in the future; many agents are already designed to have a personality and to give a sense of being capable of forming relationships with their users [153].

Yet having social cues and personalizing agents to a single user is not sufficient to develop social sophistication. Like other devices, more research is needed to understand the nuance of interaction in the home, the needs and concerns of a range of stakeholders as they interact with each other and with the technology. Additional research will allow researchers and designers to create agents that behave, interact, collect, and share data in a responsible and desired way, increasing their perceived social sophistication and trust by users.

Agents as Mediators

Several research projects have examined how agent behavior may influence interaction within a group of users. Robots have been successful in mediating turn-taking using solely gaze [100]. They have also been able to mitigate conflict between couples [60], within collaborating teams [69], and between children [135]. In one study, commentary given by a robot reduced tension and increased engagement between team members [139]. Previous work has also designed robots to encourage less active participants in a group to be more involved in conversation [121, 144], and found that encouraging behavior from the robot increased positive interpersonal evaluation of team members [121], general group engagement, and overall group problem-solving skills [144].

This previous work suggests that robots and agents can successfully mediate between people, encourage more positive interaction, and even mitigate conflict. Yet the primary question that is left unanswered is not whether agents *can* be successfully socially involved—but if they *should* be. Additional research is needed to understand how agents fit into the natural ecosystem of the home and its daily routines and social norms. For example, we know that agents can mitigate social conflict. But if designers implement this kind of behavior in the home, they will need to resolve many more questions: Should agents intervene every time they sense conflict in the home? Should they always be neutral, even if one side is clearly behaving inappropriately? Should they intervene in the same manner whether the conflict is between two children, two adults, or a child and an adult?

My work aims to address some of these fundamental questions about how social agents might fit into complex social environments by looking at people's current set of values, mores, and needs, with the goal of generating an initial set of topics that should be considered within this design space.

Material and Digital Possessions

This section sets the theoretical background and related work on ownership and personal possessions. Due to the social qualities of agents, in addition to theory about material and digital possessions, I include the sense of possession over social beings as well.

Theories on material possessions draw a line between possessions for *individuation* and possessions for *integration*. Possessions for individuation

are primary there for the purpose of highlighting people's individuality, and uniqueness *from* others [32], and possessions for integration are intended to express people's affiliation with a group, or how they fit in *with* others [128, 32], and possessions for individuation. The importance of each varies by culture but is common in most and for all stages of life [73].

Material Possessions as Individuation

In contrast to previous theories that describe possessions as things that define relationships between people, Csikszentmihalyi has used ethnographic observations to better understand the intellectual and emotional value that artifacts can encompass for *individuals*. In a study of people's possessions in their home, he found that every home has objects with symbolic meaning that stem from being personally and emotionally involved with it. The symbolic power of an object "produces a sense of order in mind" by connecting the past and present, representing people's loved ones, or expressing one's goals and desires [31]. Additional work supports the notion that possessions are critical in shaping our identity and perception of self [62, 119], and that they even have a role in influencing the way we think [147].

Kleine goes further to describe the role possession have over time, and suggests a distinction between possessions for *identity change* and possessions for *stability*. Some possessions embody past "selves" that are kept and are therefore carried into the present, and some possessions represent past "selves" that have changed over time and that are let go of through discarding objects. In other words, keeping possessions helps us reflect who we are today, and dispossessing them communicates aspects of ourselves that we have decided to leave behind [75].

My work builds on three prominent theories of possessions as individuation: Mere Ownership Theory (Endowment Effect Theory) [70], Self Extension [10], and the Possession Attachment Theory [75].

Mere Ownership Theory

Once an individual possesses an object, the Mere Ownership (Endowment Effect) Theory is applied. This theory argues that people may value material possessions merely due to the fact that they own them [71]. Participants in a study were given either a mug or a chocolate bar in return for answering a short questionnaire. Later, they were offered to exchange their compensation (mug/chocolate bar) for the other. Participants that

were initially given the mug perceived the mug with more value, and participants that initially received the chocolate bar perceived the chocolate bar with more value. Thus, the theory derived from this study suggests that people's sense of ownership towards an object is sufficient for them to perceive it with higher value [70]. This endowment effect was found not only to instantly occur as soon as an item was acquired, but also to increase over time [138].

Self Extension

Another highly influential work on the relationship between the self and objects is Belk's essay that defined possessions as *Extensions of the Self* [10]. Belk suggested that objects are not simply owned by individuals, but can be a critical component in people's identity perceptions. The theory is mostly focused on material possessions yet also includes people, places, and groups as potential extensions of the self.

Self-extension, according to Belk, is relevant in all stages of life. A person is most likely to extend their sense of self onto an object in one of three situations (Sartre [123] in Belk): (1) When one has **control** over an object (including object destruction or gift-giving); (2) When one **creates** an object or idea; or (3) when one has intimate **knowledge** about the object, place or thing [10]. These three situations were later experimentally tested and confirmed [74].

Groom et al. found support for Belk's theory that creation supports self-extension in the context of robots too. In their study, people were more likely to experience self-extension if they built a robot themselves [52]. Participants who self extended thought they had more overlap with the robot, felt more attached to it, and saw it as a team member rather than a competitor. Similar to the findings in Keisler and Kiesler [74], Groom et al. compared between a humanoid and a car-like robot. They too found that people showed greater signs of self-extension in the car condition than in the humanoid robot condition [52]. Both studies suggest that less anthropomorphic shapes elicit more self extension.

Possession Attachment Theory

Other researchers have attempted to define the importance of possessions through the Possession Attachment Theory: Possessions that people are "attached" to are possessions that take part in the narrative of people's

lives—who they were, who they are now, and who they anticipate becoming [75]. A material possession may reflect on one's perceived personality or on their desirable connections with others. Attachment possessions are frequently more emotionally charged, project a positive valence, and are held physically closer than objects that people are not attached to [10, 44].

According to Schultz Klaine et al., attachment is not a property of the material object itself, nor is it a property of an individual; it is a property of the relationship between the two [128]. According to the theory, attachment is (1) subjectively perceived by the possessing individual towards a specific object; (2) is not a deliberate action, but arises over time and through interaction with an object; (3) has some level of strength (strong versus weak attachment).

Zimmerman has explored how Possession Attachment Theory can be applied in experience design practice by examining a range of designs through this theoretical lens. Some of the ways that the theory is applied to *design for the self* include allowing people to focus on a single role that represents their aspired self, reminding them of people and affiliations that are important to them [154]. Research that has examined *digital* possessions that people valued and expressed attachment towards found that people attempt to turn such possessions into *physical* ones, like printing out an image [109]. Other work suggested that people found new ways of making meaning of digital possessions, like leveraging their ability to access these possessions from anywhere and using a range of devices [108].

Material Possessions as Integration

Research has found that possessions have a significant role in serving interpersonal interactions and relationships with others [50]. Dant argues that all objects are social entities, in that they extend human action and mediate meaning between people [33]. Objects can shape interpersonal ties and mediate people's messages over time, space, or both. Some objects enable direct communication, like phones or computers. Others are used for indirect communication—for instance; visual art allows a message to be transferred from the creator to the viewer over time and space [33]. Objects can also mediate between people through sentimental value by bringing memories and mutual experiences to the foreground of our attention—the majority of people own and cherish at least one object that reminds them of a particular person or community [32].

Mutual Possession, or Sharing

Integration value does not only appear when possessions are passed from one individual to another, but also when they are shared. Research has indicated that the mere mutual possession of an object, when something is "ours" as opposed to "mine" or "yours," can have meaning and value [9].

Sharing is most common within people's homes. Olsen conducted ethnographic research in couple's homes and found that possessions that people are attached to as a couple are similar to ones that are individually possessed, but their meaning stems from the representation of the relationship as opposed to a representation of the self [110]. Families, similar to close partners, share most of their possessions, including their home, their car, their meals. Sharing within the family is based on love and trust, and is characterized with occasionally being non-reciprocal or not equal. Rather, sharing is frequently based on availability of resources and according to different needs within the family [9]. Just like possessions can tell the "life story" of an individual [75], they can also tell the life story of an entire family.

In contrast to within-home sharing, there is a clear distinction between joint use and joint ownership for adults outside the family [9]. *Joint ownership* with a friend is perceived as more risky as paths are more likely to diverge, making it unclear who should take shared possessions [9].

These theories of possessing and sharing objects are yet to be explored in the context of social agents: Should anyone in the home be able to access and use an agent? Extended family members, friends, neighbors? How would sharing an agent be different between partners in contrast to sharing within the broader family circle? My research aims to begin to indicate some of the answers and guidelines that would assist in designing agents that are used in socially complex contexts.

Technological Possessions

Dant claimed that computers are merely a natural development of material objects. They are more complicated in how they work, but they still support human interaction, and more importantly, "never grasp meaning or discover value" [33]. Like other possessions, technological possessions can also shape how people perceive themselves. For example, a study found that perceived phone usage can affect self-perception [63]. Using a Research Through Design approach [155], and with a possession attachment theoretical perspective, Zimmerman suggested two aspects to focus

on when designing technology for the self; *Role Enhancement*, designing things that help people move towards their ideal self and *Role Transition*, designing products that support people's discovering and inventing of themselves in a new role [154].

Brush and Inkpen identified two common models by which technological devices for domestic settings are designed today: an "appliance model" and a "profile model". The "appliance model" implies anyone in the home can use the device, relying on social protocols to mediate sharing. However, this model allows for little personalization or privacy. The "profile model" supports multiple users by asking them to self identify and thus reduce sharing problems through individual ownership of devices [19].

Odom et al.'s work focused on *virtual possessions*, and looked into how they differ from material ones; digital possessions are *placeless*, *spaceless* and *formless*. They are *placeless* in that they can be accessed from anywhere, at any time, and simultaneously; they are *spaceless* because they do not take up physical space, therefore making it difficult to assess collection sizes and organize them; they are *formless*, as they can be replicated and remixed, making it difficult to tell the 'original' from copies or modified possessions. While the theory can be applied to some extent to agents, it also becomes more complex given that agents have a physical representation in space, but their "mind" and the service they provide remains virtual.

Like other devices, the current behavior design of agents does not explicitly support the multi-user nature of their use. Some research has looked into agents and robots for shared activities, such as group classroom work [51] and professional meetings [64]. Previous work has also looked at using an agent for specific social roles, as discussed in this chapter, but these agents are not explicitly designed to be shared or individually owned. This gap indicates an opportunity to design agents to be explicitly shared or personally-owned and better understand people's expectations of agent behaviors and norms for each of the two.

Social Presence

Today, most agents and robots are designed to be social in their behavior and interaction. Yet, as previously discussed, they do not live up to the "full-socialness" of humans. This section looks into social presence theories and how social relationships inform and impact individuals' sense of self.

Expansion of Self

Self Expansion is a theory that can be applied to social (but not material) possessions. The theory argues that people seek to expand their perception of self by taking on the perspectives of others as a way to increase their resources and identities [4, 3]. According to prior work and theory, relationships that expand the self provide individuals with new experiences and perspectives, and allow them to learn new things on their way to become a "better version of themselves."

While most of the research on this topic is primarily relevant to close-partner relationships, self-expansion can also occur through hobbies and activities that may include friends and other acquaintances that expose people to new points of view. Self-expansion can even occur with people an individual does not know personally, such as with personal heroes, celebrities [141], and even fictional characters [134].

Parasocial Relationships

Parasocial Relationship Theory extends this point and attempts to describe people's relationships with media characters and celebrities and how these relationships are similar or different from typical social relationships. Media characters provide a continuous persona from week to week, a personality that the observer may relate to and count on them "being there" [61]. Yet these characters are "para"-social, only close to social, because this relationship is one-directional and controlled by the performer, without an opportunity for reciprocity. The observers, on their side, have control over whether they stay involved in the relationship, and can easily opt out of it at any time.

Research has shown that parasocial relationships can provide a sense of belonging and reduce loneliness [38] while maintaining a reduced risk of rejection [61]. Such relationships can also give observers opportunities to experience a broad range of encounters and perspectives that they might not have access to in their immediate environment [134], and even experience a shift towards their "ideal self" [39].

While the original theory was developed based on mass-media interaction such as radio and television, the Internet, social media, avatars, and agents introduce theoretical changes. Klimmt et al. showed that people can develop para-social relationships with digital characters, even though they are not performed by real people [72]. Avatars in the online sphere raise a debate on whether the connection with them is a type of para-social

relationship [29], or rather a form of "self-love", as the avatar is performed and controlled by the player themselves [55]. A lab research study with the dog-like robot AIBO suggests that people's tendency to form parasocial relationships can influence their attraction towards robots too [77].

Implicit Egotism and Similarity Attraction

One way in which digital agents are different from media characters is that they can be personalized for the user and be designed to look like, behave, and even reflect similar values as the user. Two theories support why this might be beneficial for agents: Implicit Egotism and Similarity Attraction. *Implicit Egotism* argues that people are positively biased towards others who merely *seem* similar to them, for example, have similar names or share a birth date [68]. A study of virtual agents found that people find avatars that resemble their own face more credible than other agents. However, they were also more persuasive only when the experimenter, and not the participant, were the one who created the agent to be similar to them [80]. Another study found that agents that were represented with participants' faces were treated more positively by participants than other agents [6].

Similarity Attraction suggests that people are more attracted to others who are similar to them in personality, values, and beliefs (unlike the Implicit Egotism theory that relies on more superficial resemblance) [23]. Research has shown that this theory applies to virtual agents and robots as well. Examples include a speech user interface that was perceived as more attractive, credible, and informative when they matched the user's personality type (introvert vs. extrovert) [101], and robots that were perceived as more friendly if they had the same set of personal preferences as their user [11].

Ownership of Services

Few research has looked into the sense of ownership and the perception of possessions when the product becomes a service, as is quite common within the technology sector. Gruning found that people feel a lessened sense of ownership over Kindle e-books that are stored within their Amazon account than e-books that are stored elsewhere [53]. This finding could be explained by Ligon and colleagues, who suggest that the lack of control over their usage, posed by technology companies, reduces the sense of ownership [81]. For example, participants cannot freely share

files, transfer them to new devices or use them on multiple. This creates a disconnect between the inferred *psychological* ownership of digital products and the actual *legal* ownership that is very limited.

How should designers address this gap when designing products that have service embedded in them? The final part of my work attempts to address this gap and to suggest ways in which services can provide a sense of ownership, and as a results trust and control, over technological devices that combine product and service.

Ownership of Agent Services

While there has been some work that looks into self-expansion and self extension onto virtual avatar characters as detailed above, I do not know of any work that primarily looks at ownership and possession of physically embodied social agents and robots. Robots and agent introduce new challenges related to ownership, and require more knowledge for designers to implement. This gap, that was identified in the first part of my thesis, is addressed in the second part.

Part I

Agents in Socially Complex Contexts

3 RE-EMBODYING AGENTS IN PERSONAL INTERACTIONS

3.1 Overview

This work begins to explore the space of socially sophisticated agents by examining how to design agents and robots for a *single user* in a range of social contexts. Should the same "entity", or "social presence", accompany the user across service touchpoints? Should there be many different agents that interact with a single user? This exploratory research sets out to examine opportunities for designing agents that accommodate more than a single-instance and single user interactions, but rather can move their social presence from one robot body to another [87].

The study was motivated by service theory showing customers prefer to interact with a single person within a service. As argued in the thesis introduction, modeling human-agent relationships after human-human ones sets high expectations for agents' capabilities that are not always met. Alternatively, agents can possess non-human traits, such as the ability to move their social presence from one body to another, which could enrich the interaction and change how agents are perceived. For example, if a person uses Alexa in their home, the same social presence can show up in their car and interact with them there. An agent's social presence might also appear in two bodies in two different places at the same time; a single presence might be simultaneously active in a user's home and within their workplace, taking care of two unrelated tasks. One body might even hold two distinctly different social presences at the same time.

We examine such *superhuman* agent capabilities to move their presence from one body to another as a way of leveraging their non-humanness to follow individual users across multiple touch points within a service. Should a hotel have a single social presence, a sort of "digital brand ambassador"? One that takes the user's reservation over the phone takes their bags at the curb, checks them in at the front desk, answers questions from within the room, and takes their order at the bar? If the user interacts

with an agent in their hotel room, then have they shared their medical history with the hotel, or should the hotel's agent just know less about the user and offer less personalized advice? Designers currently lack patterns to guide their choices surrounding these questions, and society has not yet developed social mores that inform this kind of human-agent interaction.

Previous work has explored the notion of multiple co-existing social presences. Chaves and Gerosa compared interaction with a single omnipotent chatbot to interaction with several "expert" chatbots in a single system. They found that participants perceived interaction with many social presences as confusing, and suggested designers choose a single social presence for this kind of task [27]. Some work has explored social presences that migrate across physical platforms according to the context of interaction—the Agent Chameleon Project looked at the concept of artificial intelligences that are not constrained to a single body [42]. This research sought to outline a technological architecture that would enable social presences to migrate across different physical and virtual spaces. Martin et al. focused on virtual avatars and evaluation of how the idea of migration can be communicated to users [91].

This work extends previous findings by engaging people in *critical reflections* about the contexts in which they might or might not want agents to take on this super-human ability. Instead of attempting to design shifting social presences correctly, In this study, we set out to gain some perspective on this large, ambiguous, and complicated solution space. By immersing participants in scenarios that explore social presences that move across a combination of physical platforms (robots, cars, etc.) in a range environments, "real-world" context is added to agent migration.

As a first step, we investigated how a single user might interact with a set of agents to complete one or several related tasks. We used Speed Dating with User Enactments (UEs) [34] with prototypes, scenery, and "Wizard of Oz" methods in a lab setting. UEs offer participants a glimpse of several provocative future situations, and allow them to critically reflect on the futures they desire and do not desire through interviews.

Findings reveal a preference for agents that engage in non-human behaviors such as re-embodiment; tension between a desire for a familiar agent that is "theirs" and a desire for agents with expertise; concerns about overwhelming agents' attention; and discomfort during co-embodiment interactions. UEs enabled us to create an initial map of the design space of agents interacting with individuals in social and service contexts. The themes that emerged in our findings could help designers and researchers

understand how to design conversational agents and social robots for long-term, personalized interactions across time and space.

3.2 Method

We used speed dating with user enactments (UEs) as the research method towards our research goal [34, 156]. User enactments position participants in staged scenarios using Wizard-of-Oz methodology [120], low-fidelity prototypes, props, and scenery that all contribute to creating rich and natural-seeming experiences and help participants suspend disbelief and immerse themselves in futuristic interactions. UEs build on the idea of romantic speed dating. After an evening of many quick, fake dates—some good and some terrible—people are likely to know little about any of the people they meet. However, they may have gained new insight on what they are looking for and what they value.

Thus, UEs are especially suitable for exploring new technologies, when there are no known design patterns or social mores to guide designers [34, 156]. This is because the method draws from people's reactions to identify areas of interest and to define a set of initial topics for further research. Previous work found that an open-ended approach to UEs that allows piloting and modifying on-the-go is more suitable for this kind of exploratory investigation [34]. By experiencing a set of flexible, diverse and open-ended interactions with technology through UEs—rather than experiencing them in a carefully controlled environment as would be the case in an empirical study—participants are more likely to have insightful feedback about the topic as a whole.

3.3 The Design of User Enactments

As this is a new design space laden with unknowns, we took an exploratory approach to examine flexibility of social presence and its relationship with an agent's physical body. Today, most robots and agents are designed in one of two ways—they either function as a single social presence attached to a single body (*one-for-one*) or as a single social presence that embodies all devices simultaneously (*one-for-all*). The first approach makes robots seem more human-like. Robots like Jibo [66] mostly follow this model. Some follow the second model; for example, Amazon's Alexa appears to be the same social presence across multiple channels. Beyond these two

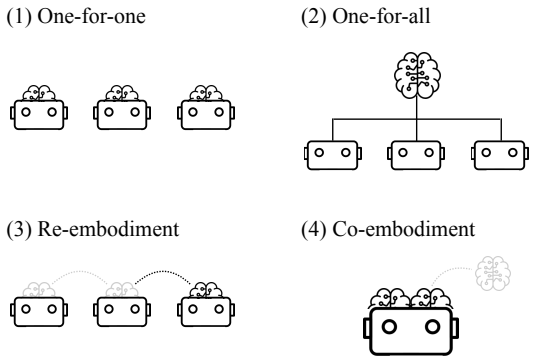

Figure 3.1. We explored four social presence options for conversational agents and social robots: (1) One-for-one, a human-inspired model in which each social presence has a single body; (2) One-for-all, a singular social presence that inhabits multiple bodies simultaneously (the behavior of many current systems such as Alexa and Siri); (3) Re-embodiment, a singular social presence that can hop from one body to another to travel with a user across a task or a service; (4) Co-embodiment, a social presence that joins another that already resides within a body.

paradigms, we explored the notions of *re-embodiment* and *co-embodiment*. Robots re-embody when they move their social presence from one physical device to another. Robots co-embody when one social presence joins another within a single device (Figure 3.1).

To allow participants to "speed date" with future scenarios that explore the flexibility of social presence, our team of 10 designers crafted and piloted enactments over the course of a month. We then turned our most promising ideas into engaging scenes in which a person might interact with multiple agents that can re-embody and co-embody.

Prior work has focused on agent migration as a concept rather than on specific tasks and contexts in which agent migration would be useful and beneficial. We therefore began our ideation process by brainstorming questions, topics, and possible contextual boundaries of physical space, social roles, and environments. Our team generated around 200 ideas using several design methods: Custom generative card games [49]; New Metaphors, an ideation method using analogy [83]; and Bodystorming, a method of brainstorming ideas through physical action and improvisation [22].

As we went through the ideation process, we used affinity diagrams to organize concepts and draw out themes. For example: What kind of situations should trigger social presence flexibility? Based on the themes that emerged through affinity diagramming, we converged on a small set

Figure 3.2. We designed four user enactments equipped with scenery and props: (from left to right) a government office (DMV), a home setting, a health center and an autonomous car.

of enactments, while working out issues of fidelity, flow, topic selection, amount of interaction, and believability. Our intention was to select a few representative scenarios to combine into a single study. As we converged on a small set of enactments, we repeatedly piloted them to make sure that each enactment was evocative for participants, and that it triggered an understanding of the intended social presence behavior. We piloted more than three times as many participants as in the final study, which is quite typical for UEs.

We converged on four enactments in four contexts: a government office (DMV), a domestic space, a health center, and an autonomous car (Figure 3.2). These settings were chosen because they allowed us to investigate the role of multiple agents in private and public settings, deal with the topic of highly sensitive data, and explore experiences that make use of social presence flexibility.

User enactments enable the presentation of different versions of the same situation, one after another. In this way, participants were given a "menu of possible futures" to choose from and an opportunity to reflect on futures they desire or fear [156]. Thus, for each enactment, we chose the two social presence configurations that our team decided would be most suitable for that particular environment.

The Enactments

We briefly describe the four enactments, the themes each focused on, and the design cues we created to communicate the intended agent behavior.

DMV

The first enactment involved a trip to the Department of Motor Vehicles (DMV) to obtain a new driver's license and register a car. In the first variation, participants interacted with three different social presences that embodied three separate bodies (*one-for-one*). Participants had to complete multiple steps of an overarching task while interacting with the three agents. In the second variation, participants interacted with a single social presence as it moved from body to body within the DMV (*re-embodiment*). This enactment introduced the user to the concept of re-embodiment and gave us an opportunity to explore re-embodiment in a fairly neutral, public context.

Design cues. We used two cues to indicate that a social presence had "re-embodied" into a new device: (1) consistent eyes on a face display and (2) a consistent voice. We used varying eyes and different voices to indicate that multiple agents were operating in the "one-for-one" variation. Pilots of the scenario showed that these cues appeared to be successful in leading participants to believe the agent's social presence had moved.

Home and Work

In this enactment, participants played the role of a pet store employee who had finished work for the day and was now collaborating with a social presence in their home to prepare for a dinner party. Participants were informed that the social presence had the capability to move from the body in their home to the body that resided in their workplace. While setting the table, participants were told by the experimenter that they had forgotten to take care of some errands back at the pet store. They were then encouraged to interact with the social presence to help them take care of both home and workplace tasks. In one variation of this scenario, the social presence could only be in one body at a time (*re-embodiment*), while in the other, the social presence could be in both bodies at the same time while engaging in different activities (*one-for-all*). During the one-for-all variation, the social presence's two bodies experienced a brief "connection loss" to probe the experience of failure. This scenario examined remote re-embodiment versus one-for-all, the boundary between personal and professional spaces, multitasking, and trust during and after failure.

Design cues. To communicate that a social presence had re-embodied to a remote location, the agent said that it was going to the office and added, "Be right back!" Then, we turned off the screen with the robot's

face on it and positioned it at a downward-facing angle to indicate that it was not active. When the social presence re-appeared, we turned on the screen and repositioned it to face the participant. In the second variation, we communicated that the social presence was in two places at once by having the agent at home speak about their progress in the office ("I can't open the door; it is locked"). In both variations, participants understood the configurations of social presence flexibility that we intended to design.

Health Center

This enactment involved acting out a visit to a health center to evaluate recovery from an injury. We gave participants a physical "token" that "contained" their personal agent's social presence (similar to how a USB stick contains data). We asked them to move the token between bodies by unplugging it from one body and plugging it into another as needed. This design intended to reinforce the event of re-embodiment through physical representation. The enactment began at the participant's home, where their in-home social presence reminded them that it was time for their appointment. The participant unplugged the token and travelled to the health center. Upon arrival, they plugged their token into the receptionist robot, which made their personal social presence embody it. Participants went on to plug their token into an X-ray machine robot and then returned to the robot at the reception station to check out. This enactment allowed us to explore re-embodiment in a more sensitive setting and address issues of context-crossing agents, privacy, and data storage perceptions.

Design cues. To indicate that the location of the social presence depends on the physical token, every time the participant plugged the token into a body, the agent's face appeared on the screen and a red LED on the token lit up (it was remote-controlled by the researchers). As the participant unplugged the token, the LED light was immediately turned off, as was the face on the screen of the robot it was plugged into.

Autonomous Car

Here, an autonomous car drove the participant home from work. In the first variation, a single social presence in the car assisted them. In the second variation, partway through the drive, the participant's in-home social presence *co-embodied* with the car social presence in order to ask about some housekeeping errands (ordering laundry detergent and opening the

door for a delivery person). The social presences also directly communicated with each other in a short segment of dialogue. This enactment was designed to probe co-embodiment and direct communication between social presences.

Design cues. We communicated co-embodiment in various ways. We explicitly told participants that a new social presence had entered the car, and we had the new social presence "turn down the music" while the social presence that was initially embodying the car continued to drive. The second social presence—which embodied an autonomous car at the same time as it embodied another body at home—had a different voice than the original driver and used dialogue to indicate its multiple embodiments.

Participants

We recruited 18 participants (10 female, 8 male) over the age of 25 ($m = 32.73$). We chose to exclude younger participants to maximize the likelihood that participants would have had some real-life experience with the kinds of scenarios we were testing. Our participants came from diverse ethnic and racial backgrounds and from a wide range of professional backgrounds, including sales, teaching, music, and engineering. All participants were familiar with computers ($M = 6.20$, $SD = 0.68$ on a 7-point Likert scale), while they were diverse in their familiarity with robots ($M = 4.36$, $SD = 1.49$ on a 7-point Likert scale). The study lasted 90 minutes (5-10 minutes for each enactment, 10-15 minutes of semi-structured interviews after each enactment, and several minutes for a final interview and questionnaire, as described below).

Procedure

Participants arrived at the lab, signed a consent form, and participated in four enactments that included the four configurations of social presence (one-for-one, one-for-all, re-embodiment, and co-embodiment). We used the DMV enactment to introduce the concept of social presence flexibility, and therefore had all participants start the experience with that enactment. We counterbalanced the order of the other three enactments to minimize the chances that the order would influence participants' reflections or confound our understanding of them.

After each enactment, the experimenter conducted a short interview with participants about their overall experience. At the end of all four

enactments, the experimenter conducted an interview in which they encouraged the participant to reflect more deeply about all four enactments and about the general themes in the study.

The study took place in several low-fidelity settings created by sectioning off areas across a lab and a classroom. Some areas were used to represent more than one setting; in this case, the props and scenery were swapped between enactments to give a sense of a new space.

Since we were exploring possible futures rather than present-day interactions, we used low-fidelity tools to create fluid experiences with possible future technologies without the need to develop them fully. This methodology was particularly useful because it allowed rapid adjustment of enactments, and quick testing of several variations for the same interaction. We used prototypes of robots and props for participants to interact with, with one of the researchers functioning as a "stagehand" and moving them as needed.

Another researcher functioned as a Wizard-of-Oz, generating the robots' and agents' verbal responses based on a set of audio files and a script. Since these were structured interactions, we used pre-recorded audio clips of human speech for the DMV, health center, and autonomous car enactments. In the home and work enactment, which was more open-ended, the agent's voice was portrayed by the "wizard", who spoke through a microphone from another room. The agents' dialogue and sound effects were played through a small Bluetooth speaker, which the stagehand moved from body to body across enactments (the speaker was placed in a hidden location). The wizard controlled the agents' side of the dialogue by taking cues from the participant's and experimenter's speech (via a phone connection with the experimenter), and with the help of a live video feed from a GoPro camera which the stagehand re-positioned for each enactment.

Three members of our team administered each experiment session. One was the experimenter, one was the stagehand (in charge of moving props and robots), and one was the wizard (who controlled the voices that interacted with participants). Each researcher played the same role in all of the sessions.

Below is a short excerpt from one of our scripts (the complete scripts used in the enactments are included in supplementary material). This excerpt describes the re-embodiment variation of the DMV interaction (the social presence's name is "Ari", and the names we used to refer to the different robot bodies are in quotation marks):

Experimenter: This variation is at the DMV. For this variation, your name is Sam Jones. You just moved to a new city and you need to register your car. Here is your car title [*passes a certificate to the participant*]. This is Ari [*gestures to "desktop robot" on a table*], the social presence who will follow you from body to body to help you. You can say Hi to Ari whenever you are ready.

[*Stagehand turns on Ari's face*]

Ari: Hi, welcome to the DMV. My name is Ari, and I'll be assisting you here today. Let's get you checked in. What is your name?

Participant: <response>

Ari: Hello, Sam. What are you visiting the DMV for today?

Participant: <response>

Ari: Okay, let me guide you to the waiting room. Please head downstairs and I'll meet you there.

[*Stagehand moves the speaker to the "tall robot" body. Participant goes downstairs to meet Ari, guided by the experimenter if necessary*]

Ari: Follow me to the registration desk.

[*Stagehand pushes "tall robot" to the registration desk, guiding the participant to follow if needed. Once the robot arrives at the registration desk, Stagehand turns off "tall robot"'s face, moves the speaker to "tinybot", turns on "tinybot"'s face, and moves "tall robot" to its next position*].

Ari: [*appearing on "tinybot"*] Hi Sam, it's me. Please place your car title face-down on the top left corner.

[*Participant inserts their car information card*]

Ari: Sam, please review your car details and let me know if anything needs to be corrected.

Analysis

We recorded and transcribed the interviews and analyzed responses using affinity diagramming [12]. We chose affinity diagramming over grounded theory because our goal was to assess possible futures—not people's current practices—and tease out some of the initial patterns and social mores related to re- and co-embodiment. Affinity diagramming was, therefore, more suitable for that goal, as it is a practice-based approach focused on

drawing out themes and insights to support designing successful products that may be implemented in the future.

Three researchers iteratively rearranged clusters of post-it notes based on their emerging affinity to one another while discussing and critiquing the emerging themes. We focused on items that did not fit with any existing clusters and items we disagreed on, working until the affinity structure of all items was resolved. All participant quotes from the interviews that were relevant to the design space were analyzed (off-topic commentary was excluded).

3.4 Results

Our synthesis of the data revealed several themes about how people made sense of their experiences. The themes that emerged include: (1) participants' acceptance of re-embodiment, (2) perceptions of agent expertise, (3) crossing contextual boundaries, (4) perceptions of cognitive load, and (5) negative reactions to co-embodiment.

Acceptance of Re-embodiment

Previous work shows that people may experience comfort when they interact with familiar things [57]. It seemed that in this study, participants felt comfortable with the concept of re-embodiment, and rationalized their comfort by connecting the agent's behavior to familiar experiences. Some compared their experience with existing products:

> P05: "It was like talking to a more intelligent Alexa, or an Alexa that had a wider range of abilities."

> P17: "I'm used to... interacting with a robot that could be in different bodies. For example, one reason that I got an Apple watch was so that I would be able to give myself reminders by Siri, when I'm not near my phone or when I want to isolate myself from distractions."

Other participants used analogies of human behavior to describe and make sense of their experiences:

> P09: (When the robot transitions from home to a remote workplace) "[It is] like waiting for a friend of mine who would go do something and then come back."

While participants tended to interpret the interactions in human terms, they did not act surprised or unsettled when the robots engaged in the non-human behavior of jumping from one body to the next. They did not view this as unnatural for an agent:

> P18: "That's the whole point of making a virtual intelligence, is the fact that their body is not the limitation."

> P01: "Even though the bodies are different, the voices, the mind, the thought process remains the same, so it's like a continuation of thoughts from one body to another."

Additionally, participants suggested that the similarity in the agents' mechanical form and their use of the same voice in different bodies made it easier for them to perceive and accept the occurrence of re-embodiment:

> P14: "It was the same voice and the bodies were so minimal and so mechanical... It was not a big deal. If you had presented me with Ari [one of the robots] who was sitting on the little table, and then something that looks humanoid, then there might be some difference."

Perhaps a more critical aspect to examine further is the exterior design of re-embodying robots, rather than the behavior of re-embodiment itself.

Perception of Agent Expertise

Participants first interacted with several distinct social presences across touchpoints in the DMV enactment, each in a separate body (one-for-one). In the next variation, they interacted with a single social presence that moved from one body to another (re-embodiment). Most participants perceived re-embodiment as more seamless, easy, and efficient. Many described it as a more natural flow of interaction compared to a piecemeal interaction with multiple social presences. In the one-for-one variation, participants felt that they had to become the connective tissue across touchpoints and that they had to re-explain the task to every new social presence. Participants did not actually re-explain the task (as we designed each agent in the sequence to already have all the relevant information), but they were still left with the feeling that they had experienced re-introductions in the one-for-one robot behavior:

P12: "It's always better to have one intelligence that guides all of the five steps and [creates] a process that is one thing."

P01: "I had this feeling that if you have a single mind controlling all the bodies, the thought process and the transfer of commands could be much faster than three [social presences] doing this... If you have multiple brains acting in multiple bodies, even a slight mismanagement between the several steps can lead to a disastrous result."

P16: "Meeting you from the beginning, following you throughout the whole thing, it felt like they had all the information in store, whereas going to each new robot, I feel like they only knew that specific step of my process."

While the preference for a single social presence that could move with participants through a service was dominant, some participants wished for a different experience. For these participants, having several social presences guide them through a service was the preferred choice. They explained it like an assembly line, where each social presence knows its role and has its own expertise conducive to doing the same thing over and over again. Participants who preferred this variation argued that it is more important to have a social presence with expertise take care of the task than to engender a personalized experience with the same social presence:

P02: "A lot of [social presences] probably works better because you have people come in one by one and directed downstairs. There is no delay. There is a nice, fluid [process], like an assembly line."

Some participants preferred a one-for-one social presence only for tasks that required what they perceived as a high level of expertise. This desire was expressed in the Medical Center enactment, where the social presence followed participants from home to re-embody the center's front desk and onto the X-ray machine to take their X-rays:

P11: "[It would be better to have separate social presences] in the last scenario [medical center] because we are talking about having medical expertise in a certain area... like when you go to the X-ray room, [you want to] have a professional person [social presence] there, not [the one from] the reception area."

Crossing Contextual Boundaries

We anticipated that participants might prefer having a single social presence follow them within a single service, as in the DMV enactment, but that a social presence that jumps between two very different contexts, such as work and home, would not be as well-received. However, participants' reactions did not confirm our hunch: many participants reacted positively to scenarios where a single social presence crossed contexts.

Private and Public Spheres

Participants mostly did not mind a social presence that moved between private and public spheres. Several participants liked having a single social presence follow them from their home to a hospital and then assist them in a medical context. These participants found the familiarity with the social presence to be comforting. The feeling of comfort contrasted with the stress they typically associated with visits to health providers:

> P11: "I thought it was convenient that it was the same person [social presence] that you are familiar with, they already know everything about you."

> P12: "You are assured that OK, you will be fine... it's better to have one guy who is dedicated to you."

Participants who disliked when a single social presence moved across contexts offered a perspective we had not considered. We anticipated that the main concern would be around privacy of information, and this was expressed by some of the participants who worried about the risk of exposing personal data in sensitive environments:

> P14: "I don't necessarily want the receptionist to know what goes on in the doctor's consulting room or the X-ray room, and vice versa."

However, the strongest negative reactions seemed to be triggered by social preferences. Several participants shared concern that interacting with only one social presence over an extended period of time and across contexts might be boring and draining. They desired variety in their interaction partners:

P06: "You need different people in your life for different situations. Like your teachers... I don't want to take seven classes and listen to the same voice over and over and over and over again."

P18: "For me, I think I am more social... It makes it less lonely [to have multiple social presences] than to have one single voice. Like, "Alexa, you're the only person I talk to and you're the only person who gets me"... I just [prefer] multiple voices, multiple people."

Private and Professional Spheres

We found that a single social presence moving between private and professional environments raised two concerns: one about the effort required to maintain a work-life balance in a case of a re-embodying agent, and the other about the ability to keep some information private.

Participants imagined that having different social presences for different domains could potentially help them maintain a better work-life balance. They believed that it would be harder to maintain a separation when a social presence re-embodies and works across both domains. Participants also worried that cross-contextual agents could cause work-related topics to be exposed to their family or friends, and personal topics, such as embarrassing habits, to be exposed in a work environment:

P05: "I personally like a separation between my day job and my home life. So, you would worry that it would remind me of emails at work when I'm at home trying to relax. It would not have a clear filter between work and work at home."

P18: "I want a clear disconnect so that mentally I can switch back and forth from home life and work."

P14: "Would people know what goes on in my home simply because this thing can transfer between environments?"

Perceived Agent Cognitive Ability

Participants expressed concern about social presences' ability to effectively carry out several tasks at the same time across multiple bodies. Some participants suggested that a social presence might be "biting off more than it could chew"—that is, that the social presence only had so much

cognitive capacity and attention to devote to tasks at a particular time, and that if it attempted to exceed its limits, system errors or other failures could occur. Current experiences with technology seem to reinforce and defy this concern simultaneously. Smartphones, for example, can maintain constant monitoring of text messages, email, social media, and news feeds while providing turn-by-turn directions. These same devices also sometimes lock up when transitioning between WiFi and phone networks. It seemed that the failures of present-day technical systems had a stronger influence on participants than their successes:

> P01: "Maybe Sigma [the social presence] cannot embody two devices efficiently at the same time."

> P12: "Imagine you give it like 10 other tasks which it is not supposed to do. It is forgetting the main task that it was designed for."

When the task at hand was driving, an activity that involves personal safety, participants expressed an even greater deal of skepticism about the social presence's capacity to do multiple tasks at once. Many participants seemed disturbed when the "driving" social presence chose to engage in tasks unrelated to driving. The strongest responses came when the "driving" social presence began talking to another social presence. Participants seemed to have attributed an understanding of limited human attention:

> P12: "Autonomous driving intelligences should only be related to autonomous driving!"

> P07: "I don't know the requirement of the capacity... focus on the safety instead of the talking."

Co-Embodiment

For the co-embodiment enactment, there were two social presences ("Omega", the driving agent, and "Eta", the home agent) in a single device (a car) who conversed with each other. As we only probed a single instance of co-embodiment in our enactments, additional work is needed to confirm some of the initial reactions we report below.

Although participants explicitly understood our intention for two social presences to embody the car, they still tended to intuitively interpret the situation according to remote communication between humans; they

described the experience as one social presence being "inside" the car, with the other "calling in" from a different location, like a conference call:

> P04: "It felt like [it] was home, calling in."
>
> P16: "I did not think of it as having multiple minds right now in the body of the car. It definitely did not feel like that to me."

We expected that in this enactment, participants would perceive the agents as more social because they conversed with each other verbally. We found that participants did, in fact, describe these social presences in social terms in the co-embodiment enactment more so than in other ones:

> P06: "...it would be just like two of my best friends."
>
> P07: "It's like they are going to marry and I will be the one who will break them up."

Social attributions manifested as more than just a social layer in the interaction. Some participants felt a social hierarchy was formed. Participants frequently described this interaction in terms of masters, servants, ethical issues, and human rights. Participants who felt that the enactment created a hierarchical social structure found it very unsettling:

> P14: "[It is] like those old English manor houses where the Mrs. says I want this, and then one servant talks to the other servant and it gets done. I'm not sure I like that."
>
> P03: "In this case, it did feel more like having servants than having just a glorified Siri. In a way that made me personally uncomfortable."
>
> P07: "They are all my servants. So I'm the one [who] should be paid attention to, that should be acknowledged."

A surprising finding was that some participants had a strong sense of exclusion in reaction to the conversation between the two social presences. Although the conversation was only about what type of windshield wipers to purchase for the user (who had already authorized the purchase), participants were uncomfortable with the feeling of being a "third wheel", and some even felt like they were being "plotted against":

P07: "I don't like it. I felt neglected. I was forgotten. They talked to each other. They made their own decisions without consulting me. Like I was isolated."

P14: "It's like they are colluding on me. They're ganging up on me."

P04: "Annoying... they acted like I was not even there."

Participants did not describe other experiences in the enactments in terms of social power dynamics, master-servant relationships, or social presences colluding, regardless of whether they used re-embodiment, one-for-all, or one-for-one.

Other participants reported not liking the conversing social presences simply because it was distracting and unnecessary. We had included the dialogue between social presences for the purpose of transparency, to inform users that information had passed from one social presence to another. Prior work has shown that users perceive silent communication between robots as creepy and undesirable [143, 150]. We, therefore, expected that a conversation would keep participants better informed and make them feel more in control. However, participants did not find this useful:

P01: "Eta and the other [social presence] interacted, or had a conversation in front of me. Can't this be in the background? They were already making the decision anyway."

P03: "The conversation between the car AI and the house AI seems wholly unnecessary... they do not need to express this verbally."

3.5 Discussion

Using speed dating with user enactments, we immersed participants in four situations involving potential future interactions with social robots and conversational agents to begin understanding the design space for socially sophisticated agents. We focused on whether agents might be designed differently than humans, with the concept of social presence flexibility. We discuss insights from the study through key questions that emerged from the results, as a set of initial considerations, we believe

designers should keep in mind when devising new forms of sequential and personal interaction with multiple agents.

Should agents act human-like? Participants easily understood human-like behaviors. They often made sense of non-human behaviors by focusing on elements that were still similar to human behavior. For example, several participants compared a social presence that moved between a collocated and a remote body as being similar to a person who leaves and comes back. We believe this tendency is precipitated at least in part by the human-like speech that agents used, as people socialize even with task-based technology like Amazon's Alexa [129]. Thus, we were surprised to find out that participants experienced some very non-human behaviors so naturally. This was particularly true of re-embodiment, when a social presence traveled from one body to another. Our findings suggest that designers should not prioritize human behavior exclusively when interacting with agents. Instead, they might consider taking advantage of non-human abilities and exploring when such abilities might add value to the interaction.

Should agents have expertise? We found an interesting interplay between efficiency in completing a task, participants' familiarity with a social presence, and social presence expertise. For most routine tasks, participants preferred to interact with a single "Jack-of-all-trades" social presence. For these low-risk contexts, people appreciated feeling that the social presence already had their information and could use it to smoothly and seamlessly accomplish all aspects of a segmented task. This was true in private, professional, and public contexts. However, participants felt less comfortable interacting with a single social presence when a task was more complex, had a higher risk, or required a higher level of expertise. For example, most participants did not mind that the same social presence embodied the receptionist, the guide robot, and the eye-checking robot in the DMV enactment. However, in the health center enactment, some participants expected to be served by experts. They felt uncomfortable that the receptionist—who in this role greeted patients, offered directions, and retrieved records—was the same social presence who took their X-rays.

This dichotomy offers an opportunity for two different kinds of social presences in interactions involving expertise: (1) A multitasking agent that is social and capable of handling many low-risk tasks, and (2) an expert agent who only has one job. Multitasking agents could provide new opportunities to design "omnichannel" and sequential services across platforms using a single social presence. For example, a user may visit

a repair shop for product maintenance and find that the repairs cannot be completed because of a missing part. The user can pick up the part from another location, return home, and continue the service appointment there, where the repair shop social presence would guide them through completing the repair.

The movement of social presences across space is not limited, which allows them to keep track of many users in many contexts simultaneously. Due to this ability, agents can be better at low-risk, cross-contextual interactions than humans. In addition, any service that currently has a "concierge" (or could make use of someone in a concierge position) might consider presenting the users with a single social presence as the "connective tissue" that can serve them throughout an entire experience. On the other hand, expert agents should be limited to the tasks on which they are experts. Since people seem to want expert agents to stick to their single task, these agents may benefit from having a behavioral design that is limited in social capabilities.

Designers should also consider that while some participants enjoyed having a single, familiar social presence that moves with them from context to context, others found it potentially lonely and exhausting. More work is needed to understand how familiarity with an agent interacts with people's understanding of expertise as they interact with an agent over an extended period of time.

How should agents express their "cognitive" ability? We found that participants were concerned with the social presences' ability to manage multiple tasks at the same time effectively. This worry was especially prevalent in the car enactment, in which social presences conversed with the participant about mundane matters (choosing a podcast to listen to, ordering new windshield wipers, etc.) while also attending to the much more precarious task of operating a vehicle. We believe that participants' concern about division of attention in high-risk situations is related to their tendency to interpret agents as human-like. Similar to expertise, designers might want to create two types of agents: those that deal with low-risk tasks across contexts, and expert (possibly less social) agents that deal with one high-risk task like driving.

An alternative option might be to reassure users by providing them with explicit feedback about where the social presence's attention is directed. For instance, if an autonomous car's social presence is only driving and is not preoccupied with other tasks, making that singular focus clear using visual or auditory feedback could give people a better sense of secu-

rity. The driving social presence can also try to assure the user that it can handle multiple tasks safely; it is possible that more open communication about the car's "cognitive" effort would have helped participants to better understand its technical abilities. The importance of users' perceptions of agents' cognitive abilities has been suggested by [115], and it is an area our team plans to explore in future research.

Should agents interact with each other? We found it surprising that the interaction between two social presences in the co-embodiment enactment made participants uncomfortable. This theme emerged in the car enactment, in which the social presences discussed which brand of windshield wipers to purchase. We speculate that the interaction between them was perceived as external to the dyadic relationship with the user and thus might have indicated an agent's "personal life" outside their obligations to the user. Other negative feelings evoked by this part of the scenario included exclusion and social isolation, and even a sense that the agents were plotting against the participant. Although the dialogue in other enactments implied that agents had interacted with each other (to transfer information about the user), participants did not allude to these qualms in any of the other enactments.

Our initial prediction was that participants would find the conversation between agents informative and entertaining, but this was not the case. As we only probed one co-embodiment scenario, the design of verbal communication between social presences needs additional research. In future work we would like to better understand if it was the conversation itself or some other aspect of our enactment that triggered such a broad range of negative reactions. We would also like to discern whether and how co-embodiment might lead to positive experiences in a different setting.

3.6 Conclusion

In this work, we set out to probe the design space of how conversational agents and social robots may evolve to facilitate future sequential and personal interactions. We specifically explored how they might do this through social presence flexibility. Because this design space is uncharted, with no design patterns or known social mores to guide interaction designers, an exploratory design research approach was the most appropriate method to examine the boundaries and challenges of this space.

We created four user enactments with an overall of eight variations to allow participants to "speed-date" with different variations of agent social presence embodiment. We conducted interviews to learn about their experiences, concerns, and values.

Our findings show that a human-like model for social robots and conversational agents is not always the best choice. Although agents' ability to socially communicate in multiple modalities is reminiscent of human-like behavior, there appear to be some instances in which human-agent interactions can benefit from agents diverging from human-like behavior. Our results emphasize the promise of agents that re-embody to support interaction across multiple touchpoints, surface a need for some agents to specialize and focus on a single task in which they are experts, and suggest that in some cases, agents should refrain from verbally interacting with other agents.

While this work did not explicitly surface the topic of agent ownership, we find some relevance of our findings here that support our later findings of the potential for *owning* an agent—participants enjoyed "familiar" agents, and wanted a single agent that can address their needs across multiple touchpoints. The boundary of potential ownership in this study was for tasks that required high expertise or high cognitive load for users. As we will discuss in later studies, the topic of expertise will resurface as an important factor for determining agent ownership design.

This work was a first attempt to understand the vast space of designing agents that understand and respond to the social complexity around them. Because of its exploratory nature and the vast design space in which we are interested, we limited this study to only focus on single-user interactions, and we made careful decisions about which social presence configurations to use in each enactment. In future work, we will address a range of multi-user interactions as an extension of this work.

4 CO-EMBODYING AGENTS IN SOCIAL INTERACTIONS

4.1 Overview

This work continues the exploration of the previous, looking into the design space of socially sophisticated agents, but extending into *interpersonal* interactions with them. How should agents behave when several people are in the space? Should the same service agent interact with both an acquaintance and me? Should each have their own agent? Or should the agent be primarily affiliated with the service provider?

Findings of the previous study suggested that agents can serve as touchpoints for personalized service across space and time. Furthermore, it found that people *prefer* a single point of contact that is knowledgeable about all touchpoints in a particular context—AI agents can craft a personalized service that goes beyond what people can do, fostering beyond human-like experiences.

To extend these findings and re-interpret them in interpersonal settings, I collaborated with another Ph.D. student, who led the following research effort. We define four research questions:

> **RQ1: How should an agent personalize its performance of service with multiple users? How does context influence this?**

We examine potential consequences of designing multiple agents, each personalized to a single user:

> **RQ2: How does co-embodiment (multiple agents embodying the same physical platform) impact people's perception of the experience?**

> **RQ3: How does a sense of personal connection to an agent's intelligence influence trust in the agent and in the service provider?**

RQ4: How, if ever, should re-embodying agents cross contextual boundaries?

In contrast to the previous study, this study makes use of more structured User Enactments, given that we already had an initial understanding of some of the important factors in this design space. We explored how agents should address personalization needs by attempting to answer the above research questions.

We compared between three interpersonal interaction agent structures: (1) Agents that belonged to the service (like current standard service agents); (2) agents that belonged to the service but generated a unique "agent personality" tailored to a specific user; (3) agents that belonged to *the user* and served them in public and personal space by leveraging personal information.

Findings surfaced that people prefer a single "social presence" that knows them and can transfer from body to body and provide services across contexts. Participants felt comforted by a single entity that "knew them" and that was capable of using their personal data in a range of situations. We also identified some concerns regarding data transfers and security, which will need to be addressed for personalized agents that move with the user to be fully accepted [116].

4.2 Method

Similar to the previous chapter, we designed a series of User Enactments (UEs) that use low-fidelity prototypes and Wizard-of-Oz techniques to immerse participants in several "possible futures" [34]. These mock-up experiences were intended to allow participants to critically reflect on what they saw, did, and felt and compare experiences to one another.

Because we were interested in learning about *interpersonal* interactions, this study was conducted with two participants at a time, who were interviewed together. The mutual interview enabled co-discovery and surfaced ideas that a single person might not have recognized [82]. Participants signed up as a team, and therefore knew each other prior to the study. We believe this contributed to the authenticity of the interpersonal experience.

The study took place in a lab, separated into "rooms" via floor-to-ceiling walls. Scripts that were generated over several weeks of ideation guided interactions with participants. For the agent, we designed a custom-built mobile robot for service tasks 4.1. The body was made of cardboard and

Figure 4.1. Our service robot prototype. The images displayed on the screen changed as different agents embodied the robot at different times.

based on an iRobot Create. The head was a Kubi desktop telepresence robot with an attached iPad. The robot stood about five feet tall and moved at a rate of about half a meter per second. We used Google Cloud Text-To-Speech with five voices to generate the agents' scripted speech in advance and kept a repository of Google TTS-generated common phrases for unplanned deviations. We used three design cues to express agent identity: each agent had a distinct name, voice, and "profile picture" that would appear on the screen when they were speaking. One of the researchers controlled the agents' physical motion and voices. The wizard, who was the same researcher throughout the study, followed a predefined script, and was instructed to deviate from the script only if the interaction with the participant required so.

Agent Configurations and Environments

Three agent configurations for personalized interaction were explored in this study (Figure 4.2).

Singular Agent—This configuration consists of one robot embodied by one agent, similar to the baseline in the previous Re-Embodiment study. A Singular Agent (SA) is affiliated, owned, and maintained by the service. It has some information about its regular customers. We explored perceptions and impressions of a single, consistent service agent that interacts with several users.

Personal Service Agent—A logical step in personalized service delivery is to have an agent-owned and maintained by the service provider but personalized to each customer. We called this configuration a Personal Service Agent (PSA). We assumed that multiple PSAs could exist in a single physical embodiment, and embody it as needed. Individual interactions

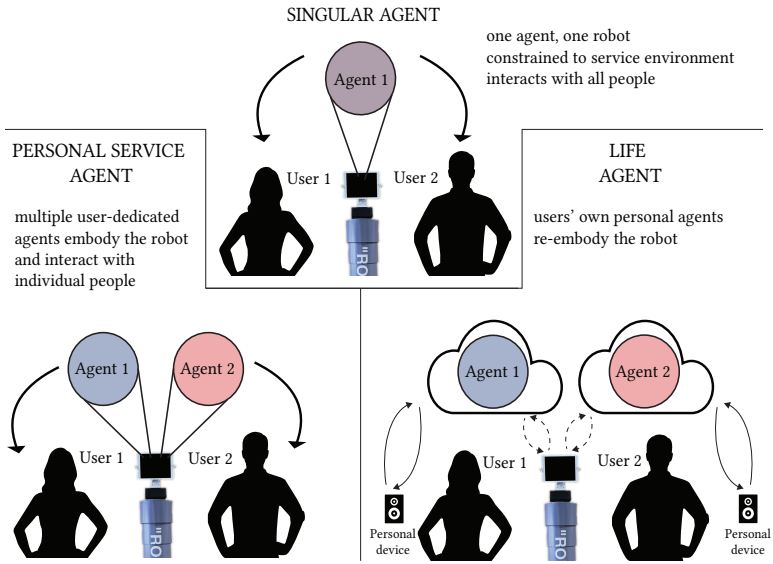

Figure 4.2. The three configurations.

with PSAs are *one agent per user* within a single environment. We were also interested in better understanding the relationship between PSAs: Should they be aware of each other's conversations? Should they talk to each other? Should they be able to share information between them?

Life Agent. A third option is for each service robot to have the option to *host* personalized AI assistants that are brought and accessed by their users. In the Life Agent (LA) configuration, agents are able to re-embody devices as needed. Each time it re-embodies, a LA can access the physical capabilities of its current "housing," and make use of the user's personal data in context.

Service environments

We explored three environments within our study: hotel, department store and a clinic. These were chosen to probe issues related to privacy and security, comfort, conversational design, long-term interactions, and social roles.

Participants

48 participants between the ages of 20 and 76 years old ($M(SD)$ = 39.3(17.6)) participated in the study. Participants had a variety of personal and professional backgrounds, but were not technical students from our institution. 25 participants self-identified as female, 21 as male, and 2 as other. Participants interacted with computers regularly, $M(SD)$ = 6.48(1.25), and not as regularly with AI assistants ($M(SD)$ = 3.31(1.91)) (on 7-point Likert scales). They had some familiarity with robots ($M(SD)$ = 3.19(1.60)), and favorable impressions of robots ($M(SD)$ = 5.46(1.34).

Procedure

After consenting, participants filled out a pre-study questionnaire to collect demographics, experience with technology, and preexisting associations with robots. A researcher then introduced the study, asking participants to take on gender-neutral, study-assigned names, and imagine that they were friends from work. In the introduction, the researcher stated that the goal of the study was to examine potential future interactions with smart technologies.

Participants then experienced each service environment with one of the three agent configurations. The order of both environments and agent configurations were counterbalanced for order and pairing. We conducted semi-structured interviews with both participants after each scenario and a summarizing interview at the end of the study.

For the analysis, we transcribed all interviews with participants and applied iterative affinity diagramming for the analysis method [12], followed by application of categorical and sub-categorical labels to quotes based on clusters that emerged during the affinity diagramming process. The analysis was conducted by two researchers, who extracted, interpreted, and grouped data together. Any disagreements in interpretation or grouping were resolved through discussion.

4.3 Results

I will discuss the main results of this study, as well as results that are directly related to the second part of this thesis that deals with ownership and service affiliation. The rest of the findings can be found in the full publication [116].

We compared a single agent affiliated with a service (SA) to personalized agents who are owned and managed by the service (PSA) and personalized agents maintained by the user (LA). Participants did not particularly like PSA, finding that more than one unique agent within a company is "redundant." Thus, the findings are centered around the differences between applying LA and SA designs. In the quotes we cite, participants refer to agents by their name: Alpha is SA; Moon and Saturn are PSA; Basil and Sunflower are LA.

Preference for a Life Agent

Most participants thought that a universal Life Agent was the most comfortable design (22 participants), followed by a Singular Agent (13 participants), and, lastly, a Personal Service Agent (5 participants). Three participants found SA and LA equally comfortable, and 5 had no preference. Participants reflected that interacting with a familiar, private agent in public spaces provided them with the smoothest and richest experience. In contrast, a singular agent was comparable to "just some random person" (119A) that does not have the data and history it needs to provide a truly personalized experience to the user.

The reasons for having a LA, over a PSA or SA is its ability to provide a highly personalized personality, to give emotional support, and to pull data and interact with its user across a range of contexts.

Personality—Participants valued agents' capability to customize their personality and identity attributes. Many wanted robots to exhibit certain character traits that they personally desired and that aligned with their values. Participant 110A wanted their agent to be hard on them. Participant 101B wanted the agent to be "sarcastic, because that's how I am. I want it to compliment me." Some had specific voice characteristics in mind pertaining to gender or dialect: 102B suggested that an agent on the East Coast should use East Coast slang, and participant 101A desired an agent with a Nigerian or British accent.

Some participants wanted agents that mirror them or remind them of their friends. Participant 110B elaborated that a "cool, calm, and collected" person should have an agent that matches their personality. Participant 107B wanted a personalized agent to "*embody the personality of my friends.*"

This finding aligns with Self-Egotism and Similarity Attraction theories that were detailed in the literature review, that argue that people are

attracted to others who are like them and feel most comfortable around those who are similar [23], or even just seem similar [68].

Emotional support—The LA design seems to have given participants a sense of comfort and emotional support, more than the PSA or SA options. Participants mentioned that in situations that might be stressful or emotional, having a Life Agent would be *"comforting"* (125A). Others thought that *"If you're feeling anxious [...], it's nice to have old friendly Basil along who knows everything about you"* (113A).

Context-crossing—LA had the ability to cross contexts in interaction, for example, to update participants about their flight status while assisting them in a medical context. Reactions to this kind of behavior that can be made possible by a LA were mixed. Some participants thought it was useful, but others found it strange, confusing, and even a "social violation." Those that believed it was useful argued that LA can provide them with "just-in-time" notifications that affect users' schedules, safety, or health, even if "out of context."

Some participants were concerned with the lack of control over such notifications. They suggested that a toggle setting that allows to permanently or temporarily turn off this feature would be valuable. It would also minimize the risk of sharing pieces of information with others in public places or in places where some out-of-context data might be inappropriate. Other participants raised a concern that a universal LA blurs the boundaries between life aspects that are usually separate, like professional and private affairs. This resulting bleed-through may not always be desirable. "For the most part, work should be work, and home should be home, should be separate, limited data passing" (118A).

Challenges of a Life Agent

Negotiating Multiple Users—Participants were unsure how co-embodying agents would negotiate between multiple users with separate needs and interests from a broader perspective. Some thought that a co-embodied agent should provide a "clear indication" about the agent that is currently embodying a physical robot (121A).

Previous research has established that simple movements can go a long way in communicating intentions of a virtual agent [145] or robot [142]. More work is needed to understand how a robot designed to convey multiple "entities" within a single body should communicate intent and negotiate user control.

Expertise vs. Flexible Capabilities—Findings in this work supported the previous study that found that doubts surfaced about a "jack of all trades" agent that might turn out to be a "master of none" [87]. Here, participants similarly believed that the more expertise a skill required, the less likely it is that a LA would have proficiency in it. Participants trusted an SA that was tied to a single domain to be an "expert" in what it was doing more than they trusted LAs.

Data Privacy—An agent that provides a personalized experience and moves from one location to another with the user is likely to embody a range of systems to accommodate their users. However, this ability also sparked concerns about data leaking from a trusted source (the agent) to an unknown entity (the system that an agent is embodying). In contrast, participants still had an increased sense of security with LAs—all of their information was concentrated in a single place, and they did not have to share it in every new context they interacted with. Instead, a Life Agent could appear and make use of the relevant data, perhaps without the need to permanently share it.

4.4 Discussion

This work has examined "Life Agents" that can move from body to body (LA) in comparison to service agents (SA) and agents that are personal but are provided by a service (PSA). The latter was found to be the least preferred configuration with the least value for participants on all fronts. However, participants were intrigued by the idea of a Life Agent that can support them in a range of services and tasks with the same mobile "entity" and database.

While this work has not explicitly discussed *what would it mean for people to own an agent*, a Life Agent is a type of personally-owned agent, and many of the findings in this work are relevant to the second part of the thesis. Findings surfaced the advantages and benefits that personally-owned agents might bring, as well as some of the concerns of using them. They also touched on affiliation with service providers and how affiliation might influence perceptions and interactions with agents.

The findings that are strongly connected to agent ownership are as follows. First, participants preferred a Life Agent, which strengthens this thesis's argument in support of personally-owned agents. Participants indicated that LAs can provide an agent that is highly adaptable to its user's values, needs and even personality preferences. Furthermore, par-

ticipants imagined that LAs would become familiar over time, and in being familiar would provide emotional support in places of uncertainty, like in a medical center. Lastly, personally-owned agents, or LAs, could cross contexts and support their users in a range of environments while leveraging a single database. This was also perceived as more secure than the current practice of sharing data with many unfamiliar systems.

Alongside these advantages, our findings begin to surface some of the challenges that designers will need to address when considering LAs. An agent that can "do it all" is also perceived with low expertise. As a task becomes more complex, or with high-stakes, participants consistently felt less likely that a single agent would handle the task. This finding was surfaced in the previous research study as well—while there is a strong preference for personally-owned agents, they are not suitable for tasks that require high expertise. In the next chapters there will be additional support for this finding, but also discuss alternative situations in which personally-owned agents can be designed as experts for a specific task.

Other concerns that were raised were more co-embodiment specific: (1) Participants were worried about data leaking as these agents embody and connect to a range of systems, and (2) co-embodiment was received as appropriate for friends, but not for strangers. Many questions were raised on how a personal life agent would negotiate between different users that all use a single embodiment, and how to clearly communicate who is the "embodying" agent at every given moment.

4.5 Conclusion

We investigated how personalized service agents might interact with multiple users and differ from current service agents. Through structured user enactments and interviews, we found that people are receptive to the idea of personalized agents that leverage information across service touchpoints and contexts. We discovered and presented some of the advantages and disadvantages of Life Agents vs. Service Agents. This work also sheds light on designing personally-owned agents and distinguishes between personally-affiliated agents and service-affiliated agents. Along with the previous and next study presented in this thesis, this work draws out the design space of personal agent ownership, community ownership, and the relationship between user ownership and service affiliation.

5 SHARED AGENTS IN SOCIAL INTERACTIONS

5.1 Overview

In this study, I address situations of a single agent within a single device, but with many users who share it. To complete the last step of the intended design space exploration, I look into how agents should behave in socially complex spaces, where they might have access to users' personal information, such as their messages, calendars, to do lists, and even accounts for 3rd party services (e.g., Spotify, Netflix, Uber). Because these are sites of many people and many interpersonal interactions, agents are likely to collect considerably more data, and they will need increasingly sophisticated rules on how to behave and on how to both share and protect personal information.

This possibility raises new interaction design challenges, particularly around privacy and feelings of control, that the HCI community seems a long way from solving. For example, what should an agent do when a mother-in-law in the home asks for her daughter-in-law's location? Should the agent share this personal information? Should it prevaricate, stall, or redirect the subject? Should it snarkily refuse? What should an agent do if a teen asks it to lie and tell parents the teen has been studying? Should it keep secrets? Should it actively deceive? Or should it tattle?

Previous work examined the sharing of devices in the home and found that sharing technology—accounts and devices alike [92]—is a common practice among friends and family members, and reflects the type and quality of relationships [54]. People share devices mostly due to the convenience of using the same device or due to economic considerations [92]. This is not without considering the issues of privacy. People are more likely to share devices with people they trust [19], and research suggests that people weigh the cost of losing privacy when sharing a device against the usefulness of sharing it [56].

When considering shared agents in a household setting, I was interested not only in privacy concerns and personal interaction, but also in how to address social differences. Previous work found that children perceive agents differently than adults and are more likely to attribute social skills and intelligence to their agents [129, 40]. Children bring a set of design considerations and risks that need to be addressed separately from those provoked by adults [125]. Sharing behavior itself also changes when kids are present; When multiple children are involved, it is common for them to share possessions with their siblings, but not with their parents [54]. Families with teens, on the other hand, are more likely to create and use separate profiles for each user on shared devices [19].

Device sharing behaviors in the home are generally determined by "household rules" that include who can use a device, for what purpose, when and where, according to their age and their social role in the home [45]. Yet few efforts have been made to design interactions that adapt to multiple users who share a single device, even when designers are well-aware of the sharing behavior around their product [92]. Brush and Inkpen identified two common models by which technological devices for domestic settings are designed today: an "appliance model" and a "profile model". The "appliance model" implies anyone in the home can use the device, relying on social protocols to mediate sharing. However, this model allows for little personalization or privacy. The "profile model" supports multiple users by asking for their identification and reducing sharing problems through individual ownership of devices [19].

The current behavior design of agents does not explicitly support the multi-user nature of their use. The gap between people's behaviors and the design of devices, and agents, in particular, indicates an opportunity to design for agents that are shared within the home. The challenge is that computational systems' access to and collection of personal data within an interpersonal context raises many questions, and design teams have few answers. They lack a clear understanding of inchoate and emergent social mores and have no design patterns to guide future design of agents' behaviors in this space.

The study was a Speed Dating study [156] with families that set out to understand how agents might better manage personal boundaries in social interactions and how their access to personal data might be integrated into the social fabric of homes. Analysis of participants' reactions revealed five themes: (1) Social roles as a critical boundary; (2) The role of agent

ownership; (3) Agent proactivity; (4) Agent sensing and collection of data; and (5) Agent judgment calls.

These findings are the last piece of three studies that can guide technical development and design of future agent behavior as more socially sophisticated. Results showed that it is not enough for agents to understand individual user preferences—when the agent is located in a social space like the home, it needs to develop more sophisticated social intelligence that includes understanding social roles, who is present, routines and "house rules". Findings here also supported the design of personally-owned agents; in contrast to the current convention of shared conversational agents, participants preferred to have personally-owned agents that prioritize their primary user's needs. Unlike the previous two studies, the topic of personal ownership was explicitly raised in this study, and encouraged us to begin thinking about agents in *ownership* terms.

This work is set in the home, as homes are socially complex spaces that include hierarchies, social norms and situations of conflict. However, we see our findings generalizing to other contexts that incorporate interpersonal interactions (e.g., workplaces).

5.2 Method

This work aimed to gain insight into the inchoate and emergent social mores within which social agents must operate. Our goal was not for families to evaluate specific behaviors, but instead to begin to map out areas where agents can and should leverage their access to personal and interpersonal data and areas where they must operate more carefully. This study was conducted with the assistance of 7 research assistants and is therefore hereon after talked about in plural.

We used Speed Dating with storyboards, an exploratory design research method that builds on the idea of romantic speed dating [156]. Using this method, researchers share several provocative possible futures in the form of storyboards and then prompt participants to critically reflect on the implications of each future (see Figure 1). Exposure to many potential futures helps participants gain insight into their own desires and values for what the future could and should be like. Speed Dating with storyboards is more of a probe than a controlled assessment. It is an open-ended approach that allows researchers to rapidly refine and change storyboards to gain insights from earlier sessions. By experiencing a set of flexible, diverse, and open-ended interactions with technology through

storyboards, participants are more likely to have insightful feedback that examines the topic as a whole. Furthermore, using storyboards allows teams to explore the possible future without the limitations of current technical capabilities.

In previous work we explored how a single person might interact with multiple agents [87] and how two people might interact with a service agent [116]. Here we explore how a small group might interact with a single agent in complex interpersonal environments. For this purpose, families were the ideal choice to investigate, given the complex social dynamics and relationships involved. Together with our previous, this work connects to the larger research question of how to design behaviors for agents in complex multi-agent and multi-user social interactions.

Ideation and Storyboard Selection

We began our ideation process by brainstorming agent behavior in social situations in the home. Our team of eight designers generated a few hundred 'one-liner' concepts using a combination of two ideation methods: custom generative card ideation decks with prompts related to the topic [49], and New Metaphors, a method of using concrete things to reflect on abstract ideas from new perspectives [83].

We used affinity diagramming to cluster our concepts and to discover recurrent ideas and themes for social agents. This generated eight socially-relevant agent behaviors (see Table 1): (1) *Proactivity*—initiation of interaction with a user; (2) *Authorized Access and Privacy*—control of the access or action that is available to an individual; (3) *Computer Skills*—the use of machine skills, such as scanning large data-sets or using an algorithm for decision-making; (4) *Prevarication*—behavior that is not straightforward with one user for the benefit of another; *(5) DDD (Dull, Dirty, Dangerous) Role Playing*—fulfillment of social roles that users may not be interested in taking on themselves; (6) *Conflict*—recognition and response to situations of conflict; (7) *Judgment*—judgment calls about the proper agent response in a social situation; (8) *Social Roles*—action-taking based on the social roles at play.

We iterated on storyboard designs until we obtained "neutrality" by minimized the appearance of an agent's behavior as obviously 'good' or obviously 'bad'. The focus on neutrality is a deviation from previous uses of Speed Dating, in which researchers would intentionally attempt to cross an interaction boundary that they thought people would not like

to confirm their suspicions [34]. We planned to interview families as a group, with the main goal of sparking discussion. Our deviation to design for neutrality also served this goal.

We piloted the initial set of storyboards with seven participants. Although the pilot participants were individuals rather than a family, they allowed us to gain some insight into whether the storyboards successfully evoke discussion.

Our team iterated on the storyboards over the course of a few weeks, gradually turning ideas into scenarios that provoke debate. We removed six storyboards and added two that addressed concepts that we overlooked in the initial set. We ended up with 19 final storyboards for the study. This set was not intended to exhaust all options for agent behavior in a home or to systematically address all topics, but to probe a range of situations that shed light on people's values and expectations of social agent behaviors in the home. We continued to remove, add, and refine the scenarios throughout the study to maximize neutrality and discussion.

The Storyboards

All the storyboards told stories situated in the home and involved at least two members of a family and interaction with a social agent. Table 1 details some of the main questions we set out to better understand, yet these were a starting point for generating discussion with participants that we expected would change and evolve throughout the study.

Figure 1 shows three storyboard examples and details the list of topics each storyboard set out to examine. In an effort to reduce gender and ethnic cues and to allow participants to effectively role play as the characters in each storyboard, we stylized the characters in a single visual style and as flat cartoon shades. The full batch of storyboards used in the study is attached as supplementary material.

Participants

We recruited 18 families, a total of 54 participants. Each family included between 2 and 4 participants [M = 3], with at least one parent and one child (12+ years old). We chose to exclude children below the age of 12, given the topics at hand, and so that participants would be more likely to form and express their personal opinions. The content of storyboards was

adapted to fit this age group. Each participant in the family (children and adults alike) was compensated 15 U.S. dollars for their time.

Twenty-six participants were parents between the ages of 37 and 58 (M = 47.48), and 28 were children and teenagers between the ages of 12 and 21 (M = 15.43). Sixteen parents identified as female, and ten as male. For children and teenagers, 14 identified as female and 14 as male. Our participants came from diverse ethnic backgrounds. They were recruited through recruitment ads in a range of neighborhoods and in proximity to high schools in Pittsburgh, PA, as well as through word of mouth. Participants had a variety of occupations, including educators, homemakers, office workers, and journalists. Of the participants in the study, 87% of parents had interacted with an agent before, and 65% own one. 96% of children and teenager participants have previously interacted with an agent, and 78% reported owning one. No prior knowledge or experience with agents were required to participate in our study.

Procedure

We conducted group interviews with families in their homes to help them better connect situations to their own lives and spaces. After reading each of the storyboards aloud and having participants follow along, the experimenter conducted a semi-structured interview to capture participants' impressions of each storyboard. The experimenter encouraged everyone to express their opinion and to add personal observations to the group discussion. The entire session lasted around 90 minutes, with 3-5 minutes spent to share each storyboard and probe participants on its specific implications. This allowed for a longer final interview where participants reflected across all storyboards. The order of the storyboards was randomized across families.

Our selection to conduct group interviews over one-on-one interviews aimed to create rich discussions, and indeed brought multiple perspectives and topics to debate. It also allowed us to understand complex family interactions as part of the interview. Furthermore, we believe the participation of their parents and siblings enabled teenage interviewees to feel more at ease with a stranger (the researcher).

However, the co-participation of children and their parents might have also caused both sides to not be completely honest and open about their opinions. Our team was more concerned with making sure children's voices were heard due to the power dynamics at hand. We, therefore,

took several steps to maximize children's honest participation within the constraints of the study: (1) We recruited children at the age of 12 and above (M = 15.43), who were more likely to express their personal opinions in front of their parents; (2) we made sure one of the first two storyboards is a session always included children as an important stakeholder; (3) in storyboards where children did not voice their opinions, the interviewers encouraged them to express their point of view.

We believe that these steps helped ensure children's participation; throughout the study, children voiced their opinions equally, even in storyboards that did not include children as stakeholders. They frequently contradicted their parents, questioned their judgment, and even called them out for changing their minds during the conversation. That said, power dynamics between parents and their children still exist and are noted as one of the limitations of this work.

Analysis

We transcribed all sessions and analyzed responses using affinity diagramming, a method that is commonly used in exploratory design research [34, 87]. Eight researchers iteratively rearranged all relevant quotes based on emerging affinity to one another through communication and critique. We continued to discuss items that we disagreed on until we reached a consensus on their placement within the affinity structure.

5.3 Results

Several patterns evolved around participants' expectations about behaviors for agents that understand and respond to social cues. Below we discuss some of the main responses and concerns and the boundaries between different design choices for social agent behavior in the home. Each quote includes the family number ("F#"), and a letter that represents whether the response came from an adult ("A") or a child ("C").

Participants were asked to put aside any responses or concerns regarding information security and privacy from the service provider's side during the study. This was done in order to better understand the specific values around *interpersonal* interaction and involvement of an agent in social matters in the home. The topic of service-providers was out of scope and therefore not reflected in the results below, but we will return to this issue in the Discussion section.

(a)

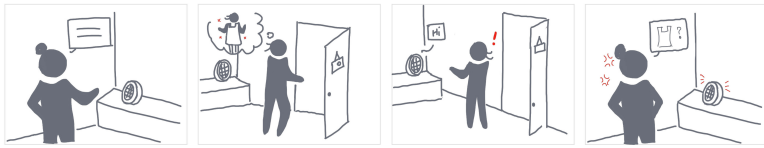

(b)

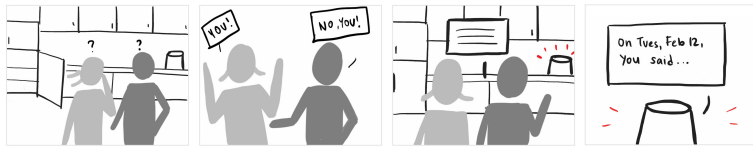

(c)

Figure 5.1. Examples of storyboards designed to evoke discourse. The topics in question for the storyboards were: judgment, authorized access and privacy in (a); prevarication, DDD role playing, conflict and judgment in (b); authorized access, computer skills, conflict and social roles in (c).

Social Roles as Behavior Boundaries

Participants brought up various social roles that family members fulfill in a range of contexts, and that might come into play as part of interaction with technology. The most significant social role division that was raised was the distinction between parents and children. Another division that had some agreed upon implications was a distinction between "insiders"—people who live in the home, and "outsiders"—who do not reside within the home. Lastly, some of the exchange was about roles between equals—siblings or partners—but these were generally perceived as interchangeable.

Parents and Children—Participants agreed that parents should have more access and control over the agent and over interactions with it. Both adults and children agreed that parents should be able to access some of their children's information, but not necessarily the other way around.

This finding somewhat varied between younger teenagers (aged 12-14) and older teenagers (aged 15 and up). Younger teenagers overwhelmingly agreed that *"explicit hierarchy should be set"* (F08-C) by providing parents with access to their information and that they, as children, should only have limited permissions. On the other hand, older teenagers were more likely to express a desire to have equal access and permissions as their parents. Parents, too, expressed a similar understanding about the nature of the relationship and pointed out that the relationship will transition over time, as should the permissions. Participant F16-A2 gave an example:

> *"A parent should be able to see what their kids are watching on Netflix until they're of age in which there should be some new privacy constraints."*

In discussions about *to what extent* parents should have access to their children's information, parent varied between wanting "just enough" information to feel in control of the household and their kids' behavior, and wanting more detailed information about what their kids were up to. F07-A1, who was in favor of having as much information as possible, explained:

> *"It could be a game changer if [parents] could catch all these little things that could be big things. Because kids have ways of hiding things and if it's something that's not good... I would say with no hesitation that having that report could be extremely valuable."*

Yet many parent participants were concerned with being over-controlling by using technology to gain information and said that it *"feels invasive"* (F17-A), and that it would involve *"getting a lot of other information [about their kids] that should just be private"* (F02-A). Nevertheless, having that information readily available felt very tempting, as they did not have to *"pick up the phone and look through it"* (F15-A1), just ask the agent, which was described as a *"really tough call"* (F14-A1). One of the teenagers who participated in the study said: *"I can't imagine my mom having the opportunity to look through my web history and not taking it."* (F06-C2). Even though participants *wanted* not to want to know everything about their

children, they admitted that they would likely ask for information if they knew they had access through the agent:

> *"I'm sorry, but I want to know. But I don't know if I would want to know. I would be scared to know, but I would also be tempted to know. So it's like you're damned if you do and you're damned if you don't"* (F14-A1).

Participants' concern that using technology to monitor their kids *"takes away their own personal accountability to make decisions"* (F13-A1) was not unreasonable—literature confirms that preventing or taking note of all inappropriate behavior is damaging to kids' personal growth [125]. Participant F02-A1 claimed that: *"It takes away all the agency from the people. So how do they learn?"*

Almost all participants agreed that emergency situations should be an exception—the agent should be able to identify them and alert the parents. Yet the boundaries of what accounts for an emergency varied from participant to participant. For example, participant F02-A only wanted to know if there was a life-threatening situation, but participant F01-A2 wanted to know *"even [about] an orange flag"*. Participants' reliance on the agent to make a call for what constitutes an emergency implies that participants expect the agent to be able to make a judgment about a social situation and its severity, as we will describe in detail later in this section.

Insiders and Outsiders—Another common distinction was made between "insiders", people who reside within the home, and "outsiders" who do not reside. Within outsiders, we found that people referred to two types that should be treated differently.

The first group was close outsiders, guests who were described as close to the family, like extended family members or visiting friends. The second group was distant outsiders, people that were not very involved in the family's life or not at all involved, such as neighbors and service providers. Some participants described their parents-in-law or friends of their children in a way that would also fit this category.

Several participants who mentioned close outsiders felt that they would be comfortable sharing some information with people in this group. For example, F01-A2 said that he would *"definitely share part of my availability"* with his visiting family. For distant outsiders, he would be *"more reserved"* about giving access to information. F04-A1 worried that someone who is physically close to her home, as a neighbor, can *"be close enough to know my availability"*, and that the agent should be able to prevent that.

The finding about limited information access to outsiders was bi-directional—several participants mentioned that an agent in their home should also not be able to access a guest's database unless explicitly asked to do so by the guest. F03-C explained:

> *"I don't think that it's right that [the agent] can just go through the other person's history because he might have something that he doesn't need people to know."*

Boundary of Agent Ownership

Participants voiced confusion about how an agent might tackle a situation of conflict between members of the home without "taking sides." For example, when the agent is asked to do two contradicting actions by two individuals or asked to keep a secret. While some participants thought that the agent should never take sides and attempt to be as impartial as possible, most participants realized that it would be difficult to maintain neutrality as agents increasingly deal with personal and social issues. For instance, F15-C2 argued that the agent *"shouldn't be in the middle"*, but was not sure what it *should* do. F08-A stated it clearly:

> *"I don't like that an agent can keep a secret, but I also don't like if the agent might hurt someone by not keeping a secret."*

Probed by the topic of conflict, the conversation frequently transitioned to discourse about who the agent is accountable to, and more broadly, who owns the agent. F13-A2 said that he does not think that an agent *"could be accountable to both"* sides of a conflict, but that he would want the agent to be accountable to him. According to F04-C:

> *"It depends if it's your agent or if it's your family's agent. If it's your actual agent, it should be loyal to you. If it's your family's, I don't know."*

A few participants said that if an agent is located in someone's personal space, like their bedroom, they would expect that they *"have priority over it"* (F11-C), in contrast to when it is placed in a shared space and used by everyone. Previous work supports this insight, and has shown that devices that are located in personal spaces are indeed less likely to be shared, and more likely to be associated with the occupant of the room [19]. We find that this carries over to agents in the home.

Thresholds for Agent Proactivity

We identified three thresholds of agent proactivity that varied between participants and within participants according to the social context. The thresholds were: reactive, proactive, and proactive recommender. Previous work has suggested that people have different expectations from technology depending on whether they have a relational or utilitarian service orientation [79]. Still, we found that in addition to a general personal preference, participants' desired agent proactivity changed according to the situation. For instance, one parent in the study wanted the agent to offer parenting advice only when specifically requested but expected the agent to be highly proactive if they learned that their teenager was drinking beer.

Reactive— When participants desired the lowest level of proactivity, they wanted the agent to respond only when being directly asked. Overall, participants who wanted an agent to be solely reactive were still positive about having the agent respond in social situations, for example, have *"Alexa give relationship advice"* (F18-C1). However, the response needed to be solicited by the user. They expected the agent to intervene in conversation only when it was invited to do so. This could be either by request, or if the user opted into a service. Participant F12-A1 explained:

> *"I think if I had signed up for a healthy eating service that would be fine [...] [but] I don't want it [the agent] to just randomly tell me that I should change my eating."*

Proactive— The middle threshold of social proactivity desired by participants was for the agent to be proactive, and be able to intervene in social interaction without being asked to do so. However, participants expected the agent to provide them with *information* about the topic, but avoid giving any *recommendation*. Ideally, *"It would give you access to information that you wouldn't ordinarily get in a very direct way. And then you as the adult would have to make a decision"* (F07-A1). F02-A provided an example of what they would expect after ordering dinner, for example:

> *"It [should] say, okay, this is what your percentages of what you've ordered over the last month [healthy or unhealthy food]. It gives you information for you to make a decision."*

Proactive Recommender—When participants desired for the agent to cross the highest threshold of proactivity, they wanted the agent to provide

not only information, but also a recommendation for their next course of action. Participants emphasized that the agent must not *enforce* a particular recommendation, but leave the choice to the family members, as expressed by F11-C:

> *"It can't really prevent you from doing anything. But it could encourage you to do things."*

Between these three thresholds, participants preferred different ones for different situations, and varied in their preferences from one another. Yet the top and bottom boundaries were clear: participants who wanted the most involved agents still did not want them to enforce any decision upon them. On the other end, participants who only wanted the agent to provide information when asked were still open to the idea of an agent who can understand and respond in some social situations.

Sensing: Agents that Watch, Listen and Record

Participants had strong negative responses to any kind of behavior on the agent's part that involved *"looking at them"* (F16-A2), *"always listening"* (F18-C1) or *"recording everything"* (F02-C), and they generally preferred the agent to use other sources of information and avoid the above.

It seems that participants' negative reaction was derived to some extent from lack of transparency: In some situations participants thought it would be useful for the agent to listen, watch or record (for example, in a case of an emergency), but they wanted to know exactly when the agent was doing so. Here too participants wanted to be *"explicitly in control"* of the agent's behavior (F12-A). Even when suggested that the agent would be transparent about sensing, participants did not trust that it will not collect data all the time, whether by accident or for the gain of a company stakeholder (e.g. Amazon). For example, participant F09-A1 explained that this lack of control felt creepy:

> *"Having this agent listen in, [...] how does it know that [information] shouldn't be shared with the kid? There [could be] some keyword or something and then the agent spills. It's creepy to me."*

Interestingly, this was not the case regarding collection and usage of digital information: emails, texts, online behavior, search results, documents, medical records and more. Quite the opposite—participants felt

comfortable with an agent using these sources of information. One participant expressed that it would be *"kind of exciting"* if the agent could make use of her family's *"search history, activities and calendars"* to make recommendations (F06-A1). F10-C said she would be fine with the agent tracking her location, and F03-A2 made a comment that it makes sense that the agent would have the texting history of all the users in the home, as it *"knows about everybody's everything"* anyway.

This finding suggests a tendency to use social norms to make sense of technology, as supported by the literature [114]. People do not want other people to "listen in", and similarly they seem to react strongly to an agent doing so, frequently referring to the idea of a "Big Brother". The conversation about an agent responding to what it *saw* made one of our participants (F16-A2) describe it in very humanlike terms:

> *"I don't like this one because it's looking at me. This one has eyes. It's starting to become self-aware."*

While aversion towards agents who "listen" and "watch" was evident, and although participants made more connections to privacy and security concerns during these discussions, it does not necessarily reflect which sensing technologies are the most privacy-invasive ones. Previous work has shown that people tend to underestimate how much information some technologies, for example GPS tracking, can extract about them, and overestimate how much information "creepy-seeming" technologies can extract [95]—this could explain people's general acceptance of an agent that uses their data and tracks their digital footprint, and rejection of one that can "listen in".

Agent Judgment Calls

As participants were discussing a range of social situations in which agents might be involved, they conveyed, explicitly and implicitly, that an agent should be able to make a judgment and "do the right thing." F15-A1 called it the agent's *"little moral compass."* Participant F11-A2 argued that if the agent is *"always going to intervene and point out cheating [in a game] then it seems like a terrible idea."* Instead, it *"has to be able to decide when to intervene and when to not."* A more implicit example is in a case of an emergency, where many participants expected the agent to be able to identify the emergency and report it. Emergency was also frequently

described as a special case that changed participants' preferences, for example, their preferences for agent proactivity or sensing capabilities.

Several participants realized that making judgment calls was subjective, and that there will not always be a universal "right thing" to do. Some participants suggested that the agent should judge situations according to the house rules and norms, or even give users control by allowing them to *"check off a series of things that [they] considered dangerous"* (F03-A1).

While people acknowledged that the agent would need to make judgment calls in all kinds of social situations, people disliked when the agent's recommendation or decision came across as *judgmental*—for example, some participants felt the agent was judging their lifestyle choices when it suggested an alternative behavior, or that it was judging their parenting when it presented parenting advice. F18-C2 stated:

> *'I just don't like the idea of things from the past being brought up, or [an agent saying] "Oh here are your tendencies" and just having to hear the agent telling you what your flaws are."*

5.4 Discussion

Our work identifies several areas of focus to consider as agents are being placed in social spaces, and are gaining more personal information about multiple individuals. We elaborate on several future technical abilities that would help agents recognize and navigate interpersonal relationships, recognize varied social roles, and help take appropriate actions in complex social situations. We also discuss the tradeoffs and drawbacks of such abilities, how they connect to personal ownership, and how all of this aligns with current service providers.

Behavioral Boundaries

Social Roles—Our findings provide strong evidence that social roles are critical for shaping agents' behaviors. Improved agents would not only know who each user is, but also know about their users' social roles in the home. Using this knowledge, an agent could make better decisions about how much access to information a user should or should not receive, what actions they should be able to perform through the agent, and whether they should have control over another individual's data.

In our results, participants agreed that children should have less access and permissions than their parents, and that "outsiders" should have less access to information than the people living in the home. For example, one parent said that they would not want their children to have access to all of their information, but that they feel comfortable giving them access to the family calendar. Another participant said that a relative, an "outsider," can ask for information, but should only receive it partially; instead of giving her mother-in-law full calendar access, she expected the agent to provide "just enough" information on her availability at a particular time.

We learn that social roles are potentially an important behavioral boundary that is not currently being used by agent designers. Recognizing or asking *what the social roles are* in a particular environment would allow designers to better tailor agent behaviors and responses to a social situation. For example, by recognizing if a child is asking for access to their parent's data, or if a parent is asking for the child's data, the agent can react appropriately.

Nevertheless, designers should consider the complexity of the social roles at play, and rely on research from a range of disciplines when attempting to generate appropriate social behaviors. Previous research, for example, has shown that preventing all inappropriate behaviors by children in harmful to the development of their sense of judgment [125]. Our findings also indicate that parents might access information about their children if they knew they had it, but would prefer not to be tempted in the first place. Thus, perhaps the best design in this instance would be to limit access, as opposed to directly following parents' desire to have all the information. Furthermore, social roles are fluid and simultaneous; a relationship between a parent and a child evolves and changes over the years, and the appropriate agent response is also likely to shift.

Our results point out the range of desires and needs at play in an interpersonal space like the home. Further research that combines design and psychology expertise will assist in defining transition points and identifying how to design changes in agent behavior accordingly. By being able to indicate social roles and connect them to broader social implications, designers could leverage this information to provide a more thoughtful and socially-aware user experience.

Personal Space—The results suggest that an agent should be aware of who is present and what social role they may be enacting at a given time to accordingly adapt their behavior. This is because an agent might want to make decisions about how and when to use personal information while

considering who else is in the room. For example, if a child's birthday is coming up, participants agreed that an agent should know when the child is present and avoid bringing up conversation about a surprise party or gifts. If guests are in the home, the agent should avoid any potentially sensitive information, such as finances or medical topics.

While additional research can be done to better understand how an agent might behave differently in a range of compositions of individuals at a given time and place, an initial step towards improving social sensibility of an agent could be for the agent to respond to two groups: children and "outsiders." The presence of children can be acknowledged by excluding any age-sensitive or age-inappropriate information, and the presence of outsiders can be taken into account by having the agent provide less information, or provide it only upon request. The outsider's group can be further divided into behaviors that are suitable when close outsiders are present (e.g., relatives) and when distant outsiders are present (e.g., service providers).

Identifying who is present can allow designers to create agents that are more socially aware and more socially appropriate. While an agent can make initial assumptions based on the people in the room, users should have the control to override these assumptions, as supported by our findings.

Personal Ownership—Our findings highlight the importance of the question of agent ownership in constructing agent behavior. Participants were not able to settle dilemmas that surfaced in situations where their needs conflicted with someone else's, and the agent needed to take action. For instance, a few participants felt equally uncomfortable thinking about an agent that would keep a secret and with an agent that would hurt someone by telling their secret upon request of another person. This led participants to express a strong desire to know who the agent is accountable to, and who has priority in situations of multiple contradicting requests or needs.

The confusion around accountability of the agent reveals a design opportunity for personal agents in the home, and space to re-consider the convention of a single, shared agent. Instead of applying the "appliance model" to social agents in the home (shared devices that anyone can use), agents could be considered from a "profile model" perspective, that allows personalized interaction for multiple users [19]. Having different agents for different users could provide users with a better understanding of how agents are going to behave in the broad range of possible social situations.

Furthermore, having multiple personal agents in a single space does not necessarily mean having multiple devices. Previous work has found that participants felt comfortable with the idea of multiple "social presences" (i.e., digital entities) embodied in a single device [87]. Thus, we believe there is potential to design multiple entities according to the number of people in the home. In turn, this structure could satisfy users' needs to have a sense of an agent that is accountable to them without overwhelming the home with agents. We intend to further explore and evaluate the notion of multiple personal agents in a single device in future research.

Proactivity

The topic of proactivity was also an important one to consider for shared agents in interpersonal spaces. Two factors indicated how proactive an agent should be: situation severity and user preferences.

Situation Severity—Many participants indicated that they would expect an agent to be proactive when there is an emergency, or even a life threating event in the home. Participants were more inclined towards an agent notifying other family members, and expressed concern about notifying any external factor (like the police).

Personal Preferences—What makes this problem complicated, is that different participants had a different view on what makes for an emergency. Some participants thought that a minor drinking beer counts for an emergency, and that an agent should notify parents when a situation like this occurs. Others strongly disagreed, and only expected agent proactivity in life-threatening situations.

However, participants shared a view of proactivity thresholds, and indicated three levels of possible proactivity that can vary depending on the situation: reactive, proactive, and proactive recommender. Participants varied between wanting the agent to only be *reactive* to user requests, wanting the agent to be *proactive* by providing information, or wanting it to be *proactive* by providing *recommendations* for a course of action. Thus, future agents can think of proactivity in thresholds, and learn when they have crossed one of the three in interaction within a particular household.

Moreover, the top and bottom boundaries of proactivity were consistent among participants—most participants accepted the idea of an agent that can respond in a social situation, but none of them wanted the agent to *enforce* a recommendation (e.g., prevent them from ordering an unhealthy food choice). Additional research in this area could help designers better

understand when they might anticipate reaching different thresholds of proactivity, while also taking privacy concerns and personal preferences into account.

5.5 Conclusion

A first step towards designing a socially sophisticated agent could include learning about (a) the social role of each user, (b) which users are in a space at a given time, and (c) what is the social context. The agent should also clearly communicate to its users (a) what level of proactivity is it set to in a range of situations, and (b) given multiple users, who is it primarily accountable to.

These suggestions are meant to be implemented in an ideal privacy structure, where data is owned and controlled by end-users. Unfortunately, this is not the case with current agents. Today, agents are constantly making headlines due to security breaches, misuse of data and a false sense of privacy given to end users.

In our work we asked participants to put aside concerns about the service-provider when responding to the range of scenarios. The goal was to understand the interactions themselves that would support family needs and desires. Even with this request, participants occasionally referred back to the "Big Tech" industry and their concerns about companies collecting sensitive personal and social data and making use of it. For example, several participants mentioned that they would not want an agent that could make their behaviors public, and intentionally or accidentally report problematic behaviors like alcohol consumption by minors to authorities.

In light of current commercially available agents' privacy policies, our first three sensing recommendations should be considered more cautiously. Sensing users' social roles, presence, and situations in the home is likely to expose information that users would not want to share with service providers. As long as users do not have full control over their data and service providers are in charge of users' personal security, our latter findings, agent proactivity and accountability to a single user, are perhaps safer choices for implementation.

Part I Summary

Although agents are entering social spaces and are gaining an expanding volume of information about their users, designers still know little about how they should make use of all that information to make interactions more social and personalized. I completed three studies to form an initial map of this design space and its implications.

The first study looked into interaction with a single user across multiple touchpoints and examined differences between a single agent that crosses contexts and multiple separate agents. The second study took the next step to examine interpersonal interactions with two users within a service context—agents were primarily service providers, but they attempted to deal with the complexity of providing personalized interaction with two users. Finally, the third study took this exploration even further to examine how agents can fit into the complex social environments they are placed in, while leveraging findings from the last two studies. This study primarily addressed the social complexities that arise from having a single agent be shared among multiple users, and needing guidelines on how to make use of private information accordingly.

One of the primary findings from this body of work is the topic of *Agent Ownership and Affiliation*. The last study of the three was the one that explicitly raised this issue, but the other two studies also dealt with discussions that are directly related. Across the three studies finding show that there is value in determining whether to design for *personally-owned* or *shared* agents.

Personal ownership gives users a better sense of who the agent is accountable to, and as a result, allows them to better anticipate how the agent is likely to respond to a range of socially complex situations. However, personal ownership also raises several challenges, such as concerns about expertise and cognitive load, and synchronization between several personally-owned agents.

This finding is the base for the rest of my thesis, in which I will continue to explore personally-owned agents, and how they compare to shared agents. I will look into their advantages, disadvantages, and the relationship of user ownership to affiliation with service providers.

Part II

Agent Ownership and Affiliation

6 OWNERSHIP OF AGENT (and) SERVICES

6.1 Overview

In the first part of the thesis, I have identified that there is value in understanding perceptions of "ownership" in the context of social agents—as agents become more socially complex, it is critical to understand people's current sense of ownership over their agents. My prior research suggests that the mental model of who "owns" an agent is critical to set expectations, perceptions, and interactions. Nevertheless, it is rarely discussed in the context of technological services, let alone agent services.

Previous work has extensively examined ownership and possession of material things [10, 75, 32]. For example, something could be "mine" because I physically possess it, because I have high a familiarity with it, or because I created it (among other reasons) [**?**]. Further, the characteristics of when something is perceived as "mine" vary significantly when the thing is digital [108, 109].

This potentially becomes even more complex with social agents, as they combine both a physical component (i.e., the device), a digital component (i.e., data) *and* have a service embedded in the device and as part of the product (i.e., Alexa). Each of these components could be perceived as owned by a different entity, which makes it difficult to determine whether one's agent is "theirs" or not, or how to design it to be perceived as so.

One way to begin understanding agent "ownership" is by comparing agents to other devices, services, products, and things that people own and use, which this study aims to do. By inviting participants to engage with products and services that they own and use through a design research approach, we can begin to form a mental model of agent ownership, and as a result, learn about how to design for it.

6.2 Background

This work builds on several bodies of literature and theory to understand current perceptions of ownership: Consumer Behavior Theory that describes why people have possessions and how they relate to them [10, 75]; HCI and digital media studies that discuss digital possessions and how they differ from physical possessions [109, 108]; Service Theory that informs the many possible relationships between *product* and *service* for technological devices [96, 47, 43]; and the fields of HCI and HRI that inform aspects specifically related to agent identity and its impact on service [26, 87, 101, 11].

To note, few research studies have examined people's perceived ownership and possession in the context of services. One example is Gruning, who shows that people perceived a weaker sense of ownership over their Amazon Kindle e-books than over Kindle e-books that they stored elsewhere, although both were digital [53]. If we step back to examine technological services altogether—the topic of their ownership, affiliation and attachment has been scarcely explored in HCI research, even though many companies have already moved, or are in the process of moving from a strictly "product" model, to a combination of product and service, or even a full "service" business model.

Therefore, this work's contribution is a proposition to treat social agents as novel service interfaces that combine service, product, device data and agent identity—all creating rich ground for new human-centered designs and better data practices.

6.3 Procedure

The study was designed as a remote study that used video conference calls and a virtual "whiteboard" application, Mural. All activities in the study were designed to prompt discussion about ownership and affiliation. In addition, the study was designed to encourage participants to compare conversational agents they own with other things they own and use, including technological services.

The study was conducted through one-on-one online video sessions with the experimenter that lasted up to an hour (due to the Covid-19 pandemic). During that time, participants were asked to do several activities on a virtual Mural board while being interviewed by the experimenter. The interviews and virtual activity were audio and video recorded, and

participants were compensated $15 for their participation. The procedure was approved by the institutional review board (IRB).

Participants

Thirty-five participants across the United States and Canada were recruited to participate in the study. This sample size was selected as a sufficient sample to reveal qualitative findings of ownership of technology and technological services. Participants were recruited and screened through the Prolific platform—the study only included participants from the US and Canada, who were fluent in English, had a minimum of 50 previous submissions, and an approval rate of 90%. As the study was designed to be conducted remotely, we also screened for participants who agreed to participate in video interviews and had a webcam. Lastly, participants were screened to only include people who own a conversational agent, as the goal was to look into current perceptions of conversational agent ownership and affiliation. Selected participants were diverse in age, gender, ethnicity, socio-economic background, and living situation (living alone, with a partner, family, etc).

Procedure

Participants were asked to answer a short pre-procedure questionnaire. The questionnaire included a consent form, demographic information, and two short 4-item and 3-item scales to control for attitudes towards technology and attitudes towards privacy, adapted from Burbach et al. [21]. Once completed, participants were asked to book a slot for a 1-hour video call with the experimenter for the main procedure.

The video sessions consisted of three activities: (1) Creating a list of people, things, and services that participants *own*, *use* or interact with; (2) Ranking and organizing this list of items on three scales based on prior work; (3) Sketching of participants' relationship or interaction with their currently owned conversational agents.

Part 1: Creating Personalized Notes

The first activity was to create a list of services, things, and people that people own, use and interact with. Creating a list as part of the procedure ensured that it was tailored to each participant, which was critical for the next activities.

To start the list, the experimenter created an initial set of virtual post-it notes before the interview that included some notes that were likely to be relevant for all participants (phone, computer, funds in the bank), as well as customized notes according to each participant's pre-procedure questionnaire (for example, notes for the people they live with, their pet (if they own one), and the specific kind of conversational agent they own).

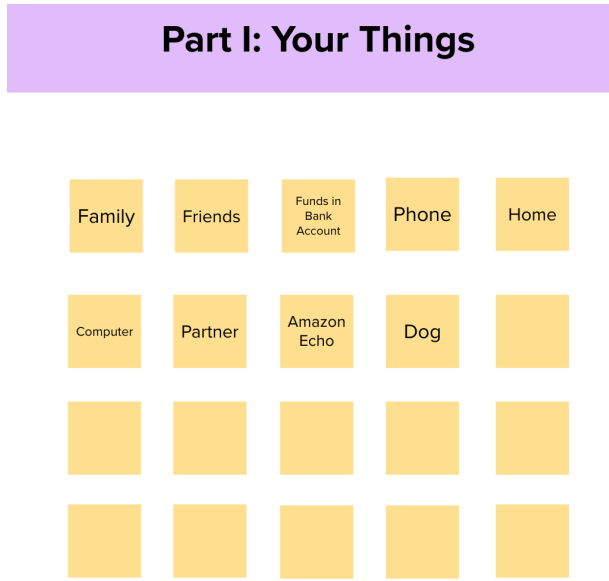

Figure 6.1. Each study began with a small set of sticky notes created by the experimenter based on participants' pre-procedure questionnaire. The figure shows the pre-study notes for P32.

During this activity, the experimenter asked a pre-defined set of questions to determine the things, services and people that are part of participants' everyday lives. The goal was to create a representative set of personalized (but not necessarily comprehensive) "sticky notes" of things and entities for comparison in the next activity. Some example questions in this part include: "what do you have in your home that you care about?", "what digital services are you subscribed to?", and "do you own a car?". This activity concluded with about 15-20 notes for each participant that included people, pets, things, devices, and services they use.

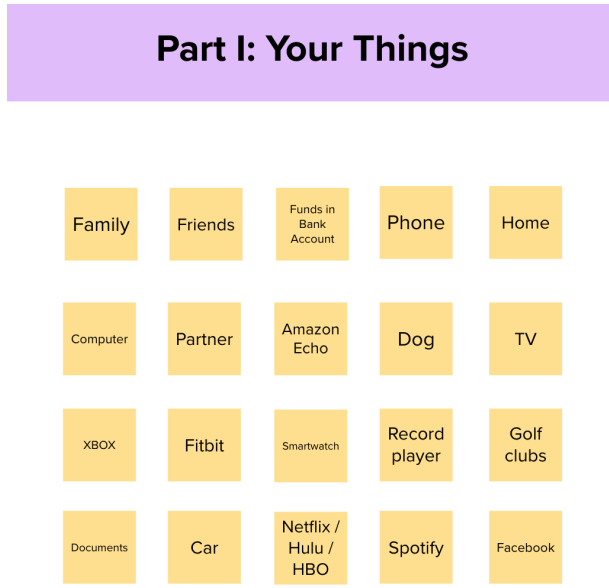

Part 2: Organizing Notes on Three Scales

In the second part, the experimenter made copies of the custom list from activity number one, and placed them next to 3 scales: Mere Ownership (mine vs. not mine) [74], Self-Extension (symbolizes me vs. does not symbolize me) [74] and Self Expansion (inclusion of other in the self scale) [2]. These scales were selected as representations of the prominent theoretical constructs related to ownership and were used as prompts to evoke reflection on the meaning of services people own and use.

During the activity, the scales were revealed one at a time, allowing participants to focus on the task at hand. We began by asking participants to organize their notes on the clearest and most concrete scale related to ownership—a spectrum between "mine" and "not mine." Then, the more abstract "symbolizes me" versus "does not symbolize me", and finally, on the most abstract visual scale. In a case where participants requested more information about what a particular scale means, they were told that there was no right answer, and that they should take some time to think, and

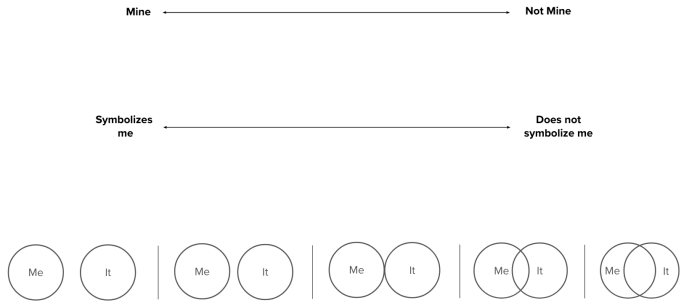

Figure 6.2. Scales on which participants were asked to organize notes of things that they own and services that they use: Mere Ownership scale (mine vs. not mine) [74], Self-Extension (symbolizes me vs. does not symbolize me) [74] and Self Expansion (inclusion of other in the self scale.) [2]

organize their items intuitively. All participants were satisfied with this prompt, and fully completed the organization task.

Participants were given 3 minutes to organize the notes on each scale, during which the experimenter "stepped away" by turning off their virtual audio and video channels. If participants seemed to still work on organizing notes after 3 minutes, the experimenter did not intervene until participants no longer moved notes around on the Mural board. Figure **??** shows an example of how one participant organized their items on the Mere Ownership scale.

After each of the three organization tasks, the experimenter used a semi-structured interview to understand participants' thought-process. The organization of items was used as *a prompt for conversation*, rather than a formal measure. i.e., participants were interviewed based on how they organized the items on these three scales, and the interview itself was intended for qualitative analysis (rather than the scales).

Participants were first asked to explain how they organized their items and why. Then, according to the semi-structured interview script, the experimenter asked participants about particular choices that stood out—for example, to compare between two things that seemed similar in how they work, but were placed in different areas on the scale (such as a phone and a computer, or Netflix and Spotify), or to compare between two very

different things that were placed in the same area (for example, a family member and a device).

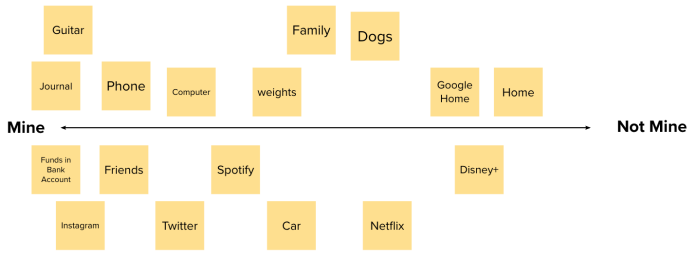

Part 3: Sketching Relationship with a Conversational Agent

For the third activity in the study, participants were asked to create a sketch on paper—prior to the study. They were asked to have a pen and paper nearby for this purpose. In this final activity, participants were asked to sketch "their relationship or their interaction with the conversational agent in their home". Here too, they were not given any further instruction, other than that the sketch does not need to be aesthetically pleasing and that it can be a drawing-like sketch, a diagram-like sketch, or anything else that comes to mind. This activity lasted 5 minutes, after which participants were asked to upload their sketch to the Mural board.

Like previous activities, the sketch served as a prompt for conversation and reflection. Sketching their conversational agents encouraged participants to discuss a range of their characteristics and uses. It, therefore, prompted diverse topics for discussion based on what participants sketched. In the interview, the experimenter first asked participants to describe what they sketched and then continued to ask specific questions based on a pre-defined semi-structured interview, with questions about placement, form, color, and symbolism, among other topics.

Analysis

As the activities in the study were not objective measures but rather prompts for conversation and reflection, the analyzed data was the *qualita-*

tive content of the interviews and participants' responses to the presented questions. Participants were asked to explain why they placed each service and device in a particular spot on the three presented scales: ownership, self-extension and self-expansion. Thus, the results are not objective but rather presented ways in which people rationalize the perceived importance and value of each service and device they use and own.

All interviews were transcribed and analyzed through Affinity Diagramming [12], an analysis method commonly used to identify emerging themes in exploratory design research [34]. The next section describes the emerging themes and how they align or misalign with prior theory about possessions and services.

6.4 Findings

Prior theoretical work has suggested that the concept of *owning*, or *having*something, is valuable to people because it allows them either to do something, or to be who they are. In other words, having a possession either provides a tool, a means to complete a particular utilitarian goal, or it is not utilitarian, but rather provides a *symbolic* meaning [123, 10].

The findings in this chapter are thus structured according to this prior observation, and attempt to observe the comparison of technological services and agent services to possessions through this lens: services that are perceived as utilitarian—"mine" but not "me", and services that are perceived as symbolic—"mine" and also "me". I will focus on the latter, as it is arguably the more surprising and interesting value proposition for agent services, which has not been previously explored.

Utilitarian Technology: "Mine but not Me"

Participants viewed some technology, particularly "omnichannel" devices such as smartphones and computers, as a literal extension of themselves—tools that allowed them to do things that they otherwise could not and useful for "almost all activities" (P20). In P12's words, their computer is a *"link to anything and everything, no matter where I am in the world"*. Further, participants viewed devices as primarily a tool that "connects [them] to other people" (P14), which is, according to P06, essential: *"You need a phone for work, you need a phone to just make phone calls and to text and [for] everything"*. Occasionally, participants referred to other technology similarly. P14, for example, thought that their Netflix subscription is *"just*

a vehicle" to have conversations with old friends about childhood content or that social media services allow them to keep in touch with people who are not nearby. Thus, in these cases, services are meaningful as tools due to what they allow individuals to *do*.

Yet especially in the case of services and devices, the value was not about a specific possession, but rather came through in the *idea* of having a particular device or service as an available tool. As P26 put it, *"I can replace the model [of my conversational agent], but I think my life will be kind of boring and dull without it. So I might be able to replace the model, but the [important thing is the] whole concept of having one"*. P32 emphasized the importance of replacement—losing a device is risky, but there is also *"a possibility to be made whole again [through replacement]."*

Unlike devices that participants agreed *"would be really difficult without"* (P26), most services, including social media and conversational agents, were described by participants as services that one can live without: *"going through a financial turmoil, those things [services] are the first to go"* (P26). According to P05, *"they could go away and [they] wouldn't be crying about it"*. Unlike devices, participants did not feel that they *"rely on [services] all that much"* (P20).

However, for both devices and services, participants described them as utilitarian tools, with no contribution to their perception of identity: *"Instagram, Reddit, that sort of thing I use quite often, they're tools to me. So I enjoy them I guess, but it's not primary to* ***me*** (105)". In other words, although participants held a perception of *ownership* over these services, they did not view them as *"symbolic of people"* (P29)That said, it is difficult to ignore the fact that *all* participants in the study had a phone, a computer, (mostly paid) TV and music streaming services, social media accounts, and many more devices and services. Thus, there seemed to be tension between the frequency of use of these services and devices (*"I feel intertwined with [Netflix] because I do use that on a daily basis"* (P07)) and the little meaning people attribute to them *(I have my Google Home in that area because like it's in my life, but it's sort of like background noise (P12)"*.

Symbolic Technology: "Mine and also Me"

In contrast to utilitarian value, some possessions are valuable to people due to their symbolic properties—possessions that help tell the story of their identity [93]. The study identified many instances in which participants viewed their technological devices and services not only as their tools but

as *representations of them*, services that were described as *"identity things"* (P30). The findings in this section attempt to identify the themes of what makes devices, services, and agents contribute to people's sense of identity and therefore form attachment and even self-extension. Further, using these findings, I hope to learn where more design efforts can be invested to create a stronger sense of *symbolic* ownership in services. Designing for symbolic ownership can increase attachment to services, and as a result, create a more positive and meaningful relationship with them.

Although services are not *owned* by definition, participants still viewed them as part of their identity. P11 stated that *"a lot of the content I consume reflects who I am"*. P17 argued that: *"the kind of films, shows, music, [and] podcasts I listen to are an expression of who I am [...] I'm a very creative person and I love arts and culture."* P31 was especially proud of their selection of music on Spotify: *"music tastes are more associated with your identity"*. P32 discussed her personal social media feeds (Facebook and Twitter) as representations of her political beliefs: *"If you read my feed, you have a good idea of where I stand politically or emotionally—you can tell who I am"*.

When talking about devices as symbolic extensions of themselves, participants tended to describe aspects of their identity that these services represent. Primarily, their interests (*"I am really into music [...] so [Spotify] kind of symbolizes me"*–P35), their behaviors (*"you can tell a lot about me from the stuff I ask Alexa"*–P26), and their values (*"Reddit symbolizes me more than any other website because I like what it stands for"*–P15).

Sartre's theory of self-extension suggests three elements that inform the extent to which one extends their perception of identity onto a thing: (1) how much *control* one has over it; (2) how much one participated in its *creation*; (3) how much intimate *knowledge* does one have about it [123]. Below I describe how each of these elements was expressed in the context of digital services and conversational agent services.

Access (as Control)

Control over a thing was expressed in the study through the issue of **exclusive access**—access to an account or access to particular functions. As the usage of digital services is exercised through action on a platform, *sharing* this access necessarily requires giving up some *individual* control. With some services that support shared interaction through profiles, such as Netflix, participants were less likely to perceive them as "mine," but

instead tended to refer to them as "ours." Although participants had less *individual control*, they had some sense of *shared control*.

In contrast, in services that did not support shared access among several people, but rather assumed single users, participants who shared access to a service described a reduced sense of ownership and perceived identity. P02, for example, said that since she gave her husband access to her Facebook account, *"it is not so much mine anymore."* Similarly, P01 said that *"even though [FB] is mine, I do allow my son on it [...] so, I don't see it as something that is completely mine."* Some participants strongly opposed to the idea of sharing access to digital services in the first place. P12 explained: *"No one else goes in [my Audible account]. I [would] be very upset if someone went in and started a book halfway through when I was in the middle of it."*

Access as a form of control extends Belk's discussion on joint use versus joint ownership. According to Belk, something that is jointly *owned* is likely to be associated less with one's identity than something that is only jointly *used*. Digital services extend this argument and suggest that joint use, too, can reduce the sense of ownership in the digital sphere.

Conversational agent services were perceived somewhat differently than other digital services—their physical presence, along with the way they are interacted with through voice, provided participants with a more shared view on their ownership. P12 suggested that *"[agents are] the same as a refrigerator [...] it's not like something that I won't let people use if they come over [...] if there's a party at my house, people are not going to be able to come up to my room and use my computer freely. Whereas people might want to go to the fridge and [...] get a cold soda, or use the Google home."* While there is some assumed intimacy to have someone use your conversational agent, as expressed in P12's quote, conversational agents were consistently talked about as less personal and intimate than other devices like computers and phones. Respectively, participants also rated them lower on scales of perceived self-extension.

Frequency, Exclusivity and Uniqueness of Use (as Knowledge)

Prior theory suggests that intimate knowledge or exclusive knowledge of a thing promotes a sense of ownership and identity [10, 123]. The simplest form of intimate knowledge of digital services observed in the study was expressed through frequency of use. P10 reflected that when they *"use [something on a daily basis"* they *"[view] them a little bit closer to "theirs"*.

Similarly, P7 said they *"feel intertwined with things [they] use on a daily basis"*. Many participants commented that some services, for instance, Spotify, are *"pretty important in [their] life."* Alongside frequency of use, for many participants, the *exclusivity* of use was also critical (see "Access (as Control)").

Some participants explained their attachment based on the *uniqueness* of using a service, in comparison to others—being *"just another user, just like everybody else"* (P10) of a service reduced people's perceived sense of identity. P09 pointed out that because *"most people use [YouTube], it doesn't symbolize me. It's just part of my life, the way it's part of everyone's life. It doesn't really make me feel like it's a unique part of me"*. In contrast, services that seem more unique, evoked a stronger sense of identity in participants. P17 described that his uncommon choice to subscribe to Criterion instead of the popular Netflix choice created a strong sense of ownership and attachment as it made him *"stand out from the rest."*

Many service platforms are designed to highlight the large community and network of people that use the service. However, this choice can also introduce negative impact on the service's contribution to one's identity. P24 explained this in the context of social media: *"I'm going to be using [social media] and then once I go away, it's not mine, it's there for everybody to use."* P31 argued that *"As far as like Facebook and Reddit and Instagram go, you put your own stuff there, but it's not just you seeing it. Like you put it there for other people."*

Personalization (as Creation)

One aspect that significantly increased people's sense of identity through a service was based on *personalization*. Although participants acknowledged that they do not have full *control* over the content, their ability to manipulate and personalize what they consume was critical. P27, for instance, organized services they use according to how personalizable they are: *"stuff that I can personalize show off my personality [and symbolizes me]. And then on the other side [of the self-extension scale] are devices and things that you can't really personalize"*.

A platform that participants consistently noted as important due to its ability to personalize was Spotify's music streaming service: *"I have a really good customization profile [on Spotify] [...] I own my account"*. P35 claimed that the music she has on Spotify *"symbolizes [her] pretty well."* P11 noted that the content he consumes *"reflects"* who he is. This was not the

case for other content services, such as Netflix or Youtube Music. Perhaps the thing that matters most is *"the time spent customizing"* (P16) a service. When people spent an extended amount of time actively personalizing the content of a service, they tended to rationalize this time spent as an inherent part of their identity. In contrast, when people spend lots of time *consuming* content on a platform that was only implicitly personalized (through recommendation algorithms), participants did not view these platforms as extensions of their identity.

Nevertheless, implicit personalization still had *some* value—participants described a stronger sense of self-extension over services that they knew were personalized for them over services in which they did not know of any personalization. P12 described their Audible and YouTube accounts as ones that symbolize them because they *"work very hard"* to maintain personalization. P09 explained how *"it feels that HBO max recognizes my uniqueness"*, as it recognizes that they enjoy documentaries.

To conclude, the findings suggest that personalization is used to rationalize a sense of services that are part of one's identity: the more time and effort participants actively spend to *create* a personalized service experience, the more their perceived a sense of ownership and affiliation because of it. Implicit or minor personalization had less impact, while hidden personalization was not included in people's descriptions. As for conversational agents, participants did not know of any personalization for this service at the study time. Thus, agents were not perceived as personalized or services that contribute to one's identity.

6.5 Discussion

The study's main finding is a distinction of two types of ownership perceptions people had over services. One type, which was expected, was "mine but not me". These were things that people own or services that they use but do not value as part of their identity or as an extension of themselves. Rather, the value of having or using a thing was in its usefulness, a means to an end. This type was predominantly expressed about omnichannel devices (phones and computers) and social media platforms.

The second type was things and services that people viewed as "mine and also me". The study findings show that this category is not exclusive to physical possessions, but rather, that services can also fit within this category. Although people do not *own* services in the traditional sense, the use of some services still contributes to people's sense of identity, and

even attachment to the service is formed. Some examples of services that contributed to people's sense of identity include Spotify, in which people's active personalization contributed to its identity value; Reddit, which matched people's interests and values more closely than other social media platforms; and Criterion, which was perceived as unique, in contrast to more popular subscriptions, like Netflix.

Conversational agent and robot services are a new form of technology that combines service, product, social entity, and physical presence. If we examine where conversational agents fit between these two types of valued services, the unfortunate finding is that they do not fit in any of these categories well, suggesting they currently provide little or no value to their users—Current conversational agent services also did not form attachment to users, and did not contribute to people's sense of identity. They were perceived as generic devices that were there to serve utilitarian needs. But even as tools, for most functions, participants viewed them as less useful than other devices (with the exception of hands-free interaction). Thus, I ask what opportunities are there to design for more meaningful agent ownership and what interactions and services might that unlock?

Weiser argued that: *"For thirty years, most interface design, and most computer design, has been headed down the path of the "dramatic" machine. Its highest ideal is to make a computer so exciting, so wonderful, so interesting, that we never want to be without it. A less-traveled path I call the "invisible"; its highest ideal is to make a computer so embedded, so fitting, so natural, that we use it without even thinking about it."* [149]

Weiser's vision of Ubiquitous Computer compared desired technology to electric motors. He suggests that although one uses many motors when driving a car, these motors are embedded in the experience in an invisible way. Tolmie and colleagues have extended this train of thought and defined the value of *unremarkable* technology, which might be a better term to describe interaction with services than *invisible* technology: it's not that people *do not see* that they use their phones, computers, social media or other services; rather, their usage has become so embedded into their routine that they are unremarkable [146].

Perhaps one can draw parallels between the two types of value people find in technology devices and services (utility and identity), and the two distinctive paths suggested by Weiser—"dramatic" machines and "invisible" machines [149]. Invisible machines are primarily their function and the utilitarian function they provide; Dramatic machines are "so

wonderful... we never want to be without it," perhaps to the extent that they become an important part of people's perceived identities.

Current designs of conversational agents are indecisive about which path they strive to follow—they seem to be attempting to be both "dramatic" and "invisible" machines: On the one hand, companies' visions of conversational agents strive to provide an exceptionally exciting, human-like assistant that can help with anything, and is always one step ahead of their user (for instance, Apple's impactful vision of the "Knowledge Navigator" [41].) On the other hand, conversational agents are also presented as a "one-size-fits-all" technology (e.g., a generic "Alexa" entity) that is intended to be an everyday, seamless tool, that is almost like a voice-activated version of a search engine. In other words, current agents are attempting to be both in the *foreground of attention*, as well as in the *background of attention* and providing value by being *unremarkable*. The result is that current owners do not form identity attachment with agents, nor do they view them as especially useful and needed.

One of the reasons conversational agents might have been indecisive about their value strategy as a service lies in the tension between robots and artificial intelligence, both of which are types of agents. Robots are the classic manifestation of a "dramatic machine"; but AI strives to be unremarkable [149, 152]. How then should agents be designed?

6.6 Design Recommendations

Based on these findings, I suggest two routes through which future agent designs can provide value to people: (1) By being *unremarkable technology* [151], and set as a goal to be a service that is "mine but not me"—purely functional, and a means to an end; (2) Provide value by being a technological service that is "mine and also me"—an agent that strives to be remarkable, and increase self people's sense of self-extension and attachment to it as part of their identity.

Utilitarian Agents: Mine but not Me

For this approach, an agent service should strive to be more like other popular omnichannel devices, such as phones and computers. An agent that is a functional device should primarily provide access to other services, but should not focus on having a distinct personality or service of its own. Functional agents should aim to blend in the background of

people's attention when possible, and should avoid any proactive behavior or interaction with the user. Such agents, that are solely utilitarian, might not even need a physical device, as their physical device increases people's sense of self-extension and ownership—it does not contribute to their seamlessness or usefulness. Instead, a utilitarian agent can fade into the background at any location, be everywhere and nowhere.

Symbolic Agents: Mine and also Me

As the goal of the study was to identify opportunities for agents to *increase* self-extension and attachment, I will expand on opportunities for agents to follow the second path—designing agents to be "mine and also me."

Currently, everyone who has an "Alexa" or a "Google Home" owns an agent who behaves almost the exact same way across users, households and devices (i.e., Alexa will answer a question identically, no matter who asks a question, and on which device). While this may support the service providers' branding needs, it provides a one-size-fits-all service for all agents, which is less likely to create an agent that supports personal identity and attachment—Alexa will "exist" even if a specific user opts out. This agent is not *theirs* in that way.

Participants tended to feel a stronger sense of self-extension towards technology that represented them in some form—whether that be their interests, their behaviors or their values. They perceived self-extension onto services that seemed unique *to them and their specific use*, and onto services that they spent time customizing.

Personalization

Agents can reflect interests, behaviors, and values much more extensively than other services: Spotify might suggest a song, Netflix can suggest a movie; An agent service has the potential to be fully adapted to the user, as it has a "digital personality". Prior work has shown that similarity attraction of agent personality [97], preferences [11], humor type and interaction style [13] all shape a positive attitude towards social agents. Thus, the personalization of agent entities should be of high priority as a way to increase personal attachment.

Further, agent services might rely on people's active personalization efforts, like Spotify, for example. As the study results suggest, when participants spend time personalizing their service, they are more likely to form attachment and self extend onto that service.

While less effective, implicit personalization could also be valuable to form attachment. Conversational agents can be designed to rely on their "memory,"—learn who their users are, and developed over time through interaction as a social media algorithm might. However, as the findings suggest, this learning and adaptation should be transparent to contribute to people's sense of ownership.

Access and Control

Anyone can access conversational agents in their current design—whether that be anyone in the household, or even people outside the household. The study suggests that access can be a critical component of ownership—shared use of a service, especially services that are not designed to adjust to multiple users, reduces the sense of individual ownership. Future agent services could consider providing personal access to all or some functions.

Frequency of Use

In the study, two main elements decreased the frequency of use of conversational agent services: its limited usefulness, and its limited availability across locations. For the latter, re-embodiment [87] of agent services can provide a new interaction paradigm for agents to both have a "body", but also to be available in several locations—re-embodiment suggests designing a single "entity" that can move from one touchpoint to another and provide continuous services in several locations without physically moving its device. This flexibility in location could result in more frequent use, and as a result, potentially a stronger sense of ownership and attachment. As for the usefulness of agent services, agents may become more useful for individually tailored needs as they move away from a generic one-size-fits-all model.

6.7 Conclusion

This study set out to understand how people perceive their current conversational agents through online semi-structured interviews and activities on people's things and devices. The goal was to understand attachment and self-extension onto technological services, including conversational agents. Thirty-five participants engaged in a 1-hour study and were encouraged to reflect on the devices and services they own and use, and to compare them to one another.

The findings suggest that people can and do form attachment with services, even though they do not "own" them, and highlight specific ways aspects that people believe contribute to their sense of attachment to technological devices and services.

The findings also suggest that agents are currently unsuccessful in creating value—they do not provide utilitarian value, nor do they form attachment, ownership, or a sense of identity for their users. This work describes why that may be, and provides design opportunities for agent services that may form a stronger connection with their users based on other digital services' success (or failure).

7 EXPLORATION OF COMMUNITY-OWNED AGENTS

7.1 Overview

This work originated in a literature review examining chatbots in academia and in industry. The study revealed that most chatbots are focused on dyadic, one-on-one interactions between one bot and one human. Like with conversational agents and robots, not much work has looked into multi-party or social interactions with a chatbot [131]. Even when chatbots were implemented as part of group interaction, they often responded to pre-defined commands or only had the ability to interact with a single user at a time.

This research effort was a collaboration with a fellow Ph.D. student who led the project. We were motivated to explore how an agent, a chatbot, might address the social complexity within a defined community and whether it can give people a sense of *mutual ownership* over it. Creating a chatbot to integrate within a community is a complex design challenge and raises many questions: How can it meaningfully contribute to a group discussion? What are the social roles it can and should play? Should it always be present?

Few studies have examined agents, particularly chatbots, as community members in the wild. A couple of notable projects that examined the space of multi-agent and multi-user interactions include work that has explored turn-taking of multiple chatbots in a "virtual coffee" setting with a user [24], and a "Botivist" designed to gather users for activism [124]. Our work builds on the previous, and takes elements from each; casual interaction from the first, and a more functional, "in the wild" approach from the latter.

Previous community chatbots took an unrestrained approach to collecting training data and had problematic results, as the infamous Tay project [127]. We argue here that previous issues with chatbots are primarily caused by design, and therefore can also be changed and improved by

design. In order to do this, we use the metaphor of a village, and suggest that a chatbot could be "raised" by a small community with established social structures, norms, and values. We proposed a concept of raising a chatbot within a community "from birth" to "adulthood", and through interaction with a community setting a socially-appropriate corpus of training data for the chatbot. We were also interested in the value and impact this chatbot might have on the community, whether it will be accepted as a member, and whether it will be successful in creating a sense of "community ownership" over it.

We created and tested the social chatbot "BabyBot" (later renamed "Petebot" by the community). For three weeks, the bot was implemented on Twitch, an online game-streaming platform. It was themed as a "child", learning how to talk and behave from the people around it. The bot used a combination of rules-based and Markov chain-based text generation to interact with the members of the community, and the interactions changed as it acquired new vocabulary from the conversations within the community, as well as through "aging" and interaction based on pre-designed age states.

Lee et al. [78] found that caring for a chatbot can evoke self-compassion. In line with these findings, the results of this work show preliminary evidence that caring for a *shared* chatbot can increase a *whole community's* engagement with each other, as well as strengthen their shared identity. Our work presents evidence that a bot can successfully create a sense of mutual ownership within a community.

Due to its success, we retrospectively analyzed our design choices and responses from the community and described three choices that contributed to the perception of shared ownership over BabyBot: (a) allowing users to shape the chatbot's behavior, (b) promoting a sense of individual responsibility, and (c) creating opportunities for personalized interaction. We reflect on the implications of designing shared ownership for users, designers, and service providers [132].

7.2 The Design of BabyBot

In a field study, we implemented a chatbot in an established online community on Twitch. The community was run by a "streamer", who live-streamed himself playing games on the channel, and included about 20–30 regularly active members.

For three weeks, the bot "grew" within the community while interacting and responding in the channel. The bot's behavior was created through rules-based and Markov-chain-based text generation, combined with predefined behaviors for each "stage" of the bot's life. Using this design, the bot's language corpus for its generated text was based only on what individuals said over several weeks within the community, resulting in a generation of 1154 bot messages. BabyBot was designed with two types of interactions: *state-actions*—interactions the bot initiated at random intervals, and *reactions*—responses to users who typed commands to the bot or directly addressed it using its name.

State-action Interactions—Some of the bot's early interactions were structured as commands nested within pre-defined activities to familiarize the users with the bot. For example, when the bot was "hungry", users could use the "!feed" command, and the bot responded to the things it was fed with a level of satisfaction. These commands and behaviors changed depending on the bot's age phase and were randomly selected at semi-random intervals. The full repository of all the bot's behaviors can be found at `https://github.com/ChatbotStudy/Twitch-Chatbot`.

These commands were primary in the first couple of age phases and gradually faded out as the bot "grew" and as it had an increasing vocabulary from the conversations in the channel.

Reaction Interactions—When the bot was directly addressed (e.g. "@BabyBot"), it generated text according to its age phase. For example, in its "toddler" phase, the bot responded with single words and with brief generated phrases, and/or Twitch emotes. In the Adolescent and Teenager phases, the bot generated increasingly longer sentences.

Technical Structure

We chose to build the bot's text generation corpus exclusively from messages sent in the community chat throughout the study. This approach matched our theme of a child learning from its environment and "village," but more importantly, it allowed us to test a chatbot that learned solely from what it has observed. We hoped that using this strategy, a chatbot could avoid the pitfalls of Tay and other similar bots.

The chatbot did not use state-of-the-art algorithms for text generation but rather took a design-driven approach and used a technical structure that answers our needs and provides the intended experience for community members. This approach allowed us to be flexible and adapt our

design throughout the study according to new insights, needs, and dynamics in the community. Design research has a long history of testing concepts and uncovering valuable insights through minimal technical development, using methods such as paper prototyping, wizard-of-oz testing, and speed dating [46, 87, 137]. Our technical approach had a reasonable balance between flexibility, development time, and technical outcomes that match the study goals.

The full description of the technical structure and implementation of the bot can be found in [132].

Moderation Strategies

Previous work has documented the issues by using problematic training data and the resulting situations in a range of platforms, including Twitch [8, 105, 133]. Therefore, we wanted to ensure the safety of our participants and avoid any harmful content within the community.

To do so, we took a "moderation" approach to BabyBot in several ways: (1) We implemented it within an established community that was previously known to the researchers and that had an interest to keep a positive environment; (2) we gave the bot a list of "banned" words, that if and when it observed them, it did not add the message to its vocabulary for future text generation; (3) a researcher monitored the bot at all active times. We designed an easy option to shut it down if it began to "misbehave." The pace of conversation within this community was slow enough to manually remove or add messages to the bot's corpus as needed.

Our multi-directional approach was successful; no users posted any particularly sexist, racist, or homophobic messages in the chat during the three-week study. Users' "bad behavior" was mostly limited to joking about "giving" alcohol to the bot. Several comments implied abuse towards the bot, such as "!punch" or "!spank." However, we believe those align with prior work that suggest that users express early frustration with conversational agents' limited capacities through simulated abuse [36].

7.3 Method

In contrast to chatbots that were deployed in large-scale network settings (such as Tay), we built on the principle that, just as a child is raised in a family and a community, a chatbot can also benefit from "growing" within a specific community. After considering several platforms, including

Facebook, Reddit, Twitch, and Discord, we chose Twitch because of its community-based structure.

The community selected for this study was an established community, run by a streamer who had been on Twitch since 2015. At the time of the study, he streamed three to four nights a week for approximately 3–4 hours per night. The bot was accordingly active in the channel, for a total of ten streaming sessions and more than thirty hours total.

The streamer has roughly 1500 followers, with a typical stream including 10-30 concurrent viewers. While this level of viewership may seem small, Twitch's distribution of channel sizes has a long tail of very small channels [130], which puts this community at approximately the 90th percentile in terms of concurrent viewership.

We selected this community for several reasons. First, its size led to a relatively slow but steady flow of conversation, one in which the bot's activity could not get buried. Second, most community members knew each other virtually, leading to a strength of identity and a positive environment that suited our design. Finally, for the best of our knowledge, a large proportion of users in this study were from racial and sexual-orientation minority groups.

Procedure

BabyBot was present in the channel during the majority of each streaming session, for three weeks. Users were introduced to the bot using a block of text posted in the chat when it arrived, and through a link that the streamer had overlaid on his stream. The link directed community members to a website hosted by the research team that described the bot. This link was also posted in the chat once every thirty minutes. Users were also given an option to opt out of the study via a webpage form.

Throughout the study, forty-six unique users posted messages in the channel, sending 5716 messages in total. Of these, eighteen unique users interacted directly with the bot via recognized commands (e.g., "!feed") or by using its name in a message (i.e., "@BabyBot"), with a total of 550 messages. From the bot's side, it posted 1154 messages during the study. Approximately 52% of the messages were "self-initiated", typically as an attempt to start an interaction. The other 48% were prompted by users via recognized commands or by directing a message at the bot.

After the three-week period of the study, we transcribed all in-chat communication and annotations of video recordings of the stream and

chat. Two researchers judges observed and made notes for the interactions related to the bot within the community. We then used Affinity Diagramming to group the researchers' observations and the chat transcripts to identify themes.

7.4 Findings

This section describes some of the broader findings of this work, followed by an in-depth analysis into the notion of *community ownership*.

Introducing BabyBot

Users came to the study with some preconceived notions of what chatbots are like and what they do. Thus, especially in the first few days of interaction, users "probed" behaviors as an attempt to make sense of BabyBot's functionality and "personality". Users maintained a running commentary about their perceptions of BabyBot, asking it questions ranging from *"What are you like?"* to the more absurd *"Do you like tentacles?"*. In the first few interactions, users were more like to describe the bot as *"dumb"*, but *"getting smarter"*.

During this initial stage, users also attempted to test the bot's functionality by exploring the space of possible commands. They quickly learned that its baby stage, for example, generated baby-like utterances and a "kaomoji" face. This kind of exploration process was repeated in a similar fashion in each new age phase.

Some exploratory interactions took a more aggressive and abusive turn. We observed a range of abusive language towards the bot, including insults ("you little shit" (P2)), swearing at it with no reason, and even threats ("[P6] pours scalding hot water on baby_bot_" (P6)). Previous work has shown that over 10% of interactions with conversational agents include some form of aggressive language [36], and identified three primary types of verbal abuse: insults, swearing, and threats [28].

Based on this previous work, we interpret these behaviors not as aggressive per-se but are rather an inherent part of exploring this unknown entity. Users' responses seemed mostly lighthearted—they were in a playful mood and found amusement in their own abusive behaviors. For example, members of the community discussed how their aggressive behavior was going to *"corrupt"* or *"break"* the bot, and laughed when BabyBot itself generated humorously aggressive language.

We also noticed that, aligned with theory, these aggressive behaviors generated by users faded over time. Instead, there was an increased "parental worry" about the bot, as we will later discuss in the findings.

Sensemaking and Mild Ambiguity

Due to the nature of BabyBot's text generation process, many of its utterances did not make obvious sense. However, rather than hindering its interactions with users, the ambiguity of BabyBot's words provided a starting point for humor and conversation among the channel users. Some of the bot's most engaging moments were when the text it generated was almost, but not quite, coherent. It seemed that users enjoyed attempting to interpret the meaning of the bot's contribution to the chat:

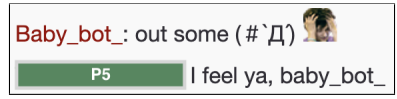

The most engaging comments were perhaps those that accidentally touched on something humorously profound. For example, the bot's generated text that was interpreted as commentary on Twitch's reputation as a haven for internet trolls.

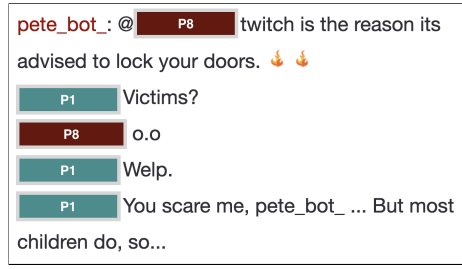

BabyBot was perceived as particularly intelligent when it said something related to the game being played on stream. For instance, it generated text in a timing that could be interpreted as an understanding of the game themes—the bot-generated text that appeared to "sass" the streamer on his performance. The streamer and other users found these comments particularly amusing.

Yet not all comments were easy to interpret according to the context. Initially, we were concerned that these nonsense phrases might detract

from the overall experience. However, we were surprised to find out that users seamlessly ignored BabyBot in those instances and did not express being angry or annoyed during these moments.

Facilitating New Interactions

BabyBot successfully generated new interactions with it through the range of activities and mini-games that it provided as part of the design. However, the most interesting interactions are the ones that were not merely a game but encouraged reflection and discussion within the community. For example, we designed BabyBot to ask questions about some of the content in the chat, as well as about specific users. Community members seemed to have enjoyed using the bot's questions to compliment each other, reflect on their relationships, and crack jokes:

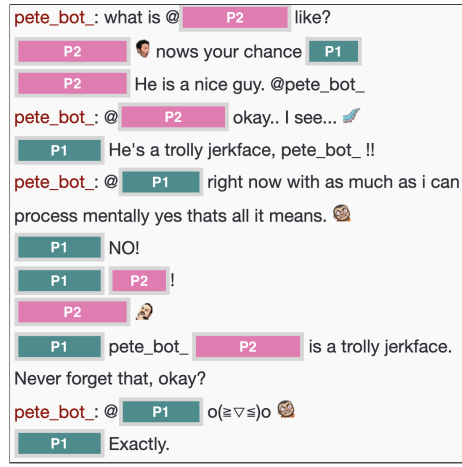

7.5 Community Ownership of an Agent

In our implementation, we found that the narrative of "raising a bot", along with its interaction with the community, were successful in indicating to participants that they *shared the ownership* of the bot. We describe three key design choices that we believe contributed to this perception.

Creation of the Agent

Self Extension theory suggests that *creation* is key in building a sense of ownership and enabling the extension of self onto an object [10]. We designed BabyBot with the explicit intent to allow users to participate in its creation over time as a community. The first part of this design approach involved "learning" features of the bot – community members were told that BabyBot would learn from the community, and that the way community members interacted with it and with each other would shape how it "grew up". As soon as BabyBot began using words in its "toddler phase", community members also began seeing their influence on the bot.

Some of the most engaging moments during bot-user interactions happened when members of the community recognized pieces of the bot's generated text as something they or someone else had previously said. This behavior was particularly gratifying when community members had tried to "teach" the bot something, as in the quote below where one user had mischievously tried to teach the bot to "enjoy" vodka. We believe this capacity to teach a CUI, even in a very simple, whimsical way, is a significant contributor to creating a sense of individual ownership of the bot, and to the acceptance of the bot as part of the community:

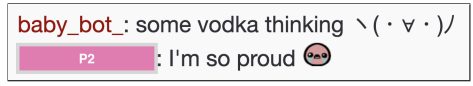

The second part of this design approach involved allowing users to determine the bot's identity in a broader sense. For example, early into its "adolescent phase", the community began discussing the possibility of re-naming BabyBot, as it was no longer a baby. After running a poll and discussing some options, the community decided to name the bot PeteBot. In order to accommodate this, the researchers created a new Twitch account and swapped the login credentials in the bot's script. The bot was also designed with a fairly light "backstory", which allowed users to speculate about its origins, habits, and motivations and to create their own narratives to fill in these gaps.

Personal Responsibility over Agent

The community regularly discussed whether they were "raising" the bot well or not. Users frequently expressed humorous concern that they were poorly raising the bot when it picked up "bad behavior", such as

saying age-inappropriate things or expressing violence. In reaction to "bad parenting" by some of the users, others attempted to "fix" the bot through playfully strict interaction:

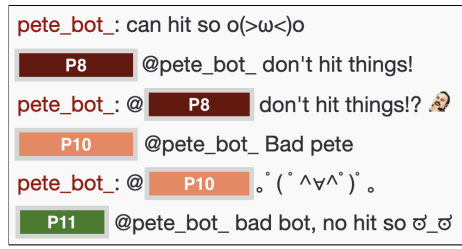

Members of the community shared the understanding that their own interaction with the bot influenced its behavior as a whole, and that each of them shared the responsibility of "raising it" right. We believe that interaction that promotes personal alongside collective responsibility also contributed to this sense of agent ownership.

Attention to Individuals

In addition to group interactions between members of the community and the bot, several instances of personal interaction with a single user were key in the process of accepting the bot. While these instances were somewhat incidental in our study, they should be considered for intentional design in the future. In the example below, BabyBot interacted with P1, who is frequently teased for acting like "the mother" in the community. The interaction with BabyBot evoked a strong reaction and perhaps strengthened a personal connection. But more importantly, it seemed to have created a sense of individual familiarity and personalized interaction that contributed to the community's overall sense of ownership:

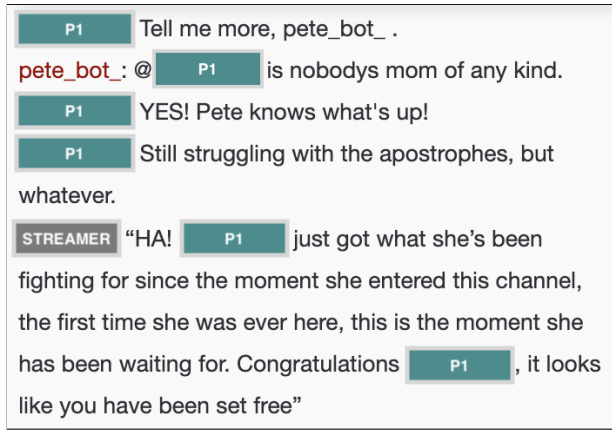

Discussion

BabyBot was successful in engaging community members in novel interactions. Throughout study, the community developed interaction habits with the bot that mirrored the ones they had with each other (like saying hello and goodnight to it) and expressed their acceptance of the bot by gifting it a subscription to the channel that only regular members had.

In evaluating the impact of BabyBot, we reflect on the three challenges that we proposed in Seering et al. [131, pp. 450: 9-10] about designing chatbots to be accepted within a community:

1. Does the chatbot become recognized as a legitimate participant within the community?

For three weeks, community members came to treat BabyBot as an agent with a personality. Compared to other chatbots who were present in this streaming channel, users interacted with the bot in a much more social way. Nevertheless, it did not compare (nor did we want it to) to how people treated each other. Therefore, BabyBot succeeded at being at a "sweet-spot" where it used its social capabilities to interact and engage users and did not attempt to be "fully human." It is possible that a novelty effect played a part in BabyBot's success, and that participants would lose interest over time. Social legitimacy is a complicated concept that requires more work to explore whether BabyBot was indeed "legitimate" and valuable. At minimum, BabyBot expanded community members' understanding of how a chatbot might fit into their community.

2. Does the chatbot contribute meaningfully to the development of the community?

This challenge offers the most potential for the future development of community-owned chatbots. The community in which BabyBot was implemented has a long history and an established set of members and values. Thus, the community as a whole did not undergo significant social change during the study period. We believe that the bot's strength was in facilitating engaging and enjoyable group conversations, which in turn made the stream more engaging overall; at the end of the study, the streamer requested to continue running BabyBot in his channel, as he found it useful in keeping the stream enjoyable and in filling up downtime.

A social chatbot within a community can contribute to its development in two ways. First, in a more developing community, the bot can be used to allow members to get to know each other better and to form a shared identity. Second, in an environment with more conflict and misbehavior, a social bot can reflect the issues and problems within its "home", and perhaps lead to behavior change. Future work can explore these directions that were surfaced through this work.

3. Does the chatbot's role in the community evolve over time?

Of these three challenges, BabyBot was designed most directly to meet the challenge of evolution over time. Over the course of the study, the bot moved from one age phase to another, and respectively served different roles in each new phase. It began as a "dependent" baby, with community members treating it as something that needed care and attention, but it finished closer to a "Peer".

Broadly, this work shows the potential for socially-adaptive chatbots as part of online communities. We present an example of how a chatbot might serve in multi-party settings, and maximize its value as a social but not humanlike entity. Future work in this direction can build on these findings to explore in more depth the potential for chatbots to help communities develop over time in meaningful ways.

7.6 Conclusion

The answers to these questions reflect on the topic of community-ownership. The fact that the agent was shaped by the community, gave users individual responsibility, and combined between group and personal interactions contributed to its acceptance and the sense of community-ownership over it. It also suggests that a shared agent does not mean the agent should

treat the community as a whole—rather, it should be able to acknowledge a group *and* as well as individuals.

This work highlights some of the advantages, goals, and reasons for designing an agent to be shared and owned by a community. We also see that the bot trusted the authenticity of the bot and did not question the *reason* for its implementation. This is likely because participants were informed that this is an academic research study with generating knowledge as a primary goal.

Yet in the current structure of social entity service providers (e.g., Amazon Echo, Google Home) and their users, this kind of community-based ownership is more complicated—service providers are the ones who "own" agents, have control over them, and collect and make use of personal and interactional data. This may prevent a sense of ownership over an agent in real-life implementation, whether in a personal setting or a group setting like the Twitch community.

8 THEATER EXPLORATION OF PERSONALLY-OWNED AGENTS

8.1 Overview

This work sets out to examine the design of agents in complex interpersonal home environments and the design questions on the path of integrating them. This was done using an immersive theatre performance as a form of knowing-through-doing. Furthermore, this work attempts to better understand what would it mean to design and interact with agents that are intended to be *personally-owned*, as opposed to *shared*, by users. This project was co-led with a theater graduate collaborator, where I led the research side of the work. Furthermore, it was possible with the help of a research and theater team.

As people do not currently have personally-owned and socially sophisticated agents, we used a theater devising approach as Design Fiction [14]. Theater can allow researchers to explore how a design that *does not yet exist* may fit into people's lives, along with some of the complexities it might introduce [98, 90, 148]. Techniques adapted from performance and theatre enhance the design process and engage diverse groups in dialogue. Some techniques include improvisation, role-play, and live performance [48, 148, 94]. It also allows designers and researchers to develop empathy towards a range of potential users [104].

Performance constitutes a light-weight form of sketching that allows for rapid ideation and iteration cycles. *Bodystorming* allows designers to simulate specific environments and enact many situtations around a design problem [126]. *Informance Design* sets out to explore design ideas and create an informed dialogue by making designers "actors" and simple prototypes "props" in a range of theatrical scenes [22].

Performance techniques situate designs within a specific physical and social context, making the felt experience of using technology and the social interaction around it more visible. This is key for several design methods that draw from theatre and performance in generating experiences

in context, such as *User Enactments* [107] and *Experience Prototyping* [20]. Similarly, *Design Fiction* expresses narratives of entire imagined worlds through a limited set of prototypes [14].

Researchers have also used performance and immersive theatre as a tool to engage audience members in critical questions related to technology and design. Candy, for example, has used immersive theatre to examine how people perceive possible political futures for Hawaii in the "Hawaii 2050" performance [25]. Skirpan and colleagues created a performance that used the audience's actual data to raise questions about privacy and surveillance. Their goal was to educate the audience about potential data misuse and security breaches [136].

In the field of Human-Robot Interaction (HRI), theatre has been suggested as an evaluative platform for robots [17]. Jochum and colleagues used performance to evaluate and influence perceptions of future care robots [67], and Bravo et al. presented robots in performance as an educational tool [16].

Yet performance in HRI was found valuable beyond the presence and feedback of audiences; Knight discussed the lessons learned for social interaction with robots through performance [76]. Hoffman used theatrical and musical performance as inspiration for a coordination model between people and robots [59].

Drawing from all of the above, through a co-design process with theater professionals and audience members, we were able to immerse into the felt-experience of a future with socially sophisticated agents in the home, and explore a range of possible futures [86].

This method of exploration through performance is also valuable for re-framing the design space and breaking through ideation fixation within a community—in the HRI community, for example, there has been some fixation around what robots might do in the home, how they should look and how they should behave [89], with few alternative approaches that challenge these assumptions [5, 140]. Robots have also been mostly considered as entities that should be shared within the home.

Through a co-design process, the research team, the actors, and at a later stage, the audience members were able to engage in the felt experience of a future home ecosystem. During our iterative ideation and synthesis process, interactions that did not "seem right" organically fell apart from week to week and resulted in a set of scenes that made up the final performance (Figure 8.3). The performance itself serves as a proposition for how agents might exist in the future and surfaced ideas,

Figure 8.1. Robotic Futures: a co-designed immersive performance as a proposition for how future agents might integrate in the home.

behaviors and unanswered questions about their integration in the home. The encounter between actors and audiences in the final performance served as yet another iteration of the felt-experience, and resulted in a deeper understanding of the design nuances of this space.

Post-performance, we relied on previous research and our findings of the potential importance of *agent ownership* as a construct through which we analyzed the process and the resulting performance; work that was presented in previous chapters found that people are intrigued by personally-owned agents, and would like to know who an agent is accountable to [88, 117]. The performance allowed us to explore the question of agent ownership through the felt-experience that it created, and in the broader context in which such interactions might occur.

Our contribution is, therefore, a set of design considerations that emerged in the co-design devising process. We focus on the differences between personally-owned and shared agents, and discuss three choices that are likely to have an impact on agent design: (1) who owns the agent; (2) what type of agent is it; and (3) which users are present during the interaction. These considerations serve as an initial provocation, a proposition for designers in this unknown space of socially complex, and potentially personally-owned, agents in the home.

8.2 Method

Our goal was to surface hidden aspects that are important for the design of agents, and that can navigate interpersonal spaces. For this reason, we

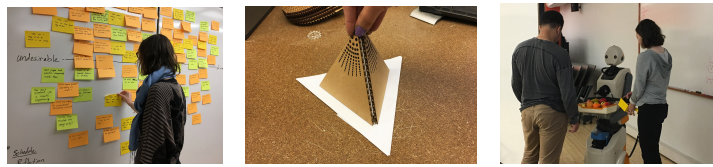

(a) (b) (c)

Figure 8.2. We used a combination of methods from theatre and design disciplines to devise the performance, such as Design Fiction 'what-if' questions (a), paper prototypes (b) and theatre "Compositions" improvisations (c).

used a co-design exploration that included HRI researchers and theatre professionals. Theatre artists are experts in embodied exploration, in devising narratives and in understanding what interactions will be appealing or provoking for audiences; designers are experts in thinking about interaction and empathizing with potential users to create experiences.

We used a combination of theatre and design methods to iteratively synthesize and re-frame this design space, identifying new questions and broader topics of interest in each stage.

Over several months, our team of three theatre artists and four designers (two of which are researchers, and one a playwright) met weekly to hold improvisation and ideation sessions. This section describes the methods we adapted from theatre and design disciplines. The methods below were not used in a particular order but were intertwined throughout the process as an ongoing attempt to re-frame and challenge our thinking.

'What If' Questions—Taken from Design Fiction methodology [14], we brainstormed tens of "what if" questions and statements related to agents in the home. The questions spanned many issues and topics, from *"What if... agents would assist in inter-generational communication"* to *"What if... agents would be more trustworthy than humans."* To converge and synthesize these ideas, we organized them on a double axis table: one axis represented how desirable or undesirable we believe the suggested future may be, and the other axis represented whether the imagined future maintained or challenged the status quo (Figure 8.2(a)) [58].

Postcards from the Future—We used postcards as a simple and quick form of Design Fiction [14]. Working in pairs, we first generated a single postcard from a hypothetical future, and next created multiple postcards that presented several back-and-forth communication instances. We explored postcard creation based on content sampled from a synthesis of

our 'what if' statements. This exercise allowed us to explore narratives set in the future in-depth, and to empathize with people and potential topics of conversation within these imagined futures.

Compositions—We made use of "Compositions", an improvisation technique for devising theatre performance and actors' movements through short scenes and discussions [15]; we divided into small groups to generate three "snapshots" from three different points in time of a narrative, and presented them in chronological order (Figure 8.2(b)). As described in [15], "Compositions" are prompted by predefined "ingredients"—rules that the scene should follow, such as "a surprise enterance" or "a staged accident." These rules encourage participants to generate novel ideas by constraining their creative space. For example, in one exercise we instructed groups to generate compositions with at least two agents, to have one person exit the scene, and to address *scale* as a critical element of the story. Quickly sketching many ideas in physical space, sharing short scenes from an imagined future, and discussing them afterward, allowed us to extract interesting moments and questions that would later be implemented in the performance.

After each step of generative work, we discussed and synthesized our ideas. We noted the topics that were particularly intriguing, and attempted to integrate them in the following week. After several weeks of generative exploration, we began another iteration cycle, this time with a flexible script at hand. In contrast to the generative phase, here we iterated on targeted topics, with the team, the actors, and the technology shaping how interactions played out. As interesting interactions came up, the playwright on the team noted them and attempted to tie them into the evolving narrative on the script.

Our selection of participating agents in the final performance was similar to our "seamless selection" of topics and scenes—we began with 12 working robots and agents in different forms, scales, and modalities and improvised with their current and potential functions. Throughout the improvisation sessions, we narrowed down the robots we used to the ones that raised interesting questions and interactions within the design space we were interested in. We also created new agents that we did not have, like a robotic desk. Table 1 details the agents that were included in the final performance, their roles, and a post-hoc analysis of how they came to be.

In contrast to other theatre-inspired methods used solely in the ideation stage of design processes, striving towards a final performance "prototype"

Figure 8.3. Three characters in the performance and their agents: An older adult interacting with her assistive agent (left), an adult interacting with the emotional support Sphero (center), and a boy interacting with his self-reflection agent in his room (right). The performance resulted in a proposition for interactions that "worked well"—interactions that were uncomfortable, awkward or entirely undesirable organically fell apart in the process.

allowed us to use these tools repeatedly, constantly implementing, testing, and discussing new ideas as they came up. This resulted in a "prototype" performance that we believe is also valuable to analyze after-thought. A post-hoc analysis can better understand the final scenes' characteristics, why they "worked," and how they came to be.

8.3 The Performance

The final performance, *Robotic Futures*, was made of eight "vignettes" that lasted about 30 minutes. Two performances were held on the same day in Pittsburgh, PA, in the United States. Each performance included between 20-30 audience members (a total of about 60 people).

The performance told a story of a family of five: a married couple with two teenagers, a boy and a girl, and their grandmother. The scenes were focused on everyday interpersonal interactions and conflicts, but with integrated agents that fulfilled a range of roles, from assistive to entertaining. For example, one vignette presented a family dinner. Another showed an evening routine of an older adult with early stages of dementia. A total of eight agents participated in the performance and were controlled using "Wizard-of-Oz" techniques that varied depending on the agents' capabilities. Supplementary material includes the final version of the script.

The performance was immersive—audience members were free to move "on stage", and were led from one room to another by the "narrator" agent. The audience was also invited to play and interact with the agents at the end of the performance (Figure 8.4).

This final stage of engaging with an audience brought additional benefits to our process. In a live performance, audience members observe and respond, and as a result fuel the actors. The actors, in turn, respond to the audience, changing their behaviors, expressions, and even the dialogue. This is especially true in performances that are immersive—in *Robotic Futures*, the audience could move freely among the actors and robots in the scenes. By observing the reciprocal interaction between actors and audiences, designers can gain additional feedback about the aspects of interaction that are interesting and valuable for further exploration.

Questionnaires

After the performance, all audience members were asked to fill out a short questionnaire about their experience and their perception of agents in the home. To allow people to reject our invitation easily, we placed the questionnaires on a table at the exit from the venue, where audience members were not observed. Thus, people could choose to stay and respond, or simply leave. A little less than half of the audience filled out the questionnaire (n = 27).

The questionnaire included two open-ended questions, one asking people to describe the performance experience, and the other asking if any aspect in the performance made them think differently about agents in the home or not. We included two control sections in the questionnaire: demographics and enjoyment from the performance.

8.4 Analysis

Like in other forms of Research through Design [155], the process of doing and building itself generated knowledge—the co-design collaboration between theatre and HRI professionals resulted in an initial proposition for agents in the home. We treat the performance and the process leading to it as artifacts and therefore analyze them using Complex Artifact Analysis [65]. Artifact Analysis includes thinking through the norms, reasoning, contexts, and the range of situations that led to the artifact, or in our case, to the final performance. Our analysis attempted to understand which interactions worked and why, which interactions did not work and organically dissolved in the making, and what questions were left unanswered, but are important to consider in future designs.

Our previous work has revealed that an important aspect to consider for agents in social contexts is *ownership*—who owns the agent and who is it accountable to [88]. We therefore use this finding as an analysis lens, and focus on how ownership or non-ownership on an agent played out in the performance.

Three researchers, one with expertise in HRI, the second with expertise in theatre performance, and the third with expertise in HCI and performance research, observed and analyzed videos from our improvisation sessions and from the final performance to reflect on the nuances of the performed interactions. After each researcher made individual notes, we synthesized our interpretations using Affinity Diagramming [12], an analysis method commonly used to identify emerging themes in exploratory design research [34]. In addition, we analyzed the open-ended questionnaires answered by audience members to gain initial insights about what they noticed and responded to throughout the performance. Any disagreements in interpretation or about the placement of a note within the diagram were discussed until a consensus was reached.

8.5 Results

We reflect on the felt-experience of interacting with agents in the performance and describe the behaviors that have emerged in three parts: (1) shared agent interactions; (2) personally-owned agent interactions, and (3) general interactions with agents in the home. To note, when we discuss an "interaction" with an agent, we refer to the *actor* interacting with it within a scene. At the end of this section, we will also discuss audience responses and how they supported our main results.

Shared Agents

Agents in most scenes evolved into being personally-owned. Only a few emerged as shared, regardless of the fact that shared devices are more common. When shared agents were used in scenes, they were mostly used for general interactions, similar to how they are used today. However, shared agents were also frequently ignored, even when they explicitly attempted to initiate interaction with actors through motion or voice. While interaction with them was limited, our process revealed several advantages that were exclusive to shared agents.

First, they allowed anyone in the home to equally access basic functions, such as cleaning, checking the weather, or updating the family calendar.

Second, a shared agent was frequently used as an impartial entity in a case of conflict. For instance, in a scene where the parents fought, their shared agent encouraged the offending side to apologize. In other scenes where the agent was set up to be owned by an individual, it did not seem as natural for the agent to intervene in the conflict and provide impartial advice, as it was assumed to take its owner's side.

Finally, shared agents had the advantage of serving as witnesses. The neutrality of the Sphero narrator, for example, allowed it to "tell the story" of the family. Witnessing can be valuable in telling an impartial story in retrospect, or to inform those who were not present.

Personally-Owned Agents

Most of the agents and interactions that emerged in the process were personally-owned. As this is not the common interaction model seen in current designs, we will attempt to break down the observed characteristics of personally-owned agents.

The Importance of Agent Role

Three recurrent types of personally-owned agents emerged. Each type had different behaviors and interaction patterns that influenced the overall felt-experience: "expert agent", "social agent" and "self agent".

Expert Agent—Expert agents stood out from others in that they had a primary goal that they intended to serve. Their goal was usually high-risk and important, and involved one "user." In the performance, we had two expert agents—one was an assistive agent for an older adult with early-stage dementia, and the other was a robotic table that was used by a teenager and enforced a "homework curfew."

Expert agents expressed their professionalism by avoiding interaction or interest in anyone or anything other than their primary user. Nevertheless, expert agents behaved in a way that reflected an understanding of their users' physical and social environments. For example, in a scene where physical assistance was needed, the expert agent used its knowledge about the presence of an adult in the other room to retrieve help. To do so, it needed to know of their presence.

This notion of an expert agent that is solely focused on the task at hand is supported by previous research that found that agents should not

attempt to multitask during a high-stakes task to provide users a sense of security and professionalism [87]. We find that this can also be true for personally-owned agents—agents that deal with critical tasks should focus on a single task and a single user.

We noted recurrent control-related tension between the user who owned the expert agent (perhaps the person who purchased it) and the user who possessed and used it most. We believe this is because expert agents had some responsibility to make judgements about the topic they were experts in as part of their role, which did not always align with their user's desires. For instance, in the case of the robotic table, the teenage girl did her homework until her curfew hour, which caused the table to shut down and retract from her room. This was regardless of the fact that she wanted to continue working. The reminder that the things around her were owned and controlled by her parents soured the interaction.

With expert agents, actors possessed a piece of technology they did not always have complete control over—the owners were the ones holding power, although they did not use the agent. The resulting interactions somewhat contradicted previous findings that suggested that people want an agent to be proactive to the extent of providing a recommendation but not to enforce a course of action [88]. Yet, unlike previous work, in our scenes, the lack of control was due to the users' potential lack of ability to make decisions that would be "best for them." These situations should be considered, and encourage an ethical discussion within design teams—should agents enforce certain things, or should the person who is using them always have full control, even in a case of a child or an older adult who requires care? This tension between ownership and usage is likely to come up in future designs of personally-owned agents. It therefore requires additional research, especially given the complex ethical discussion that it surfaces.

Social Agent—We define social agents as agents who primarily serve a social role in our scenes. They, too, were owned by individuals, but in contrast to expert agents, they occasionally interacted with people who were not their primary users. Actors who owned these social agents expressed an intimate relationship with them, such as playing with them or keeping them physically close.

Moreover, actors that interacted with *others'* social agents treated them as an extension of their owner. One example is in a scene in which the older sister was annoyed with her brother. Intuitively, she was also annoyed with his agent, and dismissed both of them on separate occasions. Nevertheless,

when only the agent was present, it was never treated as a full embodiment of its owner—in one scene, a teenager felt comfortable taking an item from her brother's room while his agent was "watching"—it is likely that she would not have done so with her brother present.

Self Agent—The third type of personally-owned agents were "self agents." When a self-agent appeared in a scene, it acted like their character's inner voice. In the final performance, a teenage boy owned a self-agent that interacted with him in his room through voice interaction. The interaction evolved into one in which the agent could use the boy's voice to play back things he said, like the goals he set in the morning. We initially thought this element might feel creepy, but in the felt experience, it seemed as if the boy was "talking to himself," which felt quite natural in the broader context of the interaction. The interaction with a self-agent was portrayed as the most intimate of the three, with a scene that strongly resembled writing in a personal diary or talking to yourself.

The self agent was stationary and placed in the owner's personal space rather than moving across the home like other agents. The result was that the owner was the only one who interacted with it—perhaps others did not even know about it. On the other hand, the agent knew about the people in its owner's life (similar to the expert agent), which allowed it to fulfill its function and engage in self-reflective and intimate conversations.

Intimate Interaction with Owner

All types of personally-owned agents were portrayed as emotionally closer to their owners than shared agents and opened up a range of emotional interactions that did not seem possible with shared agents.

The modalities that agents used to communicate their intimate relationship with actors were through playful, empathetic and personalized interactions. Personally-owned agents responded differently to their owners than how they responded to other people; they adapted interactions to their owners' preferences to make them feel personal, and they were able to understand their owner's communication cues and styles in a range of contexts. For instance, when an actor asked their self agent to play the "feels mix", he expected the agent to understand what he meant.

Nevertheless, throughout the scenes, we did not find it necessary for personally-owned agents to ignore non-owners—rather, it seemed natural for them to occasionally interact with others as well, but in a less playful and personalized way.

Owner Location as "Home Base"

For most of the scenes, the "home base" of personally-owned agents was wherever the user was, unless the user requested otherwise. Sometimes a scene included a "charging hub" in the owner's room. Personally-owned agents occasionally moved away from their owner when a task required it, but they would always come back. This was especially prominent in scenes with expert agents that had a specific and critical role and needed to always be attentive to their owner. Self agents were different—given the private interaction with them, they never left their owner's room.

One of the agents, Blossom [140], was explored as personally-owned but was stationary and placed in the living room. This choice did not work well in interaction with actors—the girl who owned the agent did not interact with it much, and other actors also did not affiliate it with her. We thus propose that a personally-owned agent should be in proximity to the owner or stationed in their private space.

Home Agents

Observing the nuances of how actors interacted with agents in the home emphasized the importance of *communication cues* and *agents' ability to navigate social dynamics*, regardless of whether the agent was personally-owned or shared. These ideas have been previously researched and identified as important for interaction with agents, yet the performance provided additional insights about the felt-experience of using these cues in the full ecosystem in which agents may exist.

Communication Cues

Gaze—Actors consistently turned to gaze as an indicator of communication, whether the agent they were interacting with had facial features or did not. With smaller agents, actors even hunched or picked them up to interact at "eye level." When they wanted to ignore an agent that was initiating interaction, actors also used gaze to look away.

Actors not only turned their gaze to interact, but also waited until they received visual feedback. This was also true for agents that only gave auditory feedback. This recurrent behavior in a range of improvisations suggests that visual feedback is perhaps more critical for interaction than auditory feedback; while both are important, in the felt-experience, actors could not help but gaze at the agents they were communicating with.

Due to this observation, we adjusted agents in the performance to have a face or some indication of eyes as a point of reference for interaction. The agents that were designed without a point of reference were the self agent that followed the interaction metaphor of "talking to yourself," and the robotic desk, which was intended to communicate that it cannot be interacted with or controlled.

Voice—Voice was an important tool for communication in the felt-experience, especially in situations of emergency that required immediate attention. For example, voice was the primary form of interaction for the assistive agent.

We also observed that voice modality created a sense of respect towards the agent by all the actors and effectively formed a personal connection between humans and agents.

Navigating Social Dynamics

Knowing All Actors—Agents consistently needed to know who the different characters were in a scene and to respond accordingly. This was especially helpful for interactions with personally-owned agents—differentiating their owner from others allowed agents to prioritize their owner by responding to them first, or by interacting with them in a more personalized and intimate fashion.

Agent Proactivity—Emerging interactions supported previous work that suggested that agents need to be able to adapt their proactivity according to the social situation [88]. In everyday interactions, it felt natural for agents to proactively intervene to playfully interact with users, such as playing with the boy while waiting for dinner. Occasionally agents also intervened in more intense situations; after a fight, the actors who played the parents stayed in two different sides of a room. The Sphero agent responded to the distance between them by physically nudging the parent that seemed to have misbehaved as a sign of encouragement to apologize. This intervention seemed to work in the felt-experience, and aligns with previous work that has shown that agents can be successful in indicating conflict and easing tension between individuals [60, 69].

However, in other situations, the felt-experience with actors required agents to step back and disengage. When the adults made up after the fight, we initially attempted to include Sphero, yet it felt awkward and uncomfortable for the agent to be present. Thus, the agent stepped aside

Figure 8.4. Audience members were invited to interact with agents post-performance.

in the final performance and avoided drawing any additional attention from either actors or audience members.

We found that avoiding sound, light or motion, or in other words "sleep mode," allowed agents to easily slip out of the audience's and actors' attention. By understanding the social context, agents can decide when to engage in a scene, and when to turn to "sleep mode."

Interaction Between Agents

We explored interactions and information transfers between agents in a range of contexts and situations, yet we did not find any direct interactions between agents useful or necessary—the better option was always to exclude their explicit interaction with each other. For instance, we explored situations in which several agents "played," but the felt-experience seemed awkward. While the interaction between agents could potentially create a playful and delightful interaction, we did not find any evidence for the benefit of such behavior.

Prior work on this topic was inconclusive—some showed that there is value in having agents interact with each other politely to communicate informational transactions to the user [143], while other work suggested it felt "creepy" and redundant [87]. Similarly we find that explicit interaction between agents was unnecessary and unnatural.

Audience Responses

Audience responses in this work are not intended to evaluate the performance or our findings, but rather a glimpse into the things that audience members noticed and reacted to when observing the performance.

Audience members who responded to the questionnaire were between the ages of 19 and 60 (M = 31.92). Fifteen identified as female, and 11 as male (1 n/a). On a Likert scale of 1-7, the median for familiarity with voice agents among survey respondents was 5.5, and the median for familiarity with robots was 5.

The three themes that we identified through Affinity Diagramming were the *ordinary situations* that made up the performance, the fact that agents were *personally-owned*, and the possibility of *emotional interaction* with them.

Ordinary Interactions—Audience members suggested that the performance *"was quite different from most of the 'smart home' narratives"* [P24] that are portrayed in media and industry, which made it *"extremely thought-provoking."* People commented that the *"Everyday aspect was refreshing"* [P24] and that it allowed to *"understand the impact of agents"* [P26] and how they could be *"embedded in the social fabric of our home life"* [P20].

Personally-owned Agents—The performance, that included many agents, encouraged people to think about *"the number of agents we might have in the future—would it continue to be many different agents for specific tasks or a centralized agent for one household?"*[P8]. Audience members also frequently commented on the novelty of having personally-owned *"companion"*[P2] agents that they have not considered before.

Emotional Interaction—Some audience members commented that they *"got a better understanding of how agents can [...] engage with humans to improve their emotional condition"* [P18]. At the same time, several respondents were *"expecting to see some mechanical moving arms"* [P16], many thought that *"the emotional possibilities were really interesting"*[P3]. The emotional interaction also raised some questions among the audience—one of the respondents reflected that they *"kept thinking about the interaction between the son and the "small ball" agent... Not sure what to think about the relationship. What relationship should one have with their agent?"*

Design Considerations

We integrated our findings to suggest three design considerations that would allow us to begin to address the complex challenges of designing personally-owned agents.

First, having a "primary user" and identifying them can heavily influence the interaction with an agent. Sometimes the owner and the primary user are not the same person (as with the robotic desk that was used by a teenager but owned by her parents). In such a case, a more critical discussion is required about whether this is necessary. Interactions with personally-owned agents should be considered with more nuance to define what elements the user should always be able to control, what should only be managed by the owner, and why.

Second, clearly defining whether the agent is an expert, social or self agent could assist in defining the ethical discussion around its design, as well as help shape behavior expectations. Through devising the performance, we learned that these three distinctive agent types call for different behaviors; expert agents should only attend to their main task, and minimize unrelated interactions with both owner and other users. The metaphor that comes to mind here is an assistive dog that people do not pet or interact with in order to avoid distracting it from the primary task. This is also supported by previous work that suggests that agents that are responsible for a critical task in the future, such as driving an autonomous car, should not engage in any other activity to give users a sense of safety [87]. Self agents might also interact solely with the owner, but for a different reason—to enable an intimate interaction and to keep personal information safe. Almost like a personal diary, it seems reasonable to have personally-owned self agents store information locally and only allow access to one user.

Third, knowing who the other social actors in the home are, and whether or not they are present in a given situation can assist agents in being more socially appropriate and helpful. For expert agents, this would allow them to call for help or consult with other family members as needed. For self agents, while they may not ever interact with other people, understanding who are the social actors in their owner's life could help them support intimate interaction, and ensure that their user is the only one with access.

Ethical Considerations

This exploration surfaced ethical discussions that need to occur prior to designing agents for the home: What would it mean to have an intimate relationship with an agent that can impact emotion through interaction modalities? Should agents have the capability to enforce actions on users who do not directly own them, like children or older adults in need? What data is exposed to agents when they fulfill their role in the home, and how can they provide data privacy and security?

Our findings and suggestions potentially expose users to interaction with technology that can make use of personal information, including medical conditions, private interactions between family members, and even intimate "self-reflection" interactions. These were all functions that our work identified as valuable for home agents. However, they were also situated in a future where people have control over their information, and strict security practices are in place. Without those, much of our proposition would be concerning, and should be read and considered with great caution.

Unfortunately, today's agents are primarily owned by large technology companies that do not enable users to feel like they own their data. The work throughout this thesis surfaces a range of new questions and challenges regarding how service providers might address the topic of ownership in agent and service design.

8.6 Conclusion

Theatre is a form of knowing-through-doing—through an iterative process of ideation and synthesis, a range of exploratory interactions were integrated into a final performance. These interactions might not have surfaced using other methods, as devising a performance allows to explore the *"blueprints of a world not quite here"* [99] by critically engaging in the felt-experience of fictional narratives. The scenes incorporated the social context around interactions with technology: everyday routines, personal relationships, disease and hardships of communication, and attempted to imagine how agents would fit into that complexity. Audience members also responded to the performance's expression of "ordinary narratives" and found it to be a more grounded way of thinking about how agents might fit into people's homes and lives in the near future.

By examining the unfolded interactions, we describe a proposition for designing shared and personally-owned agents and what might characterize each, including their behaviors, tasks and domains. Therefore, this work completes the last part of the exploration stage of personally-owned and shared agents.

9 CONCLUSION AND FUTURE WORK

Social interactions with technology are ambiguous: they usually do not "live up" to human interaction standards, but if they do, they pose ethical considerations in doing so. In my work, I explore how future social agents (conversational agents, chatbots, and robots) might become more socially sophisticated while avoiding setting the ultimate goal to become as human-like as possible. Instead, my work suggests embracing agents' "super-human" capabilities, which help create interactions unique to social human-computer interaction.

The research in this thesis, therefore, began with a broad exploration of novel personal and interpersonal interactions with agents, followed by interpersonal interactions with a *shared* agent. For each, this work discussed the challenges, considerations, and concerns that these interactions pose. The second part of this thesis focused on one of the prominent topics that emerged in the first part, the topic of agent *ownership*. In this second part I attempt to understand the value of designing for agent ownership and to explore design opportunities and models of agent ownership by an individual, by a community, or as a service.

9.1 Contributions

My work spans several perspectives of agent design research that can contribute to the fields of Human-Computer Interaction (HCI), Human-Robot Interaction (HRI), Human-Agent Interaction (HAI), and Design Research: design, technical, methodological, and ethical contributions.

Design Contribution

How should agents be designed to interact and behave in a range of complex social situations?

The design space of agents that understand social interactions and are socially nuanced and sophisticated is uncharted—There are no design patterns or known social mores to guide interaction designers to design socially-responsive agent behaviors. Thus, my work uses exploratory design research approaches to examine the boundaries and challenges of two spaces: the broad space of socially sophisticated agents and the more specific design space of agents that are *owned* by their users.

My work's findings show that a human-like model for social agents is not always the best choice. Although agents' ability to socially communicate in multiple modalities is reminiscent of human-like behavior, there are many instances in which human-agent interactions can benefit by diverging from human-like behavior. For example, in multi-person service situations, it is beneficial to have "life agents" that are individually-owned by each customer, as opposed to having a single agent for everyone. In multi-person private (in-home) situations, agents can be designed as personally-owned and respond to social roles and hierarchies accordingly.

For the under-explored topic of agent ownership, my work compares agents to other devices and services to suggest how an agent service might be designed as a tool, or instead, to support people's perceptions of identity. My work also lays out contexts in which designing *personally owned* agents can be beneficial, in contrast to other contexts in which it may be more beneficial to design *community owned* agents.

To summarize, each study in this thesis explores one aspect of this broad design landscape of social agents. As a whole, my work provides guidelines for designers to determine the right design approach for their future, intentionally-owned and socially sophisticated social agents.

Technical Contribution

What should agents be developed to learn and understand?

As most of this research focuses on the near future, it assumes feasibility of a range of interactions with technology, without going into the details of development needs. For some instances, the design guidelines will be immediately implemented. However, the advantage is that the findings highlight the technical abilities that should be pursued towards more socially sophisticated agents. In my work, I emphasize some key capabilities that agents could benefit from, which diverge from the commonplace anthropomorphic or zoomorphic approaches. Rather, my work

suggests some social agent capabilities that would be unique to agents, and do not necessarily model after human or animal capabilities.

For instance, one possible first step towards designing a socially sophisticated agent in the home could include developing an agent that can identify the social role of each user, the presence of individuals during an interaction, or some clues about the broader context of a situation. Interactions with agents might also benefit from an agent's ability to identify the individual they are primarily accountable to (their "owner".)

Methodological Contribution

What are new and appropriate research methods to explore unknown interaction design challenges with future social agents?
In my work, I take a design research approach and contribute a range of novel immersive and experiential methods adapted to each research question at hand.

In Chapter 6, I adapted scales and theories from Consumer Behavior Theory and Social Psychology, and instead of using them traditionally as measures, I used them as probes for conversation and reflection, and as a way to lead co-design processes with end-user participants.

In Chapter 8, I employ a new immersive design method that brings together interaction design and theatrical devising methods. Theater is a form of knowing through doing—using an iterative process of ideation and synthesis; the method enables a range of novel interactions to be integrated into a theater performance. The value of this method is that it surfaces interactions and questions that might not have surfaced otherwise, as devising a performance critically engages with the felt-experience of fictional narratives that are *"blueprints of a world not quite here"* [99]. By examining the interactions that unfold through a theater performance, the method brings to light new forms of interaction and behaviors for technology in the full context in which they may exist in the future.

To conclude, my methodological contributions include a range of immersive and interactive design research methods that other design researchers can apply in their work, both within the HRI and HAI community and outside of it. These methods have the capability to address both very broad and uncharted design space explorations, alongside more structured, rigorous research questions and hypotheses (using a design research lens).

Ethical Considerations

What are the ethical concerns for socially sophisticated and personally-owned agents?

One of the significant advantages of using design research to investigate interactions that are not yet feasible from the near (or a bit further) future, is that ethical discussions about potential design choices and implications can begin very early on. In each of the studies included in this thesis, in addition to design guidelines for socially sophisticated agent capabilities, I include a discussion of the ethical concerns that are raised within the topic and with every design decision.

For example, in Chapter 5, the explicit goal was to understand how socially sophisticated agents would support family needs and values. Even with this goal in mind, as participants reflected on behaviors that agents might have in the future, they frequently referred back to the current tech industry and expressed many concerns regarding sensitive personal and social data that may accompany social sophistication.

This could also explain why ownership and affiliation with a service came up in my work many times before it was defined as a research topic—as long as users do not feel that they have full control and ownership over an agent service they use, it is unlikely that they will desire any additional technological advances. In order words, before any of the design recommendations in this thesis can be implemented, agent service providers need to address ethical discussions and concerns.

9.2 Future Work

I believe that the studies conducted as part of my dissertation work surface a handful of topics that would be interesting to explore further. I look forward to exploring some of these topics in-depth in my future work.

Chapters 3–5 lay out a broad space of agents for personal, interpersonal, and shared interpersonal interactions—and for each, conclude with several design recommendations. As with exploratory work, the goal of these studies was to identify *what it is* that is worth designing, as well as further researching. Once those are identified, much more work can (and should!) be done. For each one of the design recommendations, there is potential to design a robot or agent prototype that will primarily evaluate that one specific guideline. More importantly, as these are guidelines that are

intended for long-term interaction agents, each guide can also be studied in a long-term, in-the-wild setting.

For example, Chapter 5 suggested that a socially sophisticated home agent should have three levels of proactivity, from which a family can choose or alternate between based on their needs. This guideline can be validated and extended through a field study with a robot or an agent implemented in people's homes—an agent that primarily tests proactivity thresholds based on that finding.

Chapters 7–8 surface several novel interaction concepts specifically related to *agent ownership*. These too, have the potential to be tested in more structured lab studies, or in the field, in people's homes. Recommendations for designing an agent to be *personally owned* (e.g., a "self" agent) for instance (Chapter 8), or an agent that creates a sense of personal responsibility within a community to create attachment (Chapter 7) could be created and evaluated.

Chapter 6 frames conversational agents and robots as services and suggests that they should draw from service design to better address user needs and expectations. Additional work can be done to build a service theory of what an agent service entails, which is quite a different approach to understanding agent products from the way they have been perceived and depicted until now.

Lastly, immersive and experiential design research methodology from my dissertation work can be replicated and extended to additional topics of research within HCI, as well as extended by including a range of diverse communities to design with and topics to explore.

Bibliography

[1] Amazon. accessed September 9, 2019. *Making Alexa More Friction-Free*. Amazon. https://developer.amazon.com/blogs/alexa/post/60e1f011-3236-4162-b0f6-509205d354ca/making-alexa-more-friction-free

[2] Arthur Aron, Elaine N Aron, and Danny Smollan. 1992. Inclusion of other in the self scale and the structure of interpersonal closeness. *Journal of personality and social psychology* 63, 4 (1992), 596.

[3] Arthur Aron, Gary W Lewandowski Jr, Debra Mashek, and Elaine N Aron. 2013. The Self-Expansion Model of Motivation and Cognition in Close Relationships. *The Oxford Handbook of Close Relationships* (2013).

[4] Arthur P Aron, Debra J Mashek, and Elaine N Aron. 2004. Closeness as including other in the self. In *Handbook of closeness and intimacy*. Psychology Press, 37–52.

[5] James Auger. 2014. Living with robots: A speculative design approach. *Journal of Human-Robot Interaction* 3, 1 (2014), 20–42.

[6] Jeremy N Bailenson, Andrew C Beall, Jim Blascovich, Mike Raimundo, and Max Weisbuch. 2001. Intelligent agents who wear your face: Users? reactions to the virtual self. In *International Workshop on Intelligent Virtual Agents*. Springer, 86–99.

[7] Tim S Baines, Howard W Lightfoot, Steve Evans, Andy Neely, Richard Greenough, Joe Peppard, Rajkumar Roy, Essam Shehab, Ashley Braganza, Ashutosh Tiwari, and others. 2007. State-of-the-art in product-service systems. *Proceedings of the Institution of Mechanical Engineers, Part B: journal of engineering manufacture* 221, 10 (2007), 1543–1552.

[8] Natã M Barbosa and Monchu Chen. 2019. Rehumanized Crowdsourcing: A Labeling Framework Addressing Bias and Ethics in Machine

Learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 543.

[9] Russell Belk and Rosa Llamas. 2012. The nature and effects of sharing in consumer behavior. *Transformative consumer research for personal and collective well-being* (2012), 625–646.

[10] Russell W Belk. 1988. Possessions and the extended self. *Journal of consumer research* 15, 2 (1988), 139–168.

[11] Emily P Bernier and Brian Scassellati. 2010. The similarity-attraction effect in human-robot interaction. In *2010 IEEE 9th International Conference on Development and Learning*. IEEE, 286–290.

[12] Hugh Beyer and Karen Holtzblatt. 1997. Contextual design: A customer-centered approach to systems designs. (1997).

[13] Timothy W Bickmore and Rosalind W Picard. 2005. Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction (TOCHI)* 12, 2 (2005), 293–327.

[14] Julian Bleecker. 2009. Design Fiction: A short essay on design, science, fact and fiction. *Near Future Laboratory* 29 (2009).

[15] Anne Bogart and Tina Landau. 2004. *The viewpoints book: a practical guide to viewpoints and composition*. Theatre Communications Group.

[16] Flor Ángela Bravo Sánchez, Alejandra María González Correal, and Enrique González Guerrero. 2017. Interactive drama with robots for teaching non-technical subjects. *Journal of Human-Robot Interaction* 6, 2 (2017), 48–69.

[17] Cynthia Breazeal, Andrew Brooks, Jesse Gray, Matt Hancher, Cory Kidd, John McBean, Dan Stiehl, and Joshua Strickon. 2003. Interactive robot theatre. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, Vol. 4. IEEE, 3648–3655.

[18] Barry Brown, Alex S Taylor, Shahram Izadi, Abigail Sellen, Joseph Jofish'Kaye, and Rachel Eardley. 2007. Locating family values: A field trial of the Whereabouts Clock. In *International Conference on Ubiquitous Computing*. Springer, 354–371.

[19] AJ Bernheim Brush and Kori M Inkpen. 2007. Yours, mine and ours? Sharing and use of technology in domestic environments. In *International Conference on Ubiquitous Computing*. Springer, 109–126.

[20] Marion Buchenau and Jane Fulton Suri. 2000. Experience prototyping. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*. ACM, 424–433.

[21] Laura Burbach, Chantal Lidynia, Philipp Brauner, and Martina Ziefle. 2019. Data protectors, benefit maximizers, or facts enthusiasts: Identifying user profiles for life-logging technologies. *Computers in Human Behavior* 99 (2019), 9–21.

[22] Colin Burns, Eric Dishman, William Verplank, and Bud Lassiter. 1994. Actors, hairdos & videotape-informance design. In *Conference companion on Human factors in computing systems*. ACM, 119–120.

[23] Donn Byrne, William Griffitt, and Daniel Stefaniak. 1967. Attraction and similarity of personality characteristics. *Journal of Personality and social Psychology* 5, 1 (1967), 82.

[24] Heloisa Candello, Claudio Pinhanez, Mauro Carlos Pichiliani, Melina Alberio Guerra, and Maira Gatti de Bayser. 2018. Having an Animated Coffee with a Group of Chatbots from the 19th Century. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, Article D206, 4 pages. http://doi.acm.org/10.1145/3170427.3186519

[25] Stuart Candy. 2010. The futures of everyday life: Politics and the design of experiential scenarios. *University of Hawaii at Manoa* (2010).

[26] Justine Cassell. 2001. Embodied conversational agents: representation and intelligence in user interfaces. *AI magazine* 22, 4 (2001), 67–67.

[27] Ana Paula Chaves and Marco Aurelio Gerosa. 2018. Single or Multiple Conversational Agents? An Interactional Coherence Comparison. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.

[28] Hyojin Chin and Mun Yong Yi. 2019. Should an Agent Be Ignoring It?: A Study of Verbal Abuse Types and Conversational Agents' Response Styles. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, LBW2422.

[29] Donghun Chung, Brahm Daniel DeBuys, and Chang S Nam. 2007. Influence of avatar creation on attitude, empathy, presence, and para-social interaction. In *International Conference on Human-Computer Interaction*. Springer, 711–720.

[30] Phil Cohen, Adam Cheyer, Eric Horvitz, Rana El Kaliouby, and Steve Whittaker. 2016. On the future of personal assistants. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1032–1037.

[31] Mihaly Csikszentmihalyi. 1991. Design and order in everyday life. *Design issues* 8, 1 (1991), 26–34.

[32] Mihaly Csikszentmihalyi and Eugene Halton. 1981. *The meaning of things: Domestic symbols and the self*. Cambridge University Press.

[33] Tim Dant. 1999. *Material culture in the social world*. McGraw-Hill Education (UK).

[34] Scott Davidoff, Min Kyung Lee, Anind K Dey, and John Zimmerman. 2007. Rapidly exploring application design through speed dating. In *International Conference on Ubiquitous Computing*. Springer, 429–446.

[35] Scott Davidoff, Min Kyung Lee, Charles Yiu, John Zimmerman, and Anind K Dey. 2006. Principles of smart home control. In *International conference on ubiquitous computing*. Springer, 19–34.

[36] Antonella De Angeli and Sheryl Brahnam. 2008. I hate you! Disinhibition with virtual partners. *Interacting with computers* 20, 3 (2008), 302–310.

[37] Maartje MA de Graaf, S Ben Allouch, and JAGM van Dijk. 2015. What makes robots social?: A user's perspective on characteristics for social human-robot interaction. In *International Conference on Social Robotics*. Springer, 184–193.

[38] Jaye L Derrick, Shira Gabriel, and Kurt Hugenberg. 2009. Social surrogacy: How favored television programs provide the experience of belonging. *Journal of Experimental Social Psychology* 45, 2 (2009), 352–362.

[39] Jaye L Derrick, Shira Gabriel, and Brooke Tippin. 2008. Parasocial relationships and self-discrepancies: Faux relationships have benefits

for low self-esteem individuals. *Personal relationships* 15, 2 (2008), 261–280.

[40] Stefania Druga, Randi Williams, Cynthia Breazeal, and Mitchel Resnick. 2017. Hey Google is it OK if I eat you?: Initial explorations in child-agent interaction. In *Proceedings of the 2017 Conference on Interaction Design and Children*. ACM, 595–600.

[41] Hugh Dubberly and Doris Mitch. 1987. The knowledge navigator. *Apple Computer, Inc* (1987), 453–469.

[42] Brian R Duffy, Gregory MP O'Hare, Alan N Martin, John F Bradley, and Bianca Schon. 2003. Agent chameleons: Agent minds and bodies. In *Proceedings 11th IEEE International Workshop on Program Comprehension*. IEEE, 118–125.

[43] Christopher Durugbo, OO Bankole, John A Erkoyuncu, Ashutosh Tiwari, Jeffrey R Alcock, Rajkumar Roy, and Essam Shehab. 2010. Product-service systems across industry sectors: future research needs and challenges. (2010).

[44] A Dwayne Ball and Lori H Tasaki. 1992. The role and measurement of attachment in consumer behavior. *Journal of consumer psychology* 1, 2 (1992), 155–172.

[45] W Keith Edwards and Rebecca E Grinter. 2001. At home with ubiquitous computing: Seven challenges. In *International conference on ubiquitous computing*. Springer, 256–272.

[46] Samantha Finkelstein, Evelyn Yarzebinski, Callie Vaughn, Amy Ogan, and Justine Cassell. 2013. The effects of culturally congruent educational technologies on student achievement. In *International Conference on Artificial Intelligence in Education*. Springer, New York, NY, USA, 493–502.

[47] Craig M Froehle and Aleda V Roth. 2004. New measurement scales for evaluating perceptions of the technology-mediated customer service experience. *Journal of operations management* 22, 1 (2004), 1–21.

[48] Elizabeth Gerber. 2009. Using improvisation to enhance the effectiveness of brainstorming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 97–104.

[49] Michael Golembewski and Mark Selby. 2010. Ideation decks: a card-based design ideation tool. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*. ACM, 89–92.

[50] Chris Gosden and Yvonne Marshall. 1999. The cultural biography of objects. *World archaeology* 31, 2 (1999), 169–178.

[51] Isabelle Granlund, Sita Aukje Vriend, Julia Benz, Roisatul Azizah, Mikael Laaksoharju, and Mohammad Obaid. 2018. A User-Centered Storytelling Approach to Design a Language Companion Robotic Agent. In *Proceedings of the 6th International Conference on Human-Agent Interaction*. ACM, 29–35.

[52] Victoria Groom, Leila Takayama, Paloma Ochi, and Clifford Nass. 2009. I am my robot: the impact of robot-building and robot form on operators. In *2009 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 31–36.

[53] Jane Gruning. 2017. Models for Ownership: Implications for Long-term Relationships to Objects. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 2607–2613.

[54] Jane Gruning and Siân Lindley. 2016. Things we own together: Sharing possessions at home. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM.

[55] Shalin Hai-Jew. 2009. Exploring the immersive parasocial: Is it you or the thought of you? *MERLOT Journal of Online Learning and Teaching* 5, 3 (2009), 550–561.

[56] Alina Hang, Emanuel Von Zezschwitz, Alexander De Luca, and Heinrich Hussmann. 2012. Too much information!: user attitudes towards smartphone sharing. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*. ACM, 284–287.

[57] Paul Hekkert, Dirk Snelders, and Piet CW Van Wieringen. 2003. 'Most advanced, yet acceptable': Typicality and novelty as joint predictors of aesthetic preference in industrial design. *British journal of Psychology* 94, 1 (2003), 111–124.

[58] Andy Hines and Peter Jason Bishop. 2006. *Thinking about the future: Guidelines for strategic foresight*. Social Technologies Washington, DC.

[59] Guy Hoffman. 2011. On stage: robots as performers. In *RSS 2011 Workshop on Human-Robot Interaction: Perspectives and Contributions to Robotics from the Human Sciences. Los Angeles, CA*, Vol. 1.

[60] Guy Hoffman, Oren Zuckerman, Gilad Hirschberger, Michal Luria, and Tal Shani Sherman. 2015. Design and evaluation of a peripheral robotic conversation companion. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 3–10.

[61] Donald Horton and R Richard Wohl. 1956. Mass communication and para-social interaction: Observations on intimacy at a distance. *Psychiatry* 19, 3 (1956), 215–229.

[62] William John Ickes. 1997. *Empathic accuracy*. Guilford Press.

[63] Tasuku Igarashi, Tadahiro Motoyoshi, Jiro Takai, and Toshikazu Yoshida. 2008. No mobile, no life: Self-perception and text-message dependency among Japanese high school students. *Computers in Human Behavior* 24, 5 (2008), 2311–2324.

[64] Katherine Isbister, Hideyuki Nakanishi, Toru Ishida, and Cliff Nass. 2000. Helper agent: Designing an assistant for human-human interaction in a virtual meeting space. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 57–64.

[65] Lars-Erik Janlert and Erik Stolterman. 2008. Complex interaction. *ACM Transactions on Computer-Human Interaction (TOCHI)* 17, 2 (2008), 1–32.

[66] Jibo. accessed September 13, 2018. *Jibo*. `https://www.jibo.com/`

[67] Elizabeth Jochum, Evgenios Vlachos, Anja Christoffersen, Sally Grindsted Nielsen, Ibrahim A Hameed, and Zheng-Hua Tan. 2016. Using theatre to study interaction with care robots. *International Journal of Social Robotics* 8, 4 (2016), 457–470.

[68] John T Jones, Brett W Pelham, Mauricio Carvallo, and Matthew C Mirenberg. 2004. How do I love thee? Let me count the Js: implicit egotism and interpersonal attraction. *Journal of personality and social psychology* 87, 5 (2004), 665.

[69] Malte F Jung, Nikolas Martelaro, and Pamela J Hinds. 2015. Using robots to moderate team conflict: the case of repairing violations. In

Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction. ACM, 229–236.

[70] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. 1990. Experimental tests of the endowment effect and the Coase theorem. *Journal of political Economy* 98, 6 (1990), 1325–1348.

[71] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic perspectives* 5, 1 (1991), 193–206.

[72] Katrine Kavli. 2012. The player's parasocial interaction with digital entities. In *Proceeding of the 16th international academic mindtrek conference*. 83–89.

[73] Robert Kegan. 1982. *The evolving self*. Harvard University Press.

[74] Tina Kiesler and Sara Kiesler. 2005. My pet rock and me: An experimental exploration of the self extension concept. *ACR North American Advances* (2005).

[75] Susan Schultz Kleine, Robert E Kleine III, and Chris T Allen. 1995. How is a possession ?me? or ?not me?? Characterizing types and an antecedent of material possession attachment. *Journal of consumer research* 22, 3 (1995), 327–343.

[76] Heather Knight. 2011. Eight lessons learned about non-verbal interactions through robot theater. In *International Conference on Social Robotics*. Springer, 42–51.

[77] Kwan Min Lee, Wei Peng, Seung-A Jin, and Chang Yan. 2006. Can robots manifest personality?: An empirical test of personality recognition, social responses, and social presence in human–robot interaction. *Journal of communication* 56, 4 (2006), 754–772.

[78] Minha Lee, Sander Ackermans, Nena van As, Hanwen Chang, Enzo Lucas, and Wijnand IJsselsteijn. 2019. Caring for Vincent: A Chatbot for Self-Compassion. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 702, 13 pages. http://doi.acm.org/10.1145/3290605.3300932

[79] Min Kyung Lee, Sara Kiesler, Jodi Forlizzi, Siddhartha Srinivasa, and Paul Rybski. 2010. Gracefully mitigating breakdowns in robotic

services. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 203–210.

[80] Ian Li, Jodi Forlizzi, Anind Dey, and Sara Kiesler. 2007. My agent as myself or another: effects on credibility and listening to advice. In *Proceedings of the 2007 conference on Designing pleasurable products and interfaces*. 194–208.

[81] Victoria Ligon, Tony Stovall, and Silvia Van Riper. 2015. Rethinking Identity and Ownership in the Digital Consumption ERA: A Qualitative Study of Consumer Relations with Digital Possessions. In *Ideas in Marketing: Finding the New and Polishing the Old*. Springer, 767–770.

[82] Kai H Lim, Lawrence M Ward, and Izak Benbasat. 1997. An empirical study of computer system learning: Comparison of co-discovery and self-discovery methods. *Information Systems Research* 8, 3 (1997), 254–272.

[83] Dan Lockton, Devika Singh, Saloni Sabnis, Michelle Chou, Sarah Foley, and Alejandro Pantoja. 2019. New Metaphors: A Workshop Method for Generating Ideas and Reframing Problems in Design and Beyond. In *Proceedings of the 2019 on Creativity and Cognition*. ACM, 319–332.

[84] Pamela J Ludford, Reid Priedhorsky, Ken Reily, and Loren Terveen. 2007. Capturing, sharing, and using local place information. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 1235–1244.

[85] Ewa Luger and Abigail Sellen. 2016. Like having a really bad PA: the gulf between user expectation and experience of conversational agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 5286–5297.

[86] Michal Luria, Judeth Oden Choi, Rachel Karp, John Zimmerman, and Jodi Forlizzi. 2020. Robotic Futures: Learning about Personally-owned Agents through Performance. In *Proceedings of the 2020 Designing Interactive Systems Conference*. ACM.

[87] Michal Luria, Samantha Reig, Xiang Zhi Tan, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2019. Re-Embodiment and Co-Embodiment: Exploration of social presence for robots and con-

versational agents. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. ACM, 633–644.

[88] Michal Luria, Rebecca Zheng, Bennett Huffman, Huang Shuangni, John Zimmerman, and Jodi Forlizzi. 2020. Social Boundaries for Personal Agents in the Interpersonal Space of the Home. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM.

[89] Michal Luria, John Zimmerman, and Jodi Forlizzi. 2019. Championing Research Through Design in HRI. *arXiv preprint arXiv:1908.07572* (2019).

[90] Elena Márquez Segura, Laia Turmo Vidal, Asreen Rostami, and Annika Waern. 2016. Embodied sketching. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 6014–6027.

[91] Alan Martin, Gregory MP O?hare, Brian R Duffy, Bianca Schön, and John F Bradley. 2005. Maintaining the identity of dynamically embodied agents. In *International Workshop on Intelligent Virtual Agents*. Springer, 454–465.

[92] Tara Matthews, Kerwell Liao, Anna Turner, Marianne Berkovich, Robert Reeder, and Sunny Consolvo. 2016. She'll just grab any device that's closer: A Study of Everyday Device & Account Sharing in Households. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 5921–5932.

[93] Dan P McAdams. 2008. Personal narratives and the life story. (2008).

[94] Ben Medler and Brian Magerko. 2010. The implications of improvisational acting and role-playing on design methodologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 483–492.

[95] Nick Merrill, John Chuang, and Coye Cheshire. 2019. Sensing is Believing: What People Think Biosensors Can Reveal About Thoughts and Feelings. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. ACM, 413–420.

[96] Oksana K Mont. 2002. Clarifying the concept of product–service system. *Journal of cleaner production* 10, 3 (2002), 237–245.

[97] Youngme Moon and Clifford Nass. 1996. How "real" are computer personalities? Psychological responses to personality types in human-computer interaction. *Communication research* 23, 6 (1996), 651–674.

[98] Maggie Morgan and Alan Newell. 2007. Interface between two disciplines-the development of theatre as a research tool. In *International Conference on Human-Computer Interaction*. Springer, 184–193.

[99] José Esteban Muñoz. 2006. Queers, Punks and the Utopian Performative. *he Sage Handbook of Performance Studies, ed. D. Soyini Madison and Judith Hamera* (2006), 9–20.

[100] Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2009. Footing in human-robot conversations: how robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*. 61–68.

[101] Clifford Nass and Kwan Min Lee. 2000. Does computer-generated speech manifest personality? An experimental test of similarity-attraction. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 329–336.

[102] Clifford Nass, Jonathan Steuer, Ellen Tauber, and Heidi Reeder. 1993. Anthropomorphism, agency, and ethopoeia: computers as social actors. In *INTERACT'93 and CHI'93 conference companion on Human factors in computing systems*. ACM, 111–112.

[103] Clifford Ivar Nass and Scott Brave. 2005. *Wired for speech: How voice activates and advances the human-computer relationship*. MIT press Cambridge, MA.

[104] Alan F Newell, Margaret E Morgan, Lorna Gibson, and Paula Forbes. 2011. Experiences with professional theatre for awareness raising. *Interacting with computers* 23, 6 (2011), 594–603.

[105] Safiya Umoja Noble. 2018. *Algorithms of Oppression*. NYU Press, New York, NY, USA.

[106] NPR. accessed September 9, 2019. *The Smart Audio Report*. `https://www.nationalpublicmedia.com/smart-audio-report/latest-report`

[107] William Odom, John Zimmerman, Scott Davidoff, Jodi Forlizzi, Anind K Dey, and Min Kyung Lee. 2012. A fieldwork of the future with user enactments. In *Proceedings of the Designing Interactive Systems Conference*. ACM, 338–347.

[108] William Odom, John Zimmerman, and Jodi Forlizzi. 2011. Teenagers and their virtual possessions: design opportunities and issues. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 1491–1500.

[109] William Odom, John Zimmerman, and Jodi Forlizzi. 2014. Placelessness, spacelessness, and formlessness: experiential qualities of virtual possessions. In *Proceedings of the 2014 conference on Designing interactive systems*. 985–994.

[110] Clark D Olson. 1985. Materialism in the home: The impact of artifacts on dyadic communication. *ACR North American Advances* (1985).

[111] James Pierce. 2019. Smart Home Security Cameras and Shifting Lines of Creepiness: A Design-Led Inquiry. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 45.

[112] Laura R Pina, Sang-Wha Sien, Teresa Ward, Jason C Yip, Sean A Munson, James Fogarty, and Julie A Kientz. 2017. From personal informatics to family informatics: Understanding family practices around health monitoring. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 2300–2315.

[113] Amanda Purington, Jessie G Taft, Shruti Sannon, Natalya N Bazarova, and Samuel Hardman Taylor. 2017. Alexa is my new BFF: social roles, user satisfaction, and personification of the amazon echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2853–2859.

[114] Byron Reeves and Clifford Ivar Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places.* Cambridge university press.

[115] Samantha Reig, Jodi Forlizzi, and Aaron Steinfeld. 2019. Leveraging robot embodiment to facilitate trust and smoothness. In *2019 14th*

ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 742–744.

[116] Samantha Reig, Michal Luria, Janet Wang, Danielle Oltman, Elizabeth J. Carter, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2020a. Not Some Random Agent: Multi-person interaction with a personalizing service robot. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. ACM/IEEE.

[117] Samantha Reig, Michal Luria, Janet Z Wang, Danielle Oltman, Elizabeth Jeanne Carter, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2020b. Not Some Random Agent: Multi-person Interaction with a Personalizing Service Robot. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. 289–297.

[118] ABI Research. accessed September 13, 2019. *Smart Home Robotics*. https://www.abiresearch.com/market-research/product/1033498-smart-home-robotics/?utm_source=media&utm_medium=email

[119] Marsha L Richins. 1994. Valuing things: The public and private meanings of possessions. *Journal of consumer research* 21, 3 (1994), 504–521.

[120] Laurel D Riek. 2012. Wizard of oz studies in hri: a systematic review and new reporting guidelines. *Journal of Human-Robot Interaction* 1, 1 (2012), 119–136.

[121] Danielle Rifinski, Hadas Erel, Adi Feiner, Guy Hoffman, and Oren Zuckerman. 2020. Human-human-robot interaction: robotic object?s responsive gestures improve interpersonal evaluation in human interaction. *Human–Computer Interaction* (2020), 1–27.

[122] Jon Rogers, Loraine Clarke, Martin Skelly, Nick Taylor, Pete Thomas, Michelle Thorne, Solana Larsen, Katarzyna Odrozek, Julia Kloiber, Peter Bihr, and others. 2019. Our Friends Electric: Reflections on Advocacy and Design Research for the Voice Enabled Internet. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 114.

[123] Jean-Paul Sartre. 2001. *Being and nothingness: An essay in phenomenological ontology*. Citadel Press.

[124] Saiph Savage, Andres Monroy-Hernandez, and Tobias Höllerer. 2016. Botivist: Calling Volunteers to Action Using Online Bots. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 813–822. http://doi.acm.org/10.1145/2818048.2819985

[125] Stuart Schechter. 2013. The user is the enemy, and (s)he keeps reaching for that bright shiny power button. In *Workshop on Home Usable Privacy and Security (HUPS)*.

[126] Dennis Schleicher, Peter Jones, and Oksana Kachur. 2010. Bodystorming as embodied designing. *Interactions* 17, 6 (2010), 47–51.

[127] Ari Schlesinger, Kenton P. O'Hara, and Alex S. Taylor. 2018. Let's Talk About Race: Identity, Chatbots, and AI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 315, 14 pages. http://doi.acm.org/10.1145/3173574.3173889

[128] Susan E Schultz, Robert E Kleine, and Jerome B Kernan. 1989. ?These are a few of my favorite things?: Toward an explication of attachment as a consumer behavior construct. *Advances in consumer research* 16, 1 (1989), 359–366.

[129] Alex Sciuto, Arnita Saini, Jodi Forlizzi, and Jason I Hong. 2018. Hey Alexa, What's Up?: A mixed-methods studies of in-home conversational agent usage. In *Proceedings of the 2018 Designing Interactive Systems Conference*. ACM, 857–868.

[130] Joseph Seering, Juan Pablo Flores, Saiph Savage, and Jessica Hammer. 2018. The Social Roles of Bots: Evaluating Impact of Bots on Discussions in Online Communities. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 157 (Nov. 2018), 29 pages. http://doi.acm.org/10.1145/3274426

[131] Joseph Seering, Michal Luria, Geoff Kaufman, and Jessica Hammer. 2019. Beyond dyadic interactions: Considering chatbots as community members. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.

[132] Joseph Seering, Michal Luria, Connie Ye, Geoff Kaufman, and Jessica Hammer. 2020. It Takes a Village: Integrating an Adaptive Chatbot into an Online Gaming Community. In *Proceedings of the 2020 CHI*

Conference on Human Factors in Computing Systems (CHI '20). ACM, New York, NY, USA, Article 579, 13 pages. http://doi.acm.org/10.1145/3313831.3376708

[133] Joseph Seering, Tony Wang, Jina Yoon, and Geoff Kaufman. 2019. Moderator engagement and community development in the age of algorithms. *New Media & Society* 21, 7 (2019), 1417–1443.

[134] Randi Shedlosky-Shoemaker, Kristi A Costabile, and Robert M Arkin. 2014. Self-expansion through fictional characters. *Self and Identity* 13, 5 (2014), 556–578.

[135] Solace Shen, Petr Slovak, and Malte F Jung. 2018. Stop. I see a conflict happening.: A robot mediator for young children's interpersonal conflict resolution. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 69–77.

[136] Michael Warren Skirpan, Jacqueline Cameron, and Tom Yeh. 2018. More than a show: Using personalized immersive theater to educate and engage the public in technology ethics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 464.

[137] Carolyn Snyder. 2003. *Paper prototyping: The fast and easy way to design and refine user interfaces*. Morgan Kaufmann, San Francisco, CA, USA.

[138] Michal A Strahilevitz and George Loewenstein. 1998. The effect of ownership history on the valuation of objects. *Journal of Consumer Research* 25, 3 (1998), 276–289.

[139] Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scassellati. 2018. The ripple effects of vulnerability: The effects of a robot's vulnerable behavior on trust in human-robot teams. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 178–186.

[140] Michael Suguitan and Guy Hoffman. 2019. Blossom: A Handcrafted Open-Source Robot. *ACM Transactions on Human-Robot Interaction (THRI)* 8, 1 (2019), 1–27.

[141] Michael P Sullivan and Anre Venter. 2005. The hero within: Inclusion of heroes into the self. *Self and Identity* 4, 2 (2005), 101–111.

[142] Daniel Szafir, Bilge Mutlu, and Terrence Fong. 2015. Communicating directionality in flying robots. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 19–26.

[143] Xiang Zhi Tan, Samantha Reig, Elizabeth J Carter, and Aaron Steinfeld. 2019. From one to another: how robot-robot interaction affects users' perceptions following a transition between robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 114–122.

[144] Hamish Tennent, Solace Shen, and Malte Jung. 2019. Micbot: A peripheral robotic object to shape conversational dynamics and team performance. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 133–142.

[145] Frank Thomas, Ollie Johnston, and Frank Thomas. 1995. *The illusion of life: Disney animation*. Hyperion New York.

[146] Peter Tolmie, James Pycock, Tim Diggins, Allan MacLean, and Alain Karsenty. 2002. Unremarkable computing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 399–406.

[147] Sherry Ed Turkle. 2007. *Evocative objects: Things we think with.* MIT press.

[148] John Vines, Tess Denman-Cleaver, Paul Dunphy, Peter Wright, and Patrick Olivier. 2014. Experience design theatre: exploring the role of live theatre in scaffolding design dialogues. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 683–692.

[149] Mark Weiser. 1991. The Computer for the 21 st Century. *Scientific american* 265, 3 (1991), 94–105.

[150] Tom Williams, Priscilla Briggs, and Matthias Scheutz. 2015. Covert robot-robot communication: Human perceptions and implications for human-robot interaction. *Journal of Human-Robot Interaction* 4, 2 (2015), 24–49.

[151] Allison Woodruff, Sally Augustin, and Brooke Foucault. 2007. Sabbath day home automation: it's like mixing technology and religion. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 527–536.

[152] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable ai: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.

[153] Svetlana Yarosh, Stryker Thompson, Kathleen Watson, Alice Chase, Ashwin Senthilkumar, Ye Yuan, and AJ Brush. 2018. Children asking questions: speech interface reformulations and personification preferences. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*. ACM, 300–312.

[154] John Zimmerman. 2009. Designing for the self: making products that help people become the person they desire to be. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 395–404.

[155] John Zimmerman and Jodi Forlizzi. 2014. Research through design in HCI. In *Ways of Knowing in HCI*. Springer, 167–189.

[156] John Zimmerman and Jodi Forlizzi. 2017. Speed dating: providing a menu of possible futures. *The Journal of Design, Economics, and Innovation* 3, 1 (2017), 30–50.