Not Some Random Agent: Multi-person Interaction with a Personalizing Service Robot

Samantha Reig[1], Michal Luria[1], Janet Z Wang[2], Danielle Oltman[3], Elizabeth Jeanne Carter[1], Aaron Steinfeld[1], Jodi Forlizzi[1], John Zimmerman[1]
[1]Carnegie Mellon University, Pittsburgh, PA, USA [2]Princeton University, Princeton, NJ, USA
[3]McGill University, Montréal, Quebec, Canada
{sreig,mluria}@cs.cmu.edu,janetzw@princeton.edu,daniellejoltman@gmail.com,
{lizcarter,steinfeld}@cmu.edu,{forlizzi,johnz}@cs.cmu.edu

ABSTRACT

Service robots often perform their main functions in public settings, interacting with more than one person at a time. How these robots should handle the affairs of individual users while also behaving appropriately when others are present is an open question. One option is to design for flexible agent embodiment: letting agents take control of different robots as people move between contexts. Through structured User Enactments, we explored how agents embodied within a single robot might interact with multiple people. Participants interacted with a robot embodied by a singular service agent, agents that *re-embody* in different robots and devices, and agents that *co-embody* within the same robot. Findings reveal key insights about the promise of re-embodiment and co-embodiment as design paradigms as well as what people value during interactions with service robots that use personalization.

CCS CONCEPTS

• **Human-centered computing → Laboratory experiments**.

KEYWORDS

human-robot interaction; human-agent interaction; service; groups; re-embodiment; personalization; interaction design

ACM Reference Format:
Samantha Reig, Michal Luria, Janet Z Wang, Danielle Oltman, Elizabeth Jeanne Carter, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2020. 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 (HRI '20), March 23–26, 2020, Cambridge, United Kingdom.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3319502.3374795

1 INTRODUCTION

Due to the ubiquity of smartphones, a large proportion of the population has artificial intelligence (AI) on hand at all times, but regular use of voice agents is largely question-answer in nature. We are still learning how to take full advantage of the immense computational power and diverse human-agent interaction opportunities that live in our pockets. Interactions with agents are also extending beyond our personal devices. Robots greet patients and carry supplies between stations in hospitals [25], carry goods for last-mile delivery [12, 13], lead travelers to departure gates in airports [43], and do maintenance in grocery stores [14]. These kinds of services place robots in close proximity to human-human interaction, even if social interaction is not their main purpose.

The presence of robots in service environments permits a new touchpoint for personalized service. Research shows that people increasingly prefer a single point of contact: customers wish (and expect) to interface with one agent that is knowledgeable about all touchpoints and is situationally and temporally aware [33]. This is inherently difficult for human agents. However, AI can allow a service to craft personalized experiences that go beyond what people alone can achieve, fostering human-agent service relationships that do not necessarily mimic human relationships. Companies have begun to leverage this, addressing design for the use of multiple voice assistants on the same device [1, 6]. As there is virtually no cost to instantiating multiple social agent presences, a wide array of personalized, branded, and unique agents could be deployed in future service environments.

One design that has great potential for personalization and continuity over multiple interactions is *re-embodiment* [23]—when an agent moves its *social presence* from one device to another, taking on the physical capabilities of each physical device. This allows users to interact with one familiar social presence during multiple steps of an extended interaction. Previous work showed that people generally will accept agents that re-embody robots and IoT devices [23]. Interacting with one social presence across multiple robots contributes to a seamless experience. For social robots embedded in services, re-embodiment provides some benefits over the traditional human model because users do not have to repeat what they have already said, acquaint themselves with new faces of the same service, and re-ground the relationship at every step. This begets integration, which is a key variable in service design. It also has benefits over interactions with different "expert agents" because users can come to feel personally "known" over time.

This work is part of a larger effort to explore how re-embodiment, personalization, and other "super human" AI capabilities interplay with service robots. Here, we investigate interactions among one robot, multiple AIs, and multiple people. We explore *co-embodiment*, i.e., multiple AIs residing within the same robot, as a means of giving a single robot the ability to interact in a personal way with

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

ACM ISBN 978-1-4503-6746-2/20/03…$15.00
https://doi.org/10.1145/3319502.3374795

Figure 1: Our service robot prototype. The images displayed on the screen changed as different agents embodied the robot at different times.

multiple people who may or may not know each other. We designed our study based on four open-ended research questions intended to inspire and guide scenario design and analysis. Our first research question pertains to the social norms of human-robot group interactions in service contexts:

RQ1: How should a robot personalize its performance of service with multiple users? How does context influence this?

We also explore the novel question of how multiple social agents should interact through the same physical platform (co-embodiment):

RQ2: How does co-embodiment impact people's perception of the service robot experience?

Re-embodiment gives social robots the opportunity to make use of multiple individual, personalized agents that tailor their behavior to their primary users. This opens up questions about how people develop relationships with agents, the robots they embody, and the services with which they interface:

RQ3: How does a sense of personal connection to a robot's intelligence influence trust in that robot and feelings about the services it helps to provide? What is the social role of a universal personal agent?

Finally, re-embodying agents can interact with people through different robots, in different locations, and in both related and unrelated contexts. This is a useful feature overall, but it may be inappropriate at certain times. Additionally, it is likely that the timing of these transitions between contexts should follow certain rules and that there will be some degree of nuance in their design. When the same social presence can assist a person in multiple aspects of their life, it is important to understand where social and personal boundaries lie in terms of switching from one physical or topical domain to a completely different one:

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

Given the futuristic nature of these questions, we utilized structured User Enactments [10, 32, 44] to explore how service robots should handle personalization and to attempt to address our four research questions. This methodology has a proven track record for gathering important insights on novel technologies.

2 RELATED WORK

Our work lies at the intersection of service robot personalization, social behavior of robots in groups, and flexible agent embodiment.

2.1 Personalization in robotic services

Personalization is a key aspect of a user's relationship to a service and has been said to be the most important variable in determining perceived service quality and customer satisfaction [26]. HCI research into theories of user personalization of the appearance of computers and phones suggests that while users can apply personalization to their devices of their own accord, features to enable personalization can also be built into the design of the device [5]. Recently, HRI researchers have designed robots with the explicit purpose of personalization and customization of physical appearance and behavior such that the same base platform can be used for numerous projects [36]. Critically, personalized experiences can also increase loyalty by way of enhancing satisfaction and trust [2].

Trust and personalization are often intertwined for robots. Research has shown that a single error can impact humans' trust in the robot, especially in critical situations [34]. Similarly, a robot's mishandling of personalization may have irreversible effects on a human-robot relationship; for example, a hospital robot that does not provide a patient with their desired level of privacy may destroy trust in that robot, and perhaps in the hospital. Fortunately, personalized interactions with a robot can also be beneficial. In a field study of long-term interactions with a robot embedded in a workplace, incorporating discussion of personalized topics like food preferences, frequency of use, and prior service breakdowns increased rapport and cooperation with the robot as compared with discussing social, but not personalized, topics [21].

There is also a demand for personalization: owners changed Roomba appearances to express identity or to fit in in the home environment [37], and potential users of elder care robots placed a high value on the affordance of robot personalization to meet patients' particular emotional and physical needs [27]. There has been limited work on design guidelines for adaptive robotic services. Lee and Forlizzi [20] augmented the conventional service blueprint with a *line of adaptivity*, which describes both changes in the service and changes in the user through repeated interactions. These research efforts, and the majority of work in designing for personalization, have focused on personalization for a single user and had little regard for the surrounding social context.

2.2 Agents and robots in groups

A vast literature on how robots should behave in groups and teams focuses heavily on what factors influence human attitudes and behaviors. For example, prosocial robot behavior was perceived more positively even if the team did not perform well [9]. In another study, a robot's vulnerable commentary increased engagement and reduced tension among teammates in a game [35]. Recent work has considered the degree to which social cues, such as the order in which a robot greets members of a group and subtle indications of affinity, matter for shaping longstanding human-human relationships [16]. Mutlu et al. [28] found that robots must align with existing social norms to integrate fully into an organizational workflow. Additionally, the same robot in the same complex environment can be perceived differently depending on immediate physical surroundings. For example, people perceived the robot as "getting in the way of urgent work" in busy, cluttered hallways, but not in less-crowded locations [28]. Other work has explored how agents should be designed for tasks requiring expertise. In one study, participants were confused by interactions with multiple "expert" chatbots. The authors concluded that multiple chatbots provided little added value to this sort of interaction because the challenges in turn-taking and understanding that emerged were similar to those of single-chatbot conversations [8]. Whereas that work focused on multiplicity of agents for multiple subtopics of a broader overall topic (travel), our work focuses on multiplicity of agents for unrelated topics, and for multiple users.

2.3 Flexible agent embodiment

Limited research exists on migrating agents, or social presences that move across (i.e., re-embody) physical platforms according to the context of interaction. The work has focused primarily on desires and associations pertaining to singularly-embodied agents [17, 38], the salience of the agent migration phenomenon [17], and possible future technical implementations of it [11]. One study [18] sought to examine long-term interactions with re-embodying agents by prototyping migrating intelligence in a mock smart home. The study found that as people became more familiar with how an agent moves between devices over time, it became easier for them to recognize its identity. Aside from this, there is little to no work on the concept of individual agent "personalities" that re-embody multiple robots. We build upon our prior work [23], which found that people are generally accepting of, and comfortable with, robot re-embodiment, because it creates a seamless experience. It also exposed research questions pertaining to robot expertise, discomfort with co-embodiment, and contextual boundaries. In this study, we look more closely at these questions and explore new ones.

3 METHOD

To understand how service robots can employ co-embodiment and re-embodiment to personalize multi-party interactions, we designed a series of User Enactments (UEs) [10, 44]. UEs use low-fidelity prototypes and Wizard-of-Oz methodology to immerse participants in several "possible futures". By experiencing interactions with mockups of future technologies, participants can reflect critically on what they saw, did, and felt, and compare experiences to one another. UEs work especially well in exploratory research, where social mores have not yet emerged, and where there are no existing design patterns. We ran two participants at a time and interviewed them together, which enabled co-discovery and surfacing of knowledge and ideas that one person alone might not have recognized [22]. Participants signed up together and knew each other, which improved the authenticity of the group experience.

3.1 Study setup

The study took place in a lab that was divided into four separate "rooms" by rolling floor-to-ceiling walls. We used scripts that were the result of several weeks of brainstorming and acting out service interactions. The robot was a custom-built exemplar designed for service tasks (see Figure 1). The body was made of cardboard with an exterior paper layer. The head was a Kubi desktop telepresence robot with an iPad. We used an iRobot Create as the base. 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

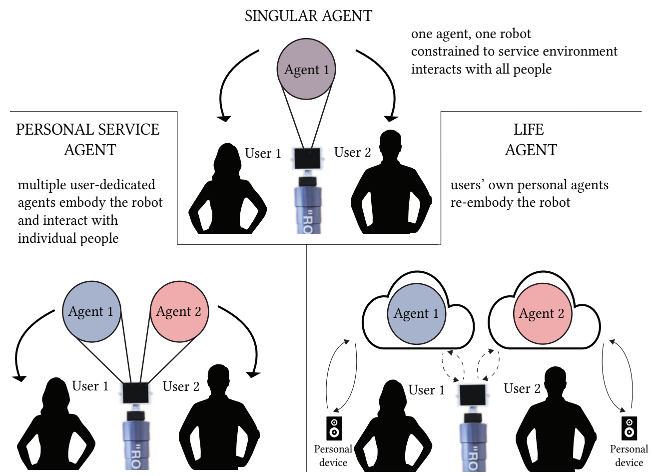

Figure 2: The three configurations.

with five different voices to generate the agents' scripted speech in advance, and we kept a repository of Google TTS-generated common phrases so that the agents could respond to unplanned deviations. We used three cues to communicate agent identity: each agent had a distinct name, a distinct voice, and a "profile picture" that would appear on the screen whenever that agent was meant to have control of the robot. The software that ran the wizard's end of the interaction can be found at https://github.com/AutonomyLab/create_autonomy and https://github.com/CMU-TBD/HRI20-Not-Some-Random-Agent-Personalizing-Service-Robot. A researcher controlled the robot and the agents' voices. The robot followed the same paths each time, so there was minimal variation in its movement. The wizard, who was the same researcher throughout the study, followed a defined script for movements and verbalizations and was instructed to deviate from the script only if the interaction with the participant required an alternative or unique response.

3.2 Agent configurations and environments

We designed three agent configurations to explore different interactions that might appear with future service robots (Figure 2). We chose these as an initial foray into the design space because they are (1) distinct enough from each other to facilitate critical reflection about ways in which public-facing robots can create a sense of personal connection, (2) conducive to social interaction with multiple people *and* multiple agents, and (3) testable with human dyads (a "single-agent, many-people" configuration limits exploration of certain questions). We utilized a structure that appears similar to a 3x3 study design to ensure good coverage of various permutations of contexts and configurations. The added structure helped us cover a vast design space relative to re-embodiment and co-embodiment and avoid overly redundant scenario combinations.

3.2.1 Agents. We iterated on concepts to reach three designs for service robots that personalize interactions.

Singular Agent. This configuration consists of one robot embodied by one agent and is essentially a baseline, i.e., the common paradigm in present-day robots. A Singular Agent (SA) is affiliated with the space(s) it is in and owned and maintained by the service. The agent has information about and "knows" its regular customers. Here, we explored perceptions and impressions of one agent that stores and uses information from multiple repeat users.

Personal Service Agent. A logical step up in service delivery is when a service agent is owned and maintained by the service provider but personalized to each customer. We call this concept a Personal Service Agent (PSA). PSAs are personalized agents assigned and curated by a company or institution. Multiple PSAs can exist in a single physical embodiment. Individual interactions with PSAs are *one agent per user* within a single environment. Because these agents are permanently affiliated with the same service and may need to say the same thing to two concurrent customers, we posited that the PSAs could speak in unison (in a "chorus") to communicate the same message to different people at the same time. With this configuration, we were interested to learn: Should co-embodying PSAs be aware of each other's conversations? How should they talk to each other? We also wanted to explore privacy concerns about agents sharing a data source.

Life Agent. A third option is for each service robot to host multiple individual, personalized AI assistants that are accessed by their users in all aspects of their lives. In the Life Agent (LA) configuration, agents are able to re-embody robots and other devices as needed. Each time the LA re-embodies, it can access the physical capabilities of its current housing and the data specific to the current environment. Thus, it can do tasks with different physical and information demands while allowing the user to interact with any number of unfamiliar devices through the same familiar social intelligence. Pertinent questions are the perceived relationship between LA "software" and robot "hardware" and the evolving social role of this type of integrated AI personal assistant.

3.2.2 Service environments. We designed three environments to examine the influence of service context. These were deliberately chosen to probe issues related to privacy and security, comfort, conversational design, long-term interactions, and social roles. We implemented personalization differently in each environment: in the hotel, it was addressed in terms of food preferences; in the department store, transaction records; in the clinic, medical history.

Quick Care Clinic. Participants entered the clinic together and the robot welcomed them each by name. Then, it guided each participant through the processes of checking in, waiting in a waiting room, and receiving a flu shot. In the LA configuration, P1's agent alerted them that a package had arrived at their home, and P2's agent notified them that an upcoming flight was delayed. The LAs used language that was more colloquial (e.g., "Have a seat" vs. "Please sit down") to connote a long-term personal relationship. In the clinic, we sought participants' impressions of agents' ability to use domain expertise and reveal potentially sensitive information.

Canton Department Store. The store environment mimicked two sections of a larger department store. The robot greeted both participants by name, asked (or, for LA, verified) what they were looking for, helped them find the items, and processed payment using a credit card on file. This allowed us to explore how robots should use and talk about personal data in a public space as well as how a robot might handle personalization in a non-personalized environment like a store.

Homestead Inn. In this scenario, we had participants ask a hotel concierge robot for nearby dinner recommendations in an unfamiliar area. Before the interaction, each participant was given a list of dietary, location, and budget requirements, with the goal of finding a restaurant that met both sets of criteria. The agent greeted participants by name and recommended restaurants based on known information about the users and general customer ratings. In the PSA and LA designs, each agent searched for a restaurant on behalf of its own user. Here, we explored how a robot utilizing co-embodiment might engage in a negotiation-like exchange to help users come to a joint decision.

3.3 Participants

We recruited 48 participants (24 pairs) via fliers, word of mouth, internet posts, and a local online recruitment tool. Participants were between 20 and 76 years old ($M(SD) = 39.3(17.6)$) and had a variety of personal and professional backgrounds. 25 participants self-identified as female, 21 as male, and 2 as other. They interacted with computers regularly, $M(SD) = 6.48(1.25)$ on a seven-point Likert scale that ranged from never (1) to multiple times per day (7). They interacted with AI assistants less frequently ($M(SD) = 3.31(1.91)$), had some familiarity with robots ($M(SD) = 3.19(1.60)$), and had relatively favorable impressions of robots before the study ($M(SD) = 5.46(1.34)$ for an average of five correlated (= .73) questions about trust and goodwill toward robots). No participants were technical students at our institution.

3.4 Study procedure

After consenting to the study, participants filled out a pre-study questionnaire to collect demographics, experience with smartphones and computers, and preexisting associations with robots. A researcher then introduced the study, asking the participants to take on gender-neutral, study-assigned first-names and imagine that they were friends from work. During the introduction, the researcher stated that the goal of the study was to have participants help the team experience and critique potential future interactions with smart technologies. Participants then experienced each service environment with a different agent configuration (three scenarios). We counterbalanced the order of both environment and agent configuration to mitigate the interference of novelty effects in participants' experiences of each of the nine environment-configuration pairings. This meant that 16 participants (8 pairs) experienced each pairing (see Figure 3 for an example). We conducted semi-structured interviews with both participants together after each scenario and a final interview at the end of the study. The study took about 90 minutes, and participants were compensated $35 USD each. The protocol was approved by our Institutional Review Board.

3.5 Analysis

We identified several hundred meaningful quotes from the interviews, during which participants had an opportunity to respond to questions, react to probes, and reflect freely on their experiences. Our qualitative approach to our data was a thematic analysis in the form of (1) iterative affinity diagramming [4] and (2) application of categorical and sub-categorical labels to quotes based on the clusters that emerged during the affinity diagramming. This

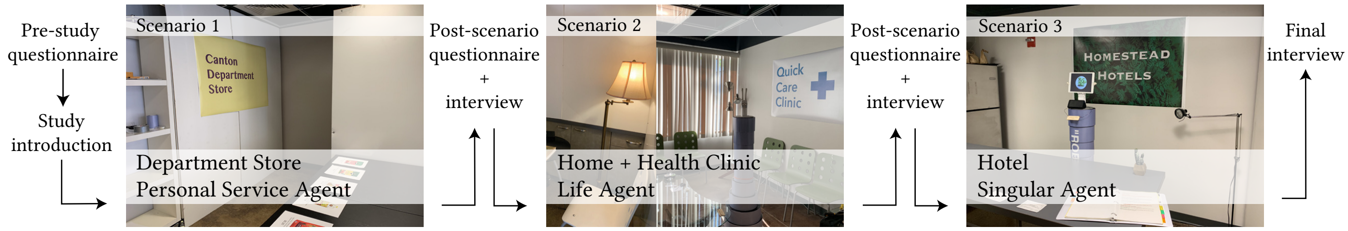

Figure 3: An example trial from the study. Participants experience the department store with Personal Service Agents first, then the clinic with re-embodying Life Agents (which follow them from home), then the hotel with a Singular Agent.

approach is used to draw out patterns and themes to explore non-existing, future interactions through UEs [10, 32]. The analysis was conducted primarily by two authors (one was personally involved in data collection, one was not) who met for multiple hours on several occasions to extract, interpret, and group the data together. They discussed with two other authors after each round of analysis and periodically consulted the remaining authors and a non-author researcher who was less familiar with the details of the scenarios.

We also took special note of responses to three specific questions about (1) acceptance of facial recognition, (2) the *chorus of agents* interaction, and (3) which configurations were most comfortable. We utilized post-scenario questionnaires to assess trust, social attributes (modified from [7] and [3]), and groupness, but results were fairly uniform across agent configurations and service settings. While our approach was primarily bottom-up, we referred back to our guiding research questions to inform the interpretation of the quotes with respect to our research focus.

4 FINDINGS

Through iterative analysis of our interview data, we uncovered insights pertaining to our research questions and discovered new themes. We compared a robot embodied by a Singular Agent (baseline configuration) with two variations of co-embodiment: agents owned and managed by the service and agents maintained by the user. Participants generally accepted re-embodiment and co-embodiment, but had some concerns about how re-embodiment might be controlled and how co-embodying agents might exchange data. They did not particularly like PSA, finding the two unique agents to be "redundant" (122B) without adding value. When participants had strong feelings about re-embodiment, these were about the personal nature of LAs. When they had strong feelings about co-embodiment, they were about the concept of multiple software intelligences within one robot. Thus, we report mostly on differences between the LA and SA designs. In the quotes we cite, *Alpha, Moon, Saturn, Basil,* and *Sunflower* refer to the five agents: Alpha is SA; Moon and Saturn are PSA; Basil and Sunflower are LA.

4.1 Preference for a Life Agent

Most people (22 participants) thought **a universal Life Agent was the most comfortable design**, followed by a Singular Agent (13 participants), and, finally, a Personal Service Agent (5 participants). Three participants found SA and LA equally comfortable, and 5 had no preference or did not answer the question. In general, participants thought that interacting with a familiar, private agent embodied in public robots would provide a smoother and richer experience. A singular agent was comparable to "just some random person" (119A) that would have neither out-of-context data nor a personal history with the user.

4.1.1 Personality. **Participants placed high value on the capability of customization of robot personality and identity attributes.** Many wanted robots to exhibit certain character traits when embodied by their own agents, sometimes focusing on traits that would align with or affirm personal values. For example, participant 110A wanted their agent to be hard on them. Participant 101B said, "I want it to be sarcastic because that's how I am. I want it to compliment me. It's like another friend." Some had specific voice characteristics in mind pertaining to gender or dialect: 102B suggested that an agent on the East Coast use East Coast slang, and 101A wanted an agent with a Nigerian or British accent.

Some participants went so far as to say that agents should remind them of their friends or themselves—even to the extent of taking on corresponding voice and speech characteristics. Participant 110B elaborated that a "cool, calm, and collected" person should have a matching robot. This idea is evocative of the well-known finding from sociology that people feel most comfortable socially interacting with people similar to themselves [19, 24].

- I'd want it to embody like a personality of my friends, just because you enjoy hanging out with your friends. (107B)
- Though I think it would be creepy, and I probably wouldn't do it, you should [...] have the choice to use your own voice. (103B)

4.1.2 Emotional support. An important function of the LA design is its ability to provide comfort and support. When reflecting on the clinic, several participants mentioned that **in situations that might be stressful or emotional, having a familiar agent would be "comforting"** (125A). Participant 123A mentioned that for someone afraid of shots, their LA should be able to "read that about [them]", and 113A said, "If you're feeling anxious [...], it's nice to have old friendly Basil along who knows everything about you." A few participants thought that robots were more flexible, less distractible, and less likely to get flustered or frustrated than humans; therefore, they were well-suited to jobs requiring patience and calmness. However, most people who alluded to empathy were more of the belief that it is a distinctly human quality that will be difficult or impossible to embed into robots' behavior (e.g., [40]). Re-embodiment has potential to augment robots that would otherwise seem impersonal or unsocial with empathetic characteristics just by virtue of feeling familiar and "known" to their users.

Robots can use co-embodiment and re-embodiment to help people feel more comfortable and at ease in unfamiliar spaces, but this raises a set of special design challenges. We discuss the two most significant of these: (1) giving users a sense of control over the interaction and (2) adapting the non-human behaviors of re-embodiment and co-embodiment to human social norms.

4.2 Context-crossing and uncertainty concerns

Because co-embodiment was novel, participants were not able to easily anticipate what an agent was going to do next. This became a problem predominantly when LAs had knowledge of participants' personal information, since it was not clear in what (potentially inappropriate) context the agent was going to make use of it in public. Some participants suggested ways to be more in control over interactions with LAs: customizing personality through a questionnaire (116A), using a settings menu to define the nature of the human-agent relationship, or adjusting the LA's conversational style on-the-fly (123B). Many people also felt that **automatic context crossing through re- and co-embodiment should be a toggle setting** such that users could decide, either permanently or for a period of time, to "turn that feature off" (105A).

4.2.1 Control over context crossing. Reactions to the context crossing behavior (i.e., getting non-health-related, robot-initiated personal notifications while at the medical clinic) were mixed. Some participants found this useful, while others thought it strange, awkward, or otherwise an unwelcome social violation. Some expressed surprise when the notification first came in but imagined adapting to such interruptions over time. Some participants noted that an agent that crosses context provides utility by leveraging instantaneous knowledge of remote situations to alert users to information that affects their schedule, safety, or health. For example, it may be appropriate for a user to receive a flight update while at a medical clinic because that can affect their plans for the day. However, inability to anticipate a Life Agent's behavior also led to concern that it might inappropriately surface "out-of-context" information in front of others, oblivious to the incongruous social setting.

Additionally, a universal LA blurs the boundaries between aspects of life that are otherwise separate, and the resulting bleed-through may not always be desirable. For example, 118A said, "There's some universal information like contact lists and stuff like that. But for the most part work should be work and home should be home, should be separate, limited data passing."

4.3 Agents are social actors in groups

"Appropriate social behavior" for social agents and robots is not a universal constant: both social context (i.e., the size and composition of the interacting party) and situational context (i.e., the space, place, and task at hand) can change how it should be defined. We found evidence that what is perceived as appropriate social behavior (or lack thereof) of a re-embodied robot may be dramatically impacted by the presence of others. We also found different impressions of social behavior in our three different environments.

4.3.1 Conversational intelligence and social norms. Following conversational norms refers to appropriate physical distance, politeness, common ground, and listening behavior. During interactions among multiple humans *and* a robot, these norms are already at play. Participants felt strongly that a robot should follow norms: 121A said, "Saturn cut me off! [...] If I don't finish, please don't speak!" The field of HCI has long known that people treat technology socially [29, 31] and expect agents to have some social intelligence [30]. However, co-embodied robots encounter special challenges in the way of appropriate conversational behavior. In our study, the coordination of multiple agents sometimes complicated conversational turn-taking, producing "unnatural" (116B) and awkward experiences. Matching or mismatching social norms can also manifest in physical behavior:

- The robot had rolled over to help Alex and then I was still over there and it just turned in my direction and sort of shouted at me instead of coming over to me to talk to me. (122B)

How co-embodied robots handle these norms can also influence or be influenced by morphology. For anthropomorphic robots, in which lifelike physical features reinforce identity, it may be more difficult to communicate the presence of multiple agents.

4.3.2 Understanding existing relationships. Participants believed that it is important for robots to acknowledge an awareness of relationships and history among human members of a group and treat them accordingly. If the humans are strangers, for instance, the robot should "give them their space" (107A). For some, **a robot's ability to exhibit an understanding of human relationships may be a determiner of acceptance of co-embodiment**, especially when interpersonal trust is critical, as in a medical setting.

- How did that agent know that we were even okay getting recognized in each other's presence? (123A)
- If we feel comfortable enough as coworkers to go to the clinic together, I think we can share the same robot body. (119B)

The behavior of agents in a group setting can also influence the way humans perceive and interact with each other—both in the short term as they navigate a conversation, and in the long term as they form lasting impressions of each other. Our interviews suggested that this mediation-like outcome is desirable in low-risk situations that already lend themselves to some degree of casual human social interaction (e.g., the hotel scenario). Through LAs, one person can "learn about the other person you're with very quickly [...] I didn't know he was vegetarian"(107A).

4.3.3 Did it work? We found that people wanted robots to use human-readable signals to continually communicate information about their status, including multitasking ability, current load, and general capacity. This was prioritized over both efficiency and humanlike social behavior. Even when information did not need to be repeated out loud for an interaction to continue, several participants wished that they had gotten some sort of confirmation that the robot had in fact heard them correctly and performed the task as it claimed it would. This was especially true when accuracy was important and perceived risk was high—e.g., when confirming that it was safe to get a flu shot or that the correct credit card had been used. This is somewhat consistent with prior work, in which people wanted robots to verbally acknowledge the receipt of personal information, even without repeating all of it aloud [41].

4.3.4 Who has the floor? There was a great deal of concern about how co-embodying agents would negotiate multiple users with independent needs and interests. Many people requested that a co-embodied robot provide a "clear indication" (121A) when one agent's interaction ends and another's begins, or when one agent has "handed off" control of the robot to another agent:

- It didn't say, like, Sunflower logging off, Basil logging on, or they didn't switch their icons or it didn't say, like, bye Sunflower, it's Basil's turn now. (109B)

Prior work established that simple movements can go a long way in communicating to users what a virtual agent [42] or robot [39] is about to do. More work is needed to understand how a robot designed to convey multiple "characters" or "personalities" at once could express intent and how the agents embodied in such a robot should negotiate control over that expression.

4.3.5 Inter-agent relationship. There were strong, polarized reactions to the PSAs speaking at the same time. Of our 48 participants, 22 were receptive to the "chorus", 20 were uncomfortable with it, and 6 did not perceive it. Negative responses were rather extreme: participants described the chorus as "an ominous flavor" (109B), "weird" (115B), "creepy and horrible" (122B), and "completely unnerv[ing]" (109A). To better understand these reactions, we affinity diagrammed 31 related quotations. Comments fell into five categories: negative feelings, positive feelings, appreciation of utility, functional complaints, and indifference. Though a few participants were excited about the agents' simultaneous speech, positive feelings mostly took the form of passive acceptance rather than enthusiasm. Many negative feelings stemmed from the fact that it is an extremely non-human behavior.

For both PSA and LA, people did not think the co-embodying agents had a social relationship to each other. They did not believe that the agents would intentionally exchange private information, but they worried that a single robot being embodied by multiple agents could lead to their personal information being "mixed up" (112B) with someone else's due to a mistake or malfunction. People found the idea of inter-agent social conversation creepy and, consistent with prior findings [23], feared the prospect of agents "talking behind their back". The exception was negotiation: if agents could coordinate to balance users' preferences or needs, they should. We observed this in the form of overall positive responses to the PSA interaction in the hotel. In other words, **if agents verbally communicate with each other the way humans do, it should only be in immediate service to the user**.

4.4 Flexible role conflicts with expertise

We observed **a belief that the more expertise a skill required, the less likely a Life Agent would be to have proficiency in that skill**. As in prior work [23], participants had doubts about a "jack of all trades" agent, fearing that it would in fact be a "master of none". In the questionnaire, ratings of trust were lower for the LA, which is intended to serve in multiple domains and embodiments, than for the SA, which is tied to one domain and embodiment and therefore may be more readily considered (and trusted as) an "expert". Beyond this, some participants generally doubted the ability of robots to have real expertise in a non-technological or human-centric domain, or one in which judgment and accuracy in the face of ambiguity are critical (this is similar to [40]). Concerns about expertise were most prominent in the clinic scenario: 9 participants commented on it in the clinic vs. 6 in the store and 2 in the hotel. Participant 119B said that an LA would be trustworthy "if it was a fairly routine problem", but with "a bunch of mystery ailments, I would definitely want a second opinion". Some people commented that upon getting wrong information in a store, "you can find it yourself" (112B), but when it comes to health, e.g., "wrong medicine" (112B), non-experts cannot correct mistakes.

4.5 Personal data and privacy

No participants reacted negatively to being recognized upon walking into the clinic. We asked about facial recognition in the clinic setting to explore recognition in the context of private and potentially sensitive information. Even though we did not ask explicitly about it in the post-scenario interviews for the store and hotel, participants took note of it in all three environments. An important characteristic of re-embodiment is that a user's data can move with an agent between robot bodies. This sparked some concern about data leaking from a trusted source to an unknown entity. On the other side of the coin, when an agent was their own, some participants had an increased sense of security—all of their information was concentrated in one place and they did not have to share it in every new context. Instead, a Life Agent could appear and make use of the relevant data. This raises an interesting design challenge: can a robot's behavior indicate that a user's data has left its hardware?

4.6 Other findings

A few participants mentioned wanting the robot to have eyes or a face, and some (109A, 111A, 112B) suggested using different modalities (e.g., voice and text input) to ensure that it can be used by older adults and people with disabilities. Many participants did not notice the agents' different voices but noted the changing "profile pictures" and distinct names. As such, voice alone is probably not a strong enough cue to signal agent identity early in a human-agent relationship. Interestingly, this contradicts the original finding from Nass' Computers Are Social Actors experiments [31], which found that different voices elicited different social attributions, even in interactions with a novel system. Another theme was societal implications of the futuristic technology we presented. Several participants noted skin tone and accent biases that exist in current face and voice recognition technology. Some expressed concerns about the roles robots will play in the future, including worry that they will not be equipped to carry out the emotional responsibilities humans do and fear that they will take away human jobs.

5 DISCUSSION

Our findings address service robot personalization and broader questions about human-robot relationships. Interpreting them requires consideration of the study's limitations: it took place in a lab, agents could not stray too far off-script, the robot was a low fidelity prototype, and only a few people were in the room. Together, these may have contributed to a lack of realism that interfered with participants' ability to fully immerse themselves in the scenarios.

We derive preliminary guidelines for designing the behavior of re-embodying agents, which are of interest to creators of robots and conversational AIs. We also contribute a new way to use UEs to acquire knowledge during an intermediate step of the design process. When a space is largely unexplored, but enough has been learned to spark specific research questions, researchers can add structure (probes, scripts, variations, etc.) to traditional enactments. Thus, they can draw comparisons but leave the experience unconstrained enough to facilitate revelation of "unknown unknowns".

We inquired as to how re-embodying agents should perform their service with multiple users (**RQ1**). We found that participants prioritized social competence and personalization during group interactions. We noted a distinction between personalization of social features and personalization of personal information. Participants in our study envisioned a Life Agent to be able to prioritize information that was specific and pertinent to them (perhaps in contrast to other users) and to build on and draw from that knowledge over the long term, regardless of whether or not its personality and social behaviors were customized. This increased their feelings of comfort interacting with the agent (**RQ3**) and made it generally desirable.

Our second research question (**RQ2**) concerned the overarching impact of co-embodiment on perceptions of social robots. Co-embodiment was received as (1) *necessarily* concerned with social signaling, and (2) appropriate for friends, but not for strangers. We draw from this two concrete design guidelines for co-embodying and co-embodyable systems. The first is opt-in co-embodiment: robots in public settings can enable co-embodiment, but should not be embodied by two agents at the same time by default; and they should be explicit on whether a third party can gain access to the data from an interaction. The second is clear indications from robots about what ("who") is in charge. Repetitive signaling regarding which agents are being accessed and which users are being helped is critically important for users to understand how to interact with a co-embodyable robot, at least in early interactions.

In response to our research question on contextual boundaries (**RQ4**), we find tension between comfort and expertise: people have difficulty with the idea of one social agent that claims to be equally adroit in all possible domains and embodiments. At the same time, they want to interact with novel robots through a Life Agent that aligns with personal identity and values. This presents a challenge of balancing quality and quantity. We conjecture that embodiment of personal agents in non-personal robots is best used for tasks that are perceived to be relatively low-risk—for example, helping people navigate a building using familiar language or making recommendations in a grocery store based on knowledge about cooking habits. In contrast, when perceived risk is high, as in a medical setting, robots need to prioritize the communication of their expertise over personal connection and emotional support. One approach to mitigating this tension might be to design an agent that communicates that it is acquiring expertise. For example, a Life Agent, upon entering a healthcare facility, might communicate that it is acquiring new expertise in support of the user's interactions with the service. But in some cases, re-embodiment of a Life Agent into a domain-specific robot may be best foregone entirely in favor of clear assurance that a robot is well-versed in the task and solely dedicated to it.

When agents *do* transition across contexts, our data suggests they should clearly express the features that constitute their identity. Defining the minimum cues necessary for users to recognize an agent is a critical part of designing re-embodyable systems. Confusion about how and when re-embodiment has occurred may be tied to discomfort with the concept and, in turn, result in lower acceptance. We used three attributes to communicate an agent's identity: image, voice, and name. In our study, image was a much stronger cue than voice or name. Of course, it is not feasible to take this as an absolute because many robots do not afford projecting an image onto a screen and because visual-only channels make robots less accessible. What we can conclude is that whenever possible, designers of re-embodyable robots should provide a means of visually indicating the presence of different agents.

Finally, our study provokes examination of and reflection on the role of robots in society. The lack of concern with facial recognition by robots in both private and public spaces likely requires a more nuanced inquiry than our study provided. The broader privacy issue of the conflicting interests of multiple stakeholders has a close and complicated relationship with feelings about facial recognition: for example, a few participants felt that facial recognition in a clinic setting would be useful or even necessary for trust, but that in a commercial setting, it would be in the interests of the company rather than their own and an inappropriate violation of their privacy. Future research into where and when it is acceptable for robots to use facial recognition, and how storage and usage of that data should be communicated, will benefit the design of service robots from a user experience perspective as well as an ethical one.

The preference for a customizable Life Agent similar to oneself raises questions about defaulting to designs that reinforce people's tendency to gravitate towards similar others. The non-human characteristics and customizable capabilities of social AIs and robots may make them conducive to designs that challenge social biases rather than conform to them. Some participants asked to have a personal agent with qualities that complemented, rather than matched, their own. This gives credence to the idea that while people value the familiarity and support of agents that are like themselves, they may also accept, and even desire, dissimilar agents [15].

6 CONCLUSION

This study investigated how future service robots can use personalization to interact with multiple users. Through structured user enactments and interviews, we found that people are receptive to the idea of robots that leverage personal information if the user has control over the information. We also discovered that service robots embodied by multiple agents can make people more comfortable with group interactions by demonstrating an understanding of pre-existing human relationships within the group. Our work sheds light on the role of flexible agent embodiment during interactions with service robots, and suggests design guidelines and directions for future research on the topics of re-embodiment, co-embodiment, and personal human-robot interactions that occur in public.

7 ACKNOWLEDGMENTS

We thank Lynn Kirabo, Stephanie Valencia Valencia, and Xiang Zhi Tan for assisting us with documenting our study and consulting on analysis, and the undergraduates in our lab who helped with piloting. This work was supported by awards NSF SES-1734456 and NASA 80NSSC19K1133.

REFERENCES

[1] Roberto Baldwin. 2019 (accessed October 1, 2019). *Alexa can mimic celebrities, starting with Samuel L. Jackson.*

[2] Dwayne Ball, Pedro S Coelho, and Manuel J Vilares. 2006. Service personalization and loyalty. *Journal of Services Marketing* 20, 6 (2006), 391–403.

[3] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics* 1, 1 (2009), 71–81.

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

[5] Jan O Blom and Andrew F Monk. 2003. Theory of personalization of appearance: why users personalize their pcs and mobile phones. *Human-Computer Interaction* 18, 3 (2003), 193–228.

[6] Dieter Bohn. 2019 (accessed October 1, 2019). *Amazon creates a huge alliance to demand voice assistant compatibility.*

[7] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. 2017. The robotic social attributes scale (RoSAS): Development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction.* ACM, 254–262.

[8] 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.* ACM, 191.

[9] Filipa Correia, Samuel F Mascarenhas, Samuel Gomes, Patrícia Arriaga, Iolanda Leite, Rui Prada, Francisco S Melo, and Ana Paiva. 2019. Exploring prosociality in human-robot teams. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* IEEE, 143–151.

[10] 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.

[11] 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.

[12] Darrell Etherington. 2019 (accessed September 5, 2019). *Postmates lands permit to test its Serve autonomous delivery robots in SF.*

[13] Chaim Gartenberg. 2019 (accessed September 5, 2019). *Amazon is now delivering packages in Southern California with its Scout robots.*

[14] David Grossman. 2019 (accessed September 5, 2019). *Googly-Eyed Robots Are Coming to Hundreds of Grocery Stores.*

[15] Katherine Isbister and Clifford Nass. 2000. Consistency of personality in interactive characters: verbal cues, non-verbal cues, and user characteristics. *International Journal of Human-Computer Studies* 53, 2 (2000), 251–267.

[16] Malte F Jung, Dominic DiFranzo, Brett Stoll, Solace Shen, Austin Lawrence, and Houston Claure. 2018. Robot Assisted Tower Construction-A Resource Distribution Task to Study Human-Robot Collaboration and Interaction with Groups of People. *arXiv preprint arXiv:1812.09548* (2018).

[17] Chyon Hae Kim, Yumiko Yamazaki, Shunsuke Nagahama, and Shigeki Sugano. 2013. Recognition for psychological boundary of robot. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction.* IEEE Press, 161–162.

[18] Kheng Lee Koay, Dag Sverre Syrdal, Wan Ching Ho, and Kerstin Dautenhahn. 2016. Prototyping realistic long-term human-robot interaction for the study of agent migration. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN).* IEEE, 809–816.

[19] Paul F Lazarsfeld, Robert K Merton, et al. 1954. Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society* 18, 1 (1954), 18–66.

[20] Min Kyung Lee and Jodi Forlizzi. 2009. Designing adaptive robotic services. *Proc. of IASDR '09* (2009).

[21] Min Kyung Lee, Jodi Forlizzi, Sara Kiesler, Paul Rybski, John Antanitis, and Sarun Savetsila. 2012. Personalization in HRI: A longitudinal field experiment. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* IEEE, 319–326.

[22] 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.

[23] 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 conversational agents. In *Proceedings of the 2019 on Designing Interactive Systems Conference.* ACM, 633–644.

[24] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27, 1 (2001), 415–444.

[25] Mattie Milner and Stephen Rice. 2019 (accessed September 5, 2019). *Robots are coming to a hospital near you.*

[26] Banwari Mittal and Walfried M Lassar. 1996. The role of personalization in service encounters. *Journal of Retailing* 72, 1 (1996), 95–109.

[27] Sanika Moharana, Alejandro E Panduro, Hee Rin Lee, and Laurel D Riek. 2019. Robots for Joy, Robots for Sorrow: Community Based Robot Design for Dementia Caregivers. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* IEEE, 458–467.

[28] Bilge Mutlu and Jodi Forlizzi. 2008. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In *Proceedings of the 3rd ACM/IEEE International Conference on Human-Robot Interaction.* ACM, 287–294.

[29] Clifford Nass, BJ Fogg, and Youngme Moon. 1996. Can computers be teammates? *International Journal of Human-Computer Studies* 45, 6 (1996), 669–678.

[30] Clifford Nass, Youngme Moon, Brian J Fogg, Byron Reeves, and Chris Dryer. 1995. Can computer personalities be human personalities?. In *Conference companion on Human factors in computing systems.* ACM, 228–229.

[31] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM, 72–78.

[32] 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.

[33] Adam A Rapp, Daniel G Bachrach, Karen E Flaherty, Douglas E Hughes, Arun Sharma, and Clay M Voorhees. 2017. The role of the sales-service interface and ambidexterity in the evolving organization: A multilevel research agenda. *Journal of Service Research* 20, 1 (2017), 59–75.

[34] Paul Robinette, Ayanna M Howard, and Alan R Wagner. 2017. Effect of robot performance on human–robot trust in time-critical situations. *IEEE Transactions on Human-Machine Systems* 47, 4 (2017), 425–436.

[35] 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.* ACM, 178–186.

[36] Michael Suguitan and Guy Hoffman. 2019. Blossom: A Handcrafted Open-Source Robot. *ACM Transactions on Human-Robot Interaction (THRI)* 8, 1 (2019), 2.

[37] JaYoung Sung, Rebecca E Grinter, and Henrik I Christensen. 2009. Pimp my roomba: Designing for personalization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM, 193–196.

[38] Dag Sverre Syrdal, Kheng Lee Koay, Michael L Walters, and Kerstin Dautenhahn. 2009. The boy-robot should bark!-children's impressions of agent migration into diverse embodiments. In *Proceedings: New Frontiers of Human-Robot Interaction, a symposium at AISB.*

[39] 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.

[40] L. Takayama, W. Ju, and C. Nass. 2008. Beyond dirty, dangerous and dull: What everyday people think robots should do. In *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI).* 25–32. https://doi.org/10.1145/1349822.1349827

[41] 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.

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

[43] James Vincent. 2017 (accessed September 5, 2019). *LG's new airport robots will guide you to your gate and clean up your trash.*

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