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 mock-ups 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

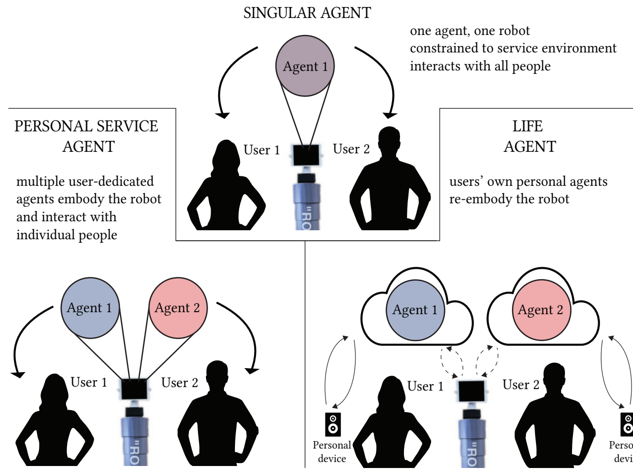

Figure 2: The three configurations.

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

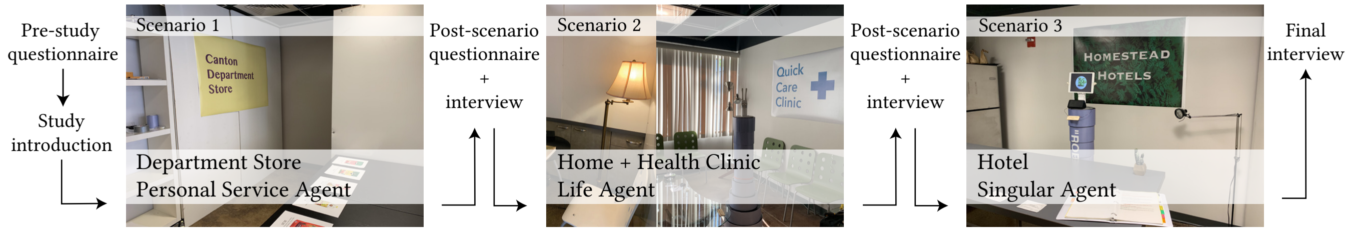

Figure 3: An example trial from the study. Participants experience the department store with Personal Service Agents first, then the clinic with re-embodying Life Agents (which follow them from home), then the hotel with a Singular Agent.

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.

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.