Iterative Design of a System for Programming Socially Interactive Service Robots

Michael Jae-Yoon Chung[1], Justin Huang[1], Leila Takayama[2], Tessa Lau[2], and Maya Cakmak[1]

[1] Computer Science & Engineering
University of Washington, Seattle, WA 98195, USA
mjyc,jstn,mcakmak@cs.washington.edu,
[2] Savioke Inc., San Jose, CA 95113, USA
leila,tlau@savioke.com

Abstract. Service robots, such as the Savioke Relay, are becoming available in human environments such as hotels. It is important for these robots to not only be functional, but also to have appropriate socially interactive behaviors. In this paper, we first present results from a formative study with service industry customers. A key demand we discover is that the robot should be aware of people present around the robot. We incorporate these lessons into the design of *iCustomPrograms*, a system for programming socially interactive behaviors for service robots. Next, we perform two field studies with *iCustomPrograms* and iterate its design. In the first field study, which took place at an airport, we witness people initiating interaction with the robot in unanticipated ways. The second field study, which took place over 2 weeks at 5 service industry properties, evaluates the socially interactive applications created with *iCustomPrograms*. Our experiences and findings from each study not only show the usefulness of our system in the field, but also provide insights for the design of future interactive applications for service robots.

1 Introduction

Today, commercial service robots such as the Savioke Relay, Vecna QC Bot, and Aethon TUG are deployed in human-populated environments such as hospitals and hotels[3]. Although these robots are designed for performing deliveries, it is important for them to be socially interactive, as they are tightly integrated into the human workplace [11]. For example, the Relay robot, built by Savioke Inc., primarily does room service deliveries to guest rooms in hotels. However, operating in the hospitality industry, it is important for the robot to have a suite of engaging, guest-facing interactions as well.

In fact, socially interactive service robots have long been the subject of interest to robotics researchers [2, 3]. Recent studies have explored socially interactive services such as guiding [1, 2, 8] and advertising [8] in uncontrolled

[3] www.aethon.com, www.vecna.com, www.savioke.com

human-populated environments such as office buildings [1], shopping malls [8], and airports [10]. As in those studies, we believe in the importance of designing for and evaluating in real-world environments.

In this paper, we present *iCustomPrograms*, a programming system, for developing socially interactive applications for mobile service robots. This paper contributes the design of *iCustomPrograms* as well as empirical findings from the deployment of interactive applications developed with *iCustomPrograms* in real-world service industry properties. We first interviewed employees in the service industry to understand what kinds of robot applications would be useful (Section 3). This information helped us to design the initial version of *iCustomPrograms*. We performed an initial field study to evaluate applications built with this system (Section 4). Based on findings from the initial study, we made enhancements to *iCustomPrograms* and conducted a larger field study (Section 5).

Our system was developed for the Savioke Relay, a 3-foot tall mobile robot with an interior bin and a touchscreen display (Fig. 2). However, *iCustomPrograms* can be used in principle with any robot that has similar capabilities [4].

2 Related Work

Researchers have long been interested in studying socially interactive service robots [3]. RHINO was one of the earliest examples, acting as a tour guide in a museum [2]. More recently, the SPENCER project investigated having a robot escort passengers in a busy airport [10], and the FROG project studied having a tour guide robot in an outdoor zoo [12]. The major contributions of this body of work addressed core technical challenges arising in the field, such as tracking human individuals or groups in crowded environments [10], or long-term outdoor mapping and localization problems [12].

Other researchers have studied human-robot interaction issues for service robots. Some investigated long-term human-robot relationships with a receptionist robot [5] and a delivery robot [9]. Researchers also investigated the commercial robots in the workplace; Mutlu et al. studied how the environmental factors in the workplace influence interaction with the robot [11].

Our work is closely related to studies that investigated integrated robotic systems for interacting with people in human-centric environments. Kanda et al. studied a humanoid robot that could guide people and advertise shops in a shopping mall [8], and Bohus et al. deployed a Nao humanoid robot that could provide directions to people in an office building [1]. Unlike the systems studied in those two, the robot we studied is not anthropomorphic and the interactive behaviors we studied were meant for cognitively lightweight interactions, e.g. short screen interactions. In addition, we focused on behaviors that used the robot's on-board sensors, rather than a network of off-board sensors in the environment. Our work is also closely related to systems for developing or generating socially interactive behaviors [4, 6]. While their evaluations were conducted in lab environments, our work presents evaluations based on field studies.

[4] (see the discussion section in [7])

Table 1. Summary information about service industry properties studied in this paper.

Property	Type	Used since	Point of contact	Requested applications	Target areas
A	Airport	2/2016*	Corporate executives, Customer satisfaction manager	People delight, Service recovery	Indoor garden, Baggage claim, Immigration hall
B	Hotel	1/2015	Hotel manager, Business consultant, Front desk supervisor	People delight, Mobile kiosk, Demo	Lobby, Bar
C	Hotel	6/2015	Guest service manager, Sales & marketing director	People delight, Service recovery, Mobile kiosk	Lobby
D	Hotel	7/2015	Hotel manager, Guest experience manager	Service recovery, Demo	Lobby, Breakfast area
E	Hotel	8/2015	IT manager, Area general manager	Mobile kiosk	Lobby

* Used since field study I

3 Formative Study

Although the Savioke Relay robot was built for a specific application (room service in hotels), it can be considered as a generic platform with a wider range of applications. The goal of our formative study is to discover potential applications that are desirable for existing Savioke customers and inform the design of a rapid programming system for creating those applications. To that end, we gathered information from five customers from the service industry.

3.1 Data Collection

Information about the properties we studied in this paper is summarized in Table 1. Property A was an airport in Southeast Asia and the rest were hotels in the San Francisco Bay Area. Unlike the other properties, Property A had not used the robot prior to the first field study. Additionally, the target areas considered by Property A were larger ($\approx 4000m^2$) and more crowded than the areas considered by the other properties ($< 200m^2$).

We analyzed meeting notes and email exchanges between Property A and Savioke employees. The meeting notes were collected from two meetings held at Savioke headquarters and one meeting held at the airport in 2015. During those meetings, Savioke employees surveyed the target areas in the airport and performed demos of the Relay doing deliveries. The two groups also brainstormed potential applications together.

For Properties B–E, we analyzed the field notes taken by a Savioke customer satisfaction manager during regular checkup visits in February 2016. As part of the visit, the customer satisfaction manager met with one or two hotel representatives individually and asked them for (i) general feedback on using Relay, (ii) their wish list of new robot applications, and (iii) feedback on support infrastructure. Our analysis focused on the wish list notes.

3.2 Use Cases

We categorized the requested applications by their use cases.

People delight: All properties wanted to provide a unique experience to visitors using the robot. 3 properties specifically proposed applications designed to make their customer experience more delightful. Property A proposed an application in which the robot would approach passengers in an indoor garden area, offering them snacks or volunteering to take their picture. Properties B and C wanted the robot to roam around in their lobby and bar areas to encourage lightweight interactions with guests. Example interactions that Property B suggested included playing a game of rock-paper-scissors or sharing a joke.

Service recovery: 3 out of 5 properties wanted to use the robot to catch unsatisfied customers before they left the building. Property A requested that the robot approach passengers in the baggage claim area whenever the unloading of baggage was delayed. They wanted the robot to explain the situation and placate potentially frustrated passengers. Previously, this task was done by the airport staff, who were often not treated well by frustrated passengers. Property C wanted the robot to approach guests who were leaving the hotel, in order to ask them about their stays. Property D wanted the robot to navigate to the hotel's breakfast area and ask guests about their stays.

Mobile kiosk: 3 of 5 properties requested applications that resembled an information kiosk. Properties B and C wanted the robot to visit a couple of highly visible locations (e.g., a location near the front entrance or the elevators), and display a series of screens encouraging interaction when people were around. They said that displaying information about the robot or the hotel would be useful, as it could trigger guests to use the delivery service or other hotel amenities in the future. Property E, which had the robot's docking station in the lobby, requested that the kiosk mode run while the robot was charging.

Demo: Properties B and D requested a guest-facing demo application. They often had to manually control the robot to show it in action to curious customers. Property B suggested that the application include a navigation demo and an introduction about its delivery service. Property D wanted control over how the application would be activated; they did not want guests to be able to trigger the demo, as it could interfere with actual deliveries that needed to be done.

Overall, we make the following observations:

- Having first hand experience with the Relay robot (except Property A), Savioke customers had realistic requests.
- Although their requests were similar and could be broadly categorized as above, they each had specific, custom requirements.
- Many of the requested applications involved interactions with humans.

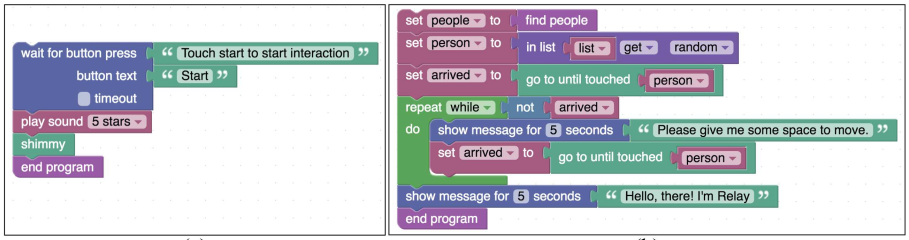

(a) (b)

Fig. 1. Example applications written in *iCustomPrograms*. (a) Simple interactive application; the robot first waits for a user to engage in interaction by pressing a button. It then plays a sound and shimmies in response. (b) Approaching a person application; the robot finds nearby people using **findPeople**, randomly selects a person, and approaches them. The **goToUntil** primitive returns *true* if the robot successfully reaches the destination and *false* if it is interrupted by a person tapping its touchscreen.

4 *iCustomPrograms*

The software for the original room service functionality of the Relay robot was developed by Savioke's team of engineers and programmers. This team could implement many of the functionalities requested by customers (Sec. 3.2). However, given the diversity of requests from customers and the time it takes for custom software to be developed and deployed, this approach would not be scalable as the number of customers increase. Instead, Huang et. al developed *CustomPrograms* [7] to enable non-technical Savioke employees (e.g. marketing representatives, customer satisfaction managers, designers) as well as customers (e.g. hotel staff) to program the Relay robot. In this paper, we extend *CustomPrograms* with an emphasis on *interactive* behaviors, which were a part of the applications requested by customers.

4.1 CustomPrograms

CustomPrograms allows users to build applications for robots by composing a set of capabilities, known as primitives, with general-purpose programming constructs like variables, loops, conditionals, and functions [7]. Applications are started manually and end when there are no more instructions to run.

Huang et al. implemented *CustomPrograms* for the Relay robot, including four categories of primitives: navigation, screen interaction, lid control, and battery state. The main navigation primitive was **goTo**, which made the robot navigate to a given location. The **shimmy** primitive was a short side-to-side swaying to convey happiness. Screen interaction primitives included displaying messages (**displayMessage**), receiving user input (e.g., **askMultipleChoice**, **askPasscode**, **askRating**), and playing non-anthropomorphic sounds (**playSound**). The other primitives controlled the robot's lid or read the battery level.

CustomPrograms can be used to program simple interactive applications. For example, Huang et al. developed a demo application in which the robot went to several predefined locations, and offered a snack to nearby people.

4.2 Supporting People-Aware Behaviors

One key capability that was needed for the applications described in Section 3.2 was the ability to find and navigate to people. In *CustomPrograms*, the robot could only go to predefined locations and wait for people to interact with it. Hence, we created *iCustomPrograms*, a modified *CustomPrograms* that included a **findPeople** primitive. **findPeople** returned a list of locations where people were detected. This enabled users to create applications in which the robot approached people or recognized when someone was walking towards it. Example applications written in *iCustomPrograms* are shown in Fig. 1.

4.3 Field Study I: Airport Trials

In February 2016, we visited Property A for a two-week period. We used *iCustomPrograms* with their staff to develop two interactive applications:

Passenger Delight: In this application, the robot visited waypoints in the airport's indoor garden area. At each waypoint, it waited and approached people around it. To encourage interaction, it played a beeping sound and displayed an on-screen greeting. When a passenger started the interaction, it played a sound, displayed greeting messages, and opened its bin with snacks inside. It also did a shimmy, as an enthusiastic gesture. It then said goodbye using on-screen text and a beeping sound, and moved to the next waypoint or detected person.

Service Recovery: This was a similar application to *Passenger Delight*, but it ran in the baggage claim area whenever the unloading of baggage was delayed. Compared to *Passenger Delight* the robot was more professional. For example, we removed the enthusiastic shimmy and changed the on-screen text to politely explain the baggage situation.

We ran four trials to evaluate these applications. *Passenger Delight* was deployed for the first two trials, which took a place in the indoor garden. *Service Recovery* was deployed for the last two, which took place in the baggage claim hall. Each trial lasted 3 to 4 hours. To maximize engagement, the trials were run during Chinese New Year weekend.

For each trial, 2 to 5 airport staff members and 1 or 2 Savioke employees were present, monitoring the robot from less than 15 meters away. We tried not to interact with the passengers; however, airport staff did intervene when unexpected events happened. Examples included children acting mischievously with the robot or encountering non-English speakers.

4.4 Findings

We recorded observations and notes from all the meetings and trials, and conducted follow-up interviews with personnel. We identified three themes:

Problems with approaching people: The robot had difficulty approaching people naturally. While navigating to the location of a detected person, curious crowds of people would often form around the robot, surrounding it. The robot was not

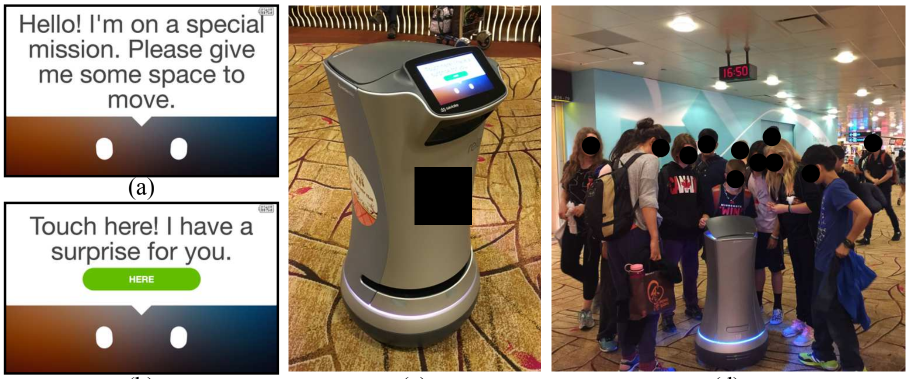

(b) (c) (d)

Fig. 2. Pictures from the trials held at Property A. (a) On-screen text when the robot was navigating and (b) encouraging people to interact with it. (c) The robot with the Chinese New Year decal on. (d) The robot interacting with passengers.

programmed to recognize this situation, and continued trying to navigate to its goal location, instead of starting the interaction. This often led to people getting confused or frustrated with the robot. We resolved the issue using on-screen text asking for a clear path (Fig. 2), which helped the robot go through crowds.

Initiating interactions via movements and sounds: The robot initiated interactions with users in unplanned ways, e.g., just by driving around. People would follow the robot, and even tap the screen while the robot was still moving. Additionally, people noticed the robot when it played sounds.

Desire for richer control over interactive elements: The airport staff wanted the robot to have a "brighter" or "more playful" personality. They asked to have more sounds and pre-programmed movements, as well as a way to choreograph them together to make the robot look "happier." They witnessed some passengers saying "Hello" and "Goodbye" to the robot, and requested text-to-speech so the robot could respond. Finally, they also requested the ability to play background music and to format text (e.g., changing font size or adding line breaks).

5 Enhancements and Evaluation of *iCustomPrograms*

5.1 System Enhancements

Based on our findings, we enhanced *iCustomPrograms* as follows:

Supporting touch-to-interact: As described in Section 4.4, the robot experienced problems with people surrounding it while navigating. This was because the robot was programmed to not respond to screen input until it was done navigating. To address this, we added the **goToUntil** primitive to *iCustomPrograms*. **goToUntil** was like **goTo**, but it stopped navigating when someone touched the robot's screen. Using this made the robot behave more naturally with crowds, as they could now get the robot's attention by tapping its screen. Fig. 1b is an example application illustrating the touch-to-interact behavior.

Richer control over interactive elements: We enabled users to format on-screen text in *iCustomPrograms* using HTML. We also updated the **playSound** primitive to play sounds asynchronously. In the original *CustomPrograms*, sounds were played synchronously, meaning that the robot could not navigate or respond to screen input while a sound was playing. With our change, *iCustomPrograms* supported playing long-running background music, as well as choreographing sounds with movement or on-screen interactions.

5.2 Improved Social Applications

We developed two new social applications using the updated *iCustomPrograms*.

People Delight was based on the *Passenger Delight* application, but was designed for use in more than just airports. It used **goToUntil** to start the interaction if a person tapped the robot's screen while it was navigating. The on-screen text was adjusted to be more property-agnostic. We also added more sounds and in-place movements to attract more attention to the robot and make the main interaction more lively.

The second application, *Mingle in Place*, was developed for smaller properties that did not want to have the robot navigating around continuously. When the application was launched, the robot navigated to a preset location and displayed three options. The first option was a demo, in which the robot described itself and its delivery service. The second option was to have the robot tell a joke. The third option was to pose for a picture, in which the robot displayed "Cheese!" on its screen. The robot played 3 to 4 different sounds during the interaction and made in-place movements. If no one interacted with the robot for over a minute, the robot attracted attention by rotating left and right while making a whistling sound. The application stopped when the battery went below a predefined threshold, or when the operator canceled it.

5.3 Field Study II: Trials at Five Properties

The first author demoed the *People Delight* application to Property A in February 2016 and provided a manual describing how to use it. Between March and May 2016, a Savioke customer satisfaction manager repeated the procedure with *Mingle in Place* at Properties B–E. For the hotels, room service deliveries continued to be the primary function of the robot. The properties ran *Mingle in Place* on the robot when they wanted; we did not ask them to do so for the purposes of the study.

We conducted semi-structured interviews with a staff member from Properties A, B, C, and E, after they had used the applications for at least 2 weeks. All interviewees said that the robot had successfully interacted with visitors. Property A reported that during Easter weekend, the robot was used from 10 a.m. to 6 p.m., interacting with about 500 passengers. They also said that children 7 and up, young adults, and group travelers interacted most with the robot.

Properties B, C, and E reported that their guests enjoyed interacting with the robot, especially on weekends. As with Property A, they noted that families with children and groups were most interested in interacting with it. Property E pointed out that their robot was often too busy running deliveries to use *Mingle in Place*. 3 out of 4 interviewees said that the sounds and movements of the robot helped initiate interaction with people. However, 2 out of the 4 interviewees wanted the robot to be even more interactive and have more sounds.

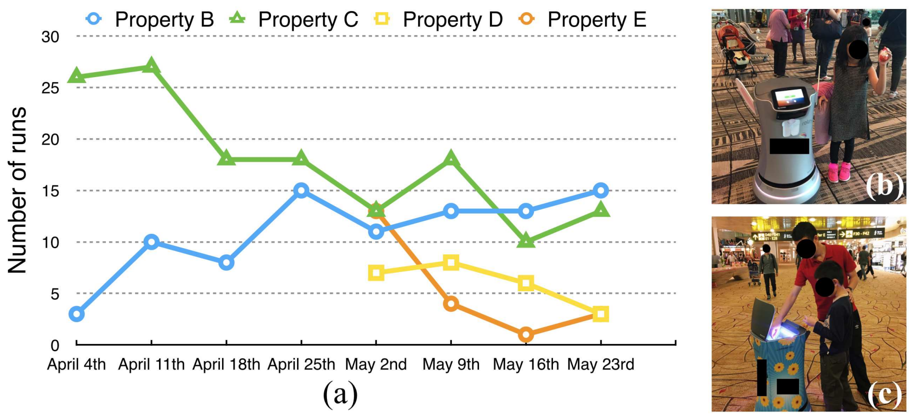

Fig. 3. (a) Weekly *Mingle in Place* usage by property. (b) Pictures taken at Property A during the field study over Easter and (c) another local holiday weekend.

We recorded the number of times the Properties B–E ran *Mingle in Place*, shown in Fig. 3a. Due to logistical problems, we could not collect any usage measurements from Property A. Properties B and C ran *Mingle in Place* the most overall. These two properties had proposed the *People Delight* application during our formative study (Section 3.2). For most properties, the number of runs peaked in the first two weeks and gradually decreased after, which could indicate a novelty effect. However, as we heard from Property E, low usage of the application could be due to the robot being busy with room service deliveries instead. And, during the study period, all of Properties B–E had run the application at least once a week.

Although we lack usage measurements from Property A, they reported that they used *People Delight* the most during Easter weekend and over another local holiday weekend in May 2016. They also said that staff members used *iCustomPrograms* to customize the contents of *People Delight* for each occasion, and applied festive decals to the robot's body (Fig. 3b,c).

6 Conclusion

This paper's formative study showed that service industry workers desired socially interactive behaviors for their robots. We presented *iCustomPrograms*, a system for developing such behaviors. In our first field study, we discovered important attributes for better interactions. Robots naturally attract attention, so they must be equipped with crowd-aware navigation and interactions. In our experience, service industry workers wanted rich control over interactive elements, like having more sounds, movements, and text formatting capabilities. With such enhancements in *iCustomPrograms*, we developed and deployed social applications to five real-world service industry properties. Our users not only actively used the applications, but also reported interesting observations about how people interacted with the robot. This information could lead to future improvements and ultimately to more socially interactive robots in the field.

References

1. Bohus, D., Saw, C.W., Horvitz, E.: Directions robot: in-the-wild experiences and lessons learned. In: International Conference on Autonomous Agents and Multi-Agent Systems (2014)
2. Burgard, W., Cremers, A.B., Fox, D., Hähnel, D., Lakemeyer, G., Schulz, D., Steiner, W., Thrun, S.: Experiences with an interactive museum tour-guide robot. Artificial intelligence (1999)
3. Fong, T., Nourbakhsh, I., Dautenhahn, K.: A survey of socially interactive robots. Robotics and autonomous systems (2003)
4. Glas, D., Satake, S., Kanda, T., Hagita, N.: An interaction design framework for social robots. In: Robotics: Science and Systems. vol. 7, p. 89 (2012)
5. Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A.C., et al.: Designing robots for long-term social interaction. In: IEEE/RSJ International Conference on Intelligent Robots and Systems
6. Huang, C.M., Mutlu, B.: Robot behavior toolkit: generating effective social behaviors for robots. In: ACM/IEEE international conference on Human-Robot Interaction. pp. 25–32. ACM (2012)
7. Huang, J., Lau, T., Cakmak, M.: Design and evaluation of a rapid programming system for service robots. In: ACM/IEEE International Conference on Human-Robot Interaction (2016)
8. Kanda, T., Shiomi, M., Miyashita, Z., Ishiguro, H., Hagita, N.: An affective guide robot in a shopping mall. In: ACM/IEEE International Conference on Human-Robot Interaction (2009)
9. Lee, M.K., Kiesler, S., Forlizzi, J., Rybski, P.: Ripple effects of an embedded social agent: a field study of a social robot in the workplace. In: ACM SIGCHI Conference on Human Factors in Computing Systems (2012)
10. Linder, T., Breuers, S., Arras, K.O.: On multi-modal people tracking from mobile platforms in very crowded and dynamic environments. In: IEEE International Conference on Autonomous Robot Systems and Competitions (2016)
11. Mutlu, B., Forlizzi, J.: Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In: ACM/IEEE International Conference on Human-Robot Interaction (2008)
12. Perez, J., Caballero, F., Merino, L.: Integration of Monte Carlo Localization and place recognition for reliable long-term robot localization. In: IEEE International Conference on Autonomous Robot Systems and Competitions (2014)