Chapter 14: A Communication Paradigm for Human-Robot Interaction During Robot Failure Scenarios

Daniel J. Brooks¹¹, Dalton J. Curtin²², James T. Kuczynski³², Joshua J. Rodriguez⁴², Aaron Steinfeld⁵³, and Holly A. Yanco⁶²

¹Toyota Research Institute ²University of Massachusetts Lowell ³Carnegie Mellon University

June 16, 2020

¹dan@tri.global ²daltoncurtin@gmail.com ³jkuczyns@cs.uml.edu ⁴Joshua_Rodriguez@student.uml.edu ⁵steinfeld@cmu.edu ⁶holly@cs.uml.edu

Abstract

When autonomous robot systems experience failures, communication about the failure to both the people responsible for the robot and to people who happen to be nearby is critically important. New robot users as well as bystanders might not be familiar enough with a system to tell whether a robot is working properly and experienced robot operators might not notice signs of trouble due to being out-of-the-loop. Thus, an important feature for robots will be the ability to communicate failure to humans when failures occur. In this paper, we describe a study conducted with a smartphone-based feedback system we designed to explore push and pull forms of communication. We found the communication methods improved participants' understanding of robots' state, increased their confidence in interacting with the robots, and allowed them to remotely monitor and control the robots. In keeping with the theme of this book, shared context between humans and machines improves performance.

14.1 Introduction

Robot systems, including self-driving cars, delivery drones, and cleaning robots, are becoming more common in public settings. Soon, people will interact with and live alongside these and other types of autonomous robotic systems on a regular basis. Due to the scale and complexity or these robot systems, we cannot expect every interaction to be flawless even after systems have matured. For example, a robot performing a necessary behavior that is perceived as in-explicable or unpredictable can have a detrimental effect on people's situation awareness and lead to negative user experiences.

Such context issues will affect not only the robots' users but also bystanders, who may have marginal awareness of or interest in the robots' capability or mission. Humans will increasingly be in situations requiring them to make decisions about unsupervised and unfamiliar systems, some of which may be critical to their own safety. To make matters worse, there are currently no standards – de facto or otherwise – to serve as a guide for allowing people to communicate with or to influence the behaviors of these machines. Therefore, we believe that robots need efficient and understandable methods for bidirectional communication, even at a basic level, which could be used for robot operators or for people who are bystanders to the robot. As a first step towards a solution to this need, this chapter describes our findings in an experiment using push and pull notifications on a smartphone to communicate robot information to minimally trained participants.

14.2 Related Work

There will likely never be perfectly reliable robots, so strategies are needed for mitigating the consequences of failure. Taxonomies have been developed in support of this challenge that categorize faults and provide insight into the many

 $\mathbf{2}$

complex ways a system could fail [3, 24]. Likewise, attributes of "dependable" systems are available [19] and there is increasing interest in translating internal representations of robot performance into human-understandable forms. Effective methods from prior work include providing advanced warning or confidence feedback, apologizing after failure, asking for help from bystanders, and failure-specific natural language requests [5, 17, 18, 22].

Unfortunately, there are also examples where recovery strategies produce negative effects. For example, inexplicable robot behavior can lead to misalignment of blame and humans often fail to recognize their own correctable mistakes [16]. Likewise, robot assignment of blame can lead to very negative reactions and lower trust [6, 14].

There has also been work on how to convey failure to humans. Failures can be expressed implicitly if the robot is able to clearly communicate its intentions in a way that can be contrasted with physical behavior when an error occurs. For example, if a drone were equipped with a light ring direction indicator (e.g., [25]) that was indicating a straight flight path while the drone was translating to the side, users familiar with the drone's normal operation would be able to immediately discern that something was not right. Likewise, a robot with an arm could point the limb where it believed a person wanted it to go, thereby providing an opportunity to intervene [8]. However, the intent which the robot seeks to convey with these methods might not be clear, especially if the recipient is a bystander not familiar with the system.

Explicitly requesting human assistance during failure can also be a valuable failure mitigation strategy in certain situations. For example, Cha et al. [4] summarized the process of asking for help as having three phases: 1) getting someone's attention, 2) indicating to the person that help is needed, and 3) conveying the request for help. Likewise, participants in a human-robot teaming study by Barber et al. [1] expressed an interest in robot failure notifications and communication of robot performance. It is worth noting that self-assessment of performance is still an open challenge for autonomous robots, although this topic is attracting increased attention and new methods are likely to be available in the near future.

For example, there has been work on introspective systems capable of explaining why a robot behaved in a particular manner by tracing and logging the flow of information through a system, and keeping track of which pieces of data were used in making progressively higher level decisions [2]. However, providing users with information about the cause of a failure could also make the situation worse. For example, Kim and Hinds [16] found that robots that attempt to explain their ambiguous actions and errors can actually decrease people's perceived understanding of the system.

There also appear to be nuances in how and when explicit requests of human assistance will be honored. For example, in one study, participants in a public kitchen area were asked for assistance with coffee preparation by an approaching robot [9]. Only half of the participants complied with the robot's request. The vast majority of the people who helped were not busy concentrating on another task (one of the experimental conditions), while most who were busy ignored the robot, tricked it into thinking they had given it the coffee so it would go away, or shut the door to keep the robot out.

Deciding who to ask for help is also important. Asking the same person for help frequently could quickly become annoying and robots which adapt to training are rated as not annoying [15]. Rosenthal et al. [22] addressed this issue by distributing the burden across an office hallway and anticipating who might be available based on prior behavior. After a few days, many people closed their office doors to avoid being asked by the robot for assistance in moving chairs, writing notes, and informing it of the room number. Similarly, Srinivasan and Takayama [23] revealed a number of important factors, like politeness, that impact human willingness to help or train robots, including resistance by participants to allocate significant effort to train robots that are developed and maintained by others. This suggests the general public may have differences of opinion between assisting and teaching robots.

There are successful interaction models where humans are willing to assist openly imperfect robots. For example, Yasuda and Matsumoto [26] hypothesized that people may relate well to imperfect robots, viewing them as similar to children or infants who try but fail in their efforts. They experimented with a robot trashcan that would sometimes spill garbage, but lacking manipulators to clean up after itself. Thus, whenever this would occur, the robot would politely ask a person to pick up the trash for it. The majority of people found the experience to be positive, even when the robot spilled trash. While robot failures could be catastrophic in safety critical applications (e.g., self-driving cars, search and rescue robots) [13], it is likely that many of the robot systems developed in the near term, and into the future as well, will experience failures and will need to ask for assistance from people or other robots.

14.3 Interaction Design

Smartphones have rapidly become a ubiquitous technology based on their value in a wide variety of tasks. Due to their market saturation, we believe these devices to be an ideal proxy through which humans and autonomous systems could communicate with each other.

To explore methods for human-robot communication, we developed used smartphones, with two types of interaction methods. The first type uses *pullstyle interactions* to enable participants to query information or communicate with the robot. Pull-style interactions wait in the background for the user to initiate interaction, rather than interrupt the user. The user can view these by tapping on a background notification which will appear as seen in Figure 14.1 and Figure 14.2. The second type used *push-style interactions* (Figure 14.3), where the robot initiated communication with nearby participants. Push-style interactions interrupted the participant and were thought to be useful when alerting about hazards or requesting help. These interaction methods were used in conjunction with a balloon popping game that we created for use in the experiment.



Figure 14.1: Pull-style notifications: The left image shows the pull notifications accessed from within the balloon game. The right image shows the pull notifications accessed by opening the smartphone app directly.

We note that robot vacuum manufacturers have been shipping their products with accompanying apps for a few years. For example, Neato Robotics' robot vacuums [21] have an app that provides four push notifications: "Done Cleaning," "Low Battery," "Base Connected to Power," and "Move Base to New Location" [20]. As another example, iRobot's robot vacuums [11] have an app that provides two types of push notifications: "Done Cleaning" and error messages [10]. iRobot's robot vacuum users can also use the "Care" section of the app to see the status of the robot, its bin, and its extractors [12]; these notifications are pull notifications. While these manufacturers have developed apps to be used with their robots, we are not aware of any published studies that investigate how these types of interactions influence human-robot interaction for the context when a robot needs assistance.

For both interaction types in the smartphone app that we designed for our experiment, robots were identified with a unique name and an icon that looked like the robot, in order to help identification of the robot(s); see Figure 14.4 for the set of robots used in the experiment. In the pull-style interactions, the robot's icon also had a small status icon overlaid on its lower right corner, to



Figure 14.2: Acting on pull-style notifications: The left image shows one robot's help screen reached from the pull notifications within the app and the right image shows the robot's help screen after the reset.

provide immediate feedback about the robot's state, without needing to open the message.

14.4 Experiment Methodology

The experiment described in this section represents our initial investigation into the use of smartphones as a platform for establishing a ubiquitous communication paradigm designed to allow participants to gain basic information from and interact with a variety of autonomous robots. In this experiment, we used several types of robot vacuum cleaners, shown in Figure 14.4), as our platforms to provide a clearly describable task and to allow interaction with multiple robots in a fairly small space. Four of the five robots were iRobot Roomba robots, each a different version to provide visual distinction between the robots; the fifth was a Neato robot, rounded on the rear side and flat on the front. The Roombas have a width of 34 cm while the Neato is 33 cm wide; the Roombas are 9.2 cm tall and the Neato is 10.16 cm tall. The enclosure in which the robots were

6



Figure 14.3: Push-style notifications: The left image shows push notifications on the smartphone's lock screen. The right image shows a push notification from a robot over the balloon game.

confined was 2.40 m by 2.40 m, as shown in Figure 14.5.

This study was approved by the University of Massachusetts Lowell's Institutional Review Board. This research has complied with all relevant national regulations and institutional policies.

14.4.1 Communication Protocol

The communication between the robots and the user's phone is performed using an underlying communication protocol built on Bluetooth Low Energy (BLE) and an Android app that makes use of this protocol to allow users to interact with the robots. One of the initial challenges addressed was the process though which smartphones and robots could discover each other. The Bluetooth V4.0 specification, also known as Bluetooth Low Energy, was well suited for this task as a relatively short range communication protocol (about 60 m, unobstructed) that utilizes broadcast messages to allow nearby devices to advertise their presence and the type of service(s) they provide. The software used to facilitate



Figure 14.4: The set of robots used in the experiment, with four Roomba models and one Neato robot. Neato (far left) and the Discovery model Roomba (far right) were disabled to include "broken" robots in the experiment design.

these communications is publicly available¹.

In order to communicate over Bluetooth, each robot received hardware upgrades, implemented such that there were no visible modifications to the robots' external appearances. The additional electronics that were added shared the same power source as the rest of the robot, thus eliminating the need for additional or modified charging systems. The primary addition was a Raspberry Pi Zero with a Bluegiga BLED112 USB Bluetooth Low Energy dongle, allowing the developers to bypass the Linux kernel's Bluetooth stack and gain greater control of the behavior of the Bluetooth radio.

Status messages consisted of formatted text, optionally attached media resources (an image or icon), a list of commands or responses a user could choose to send back, and whether or not the message should be treated as a *push-style* interaction and displayed as a pop-up. This data was organized according to a JSON definition and serialized. For the commands, the robots offered users the same controls through the Bluetooth interface as the physical button controls found on the robots. The status information sent to the users was derived from the realtime hardware state, or from simulated problems which were generated by the experiment interface.

 $^{^{1}}$ GitHub code repository https://github.com/uml-robotics/vacuum_trust contains the software used in this experiment with the exception of data loggers and peripheral programs.



Figure 14.5: One of the authors demonstrating what a participant might do in the cleaning zone to assist a robot.

14.4.2 Design

In the experiment, we asked participants to perform two simultaneous tasks: (a) manage a fleet of robot vacuum cleaners while (b) playing a simple video game on a smartphone as a secondary task.

The video game task was a simple, skill-less game of balloon popping. Animated balloons drifted up from the bottom of the screen to the top and "popped" when touched, earning the participant points. Participants needed to pay attention, as balloons marked with skull and crossbones symbols would take away earned points if popped. This game was selected because it does not require prior experience. It provided a secondary task that required monitoring, so that participants would not be able to watch the robots constantly.

In parallel, participants were asked to use three apparently standard robot vacuums, which were actually modified and programmed to exhibit the desired behaviors, to collect 6 mm x 9 mm plastic beads scattered around the floor inside an area fenced off with a 2.4 m^2 low wood frame. Participants were required to remain outside of this area and were not permitted to collect the beads themselves, forcing them to use the robots to accomplish the task. Participants could interact with the robots at the edge of the cleaning zone, as shown in Figure 14.5.

In order to track time spent physically interacting with or observing the robots, participants were only allowed to play the game when inside a marked



Figure 14.6: One of the authors demonstrating what a participant would be doing if in the game playing zone. The participant is not able to see the cleaning zone from this area, due to the cardboard wall obstruction, as can also be seen in the prior figure.

game-playing zone that did not have a view of the robots in the cleaning zone (see Figure 14.6). The game interface was disabled when exiting the game-playing zone. If present in the current experimental condition, the participant could use the app to asses the robots' status or be notified by push notification when the robots' needed assistance. Alternatively, the user would need to leave the game-playing zone to get the similar information or to interact with the robots.

The robots were secretly modified in a way that allowed them to appear as if they were still standard robot vacuum cleaners. They were programmed to intentionally exhibit various problems with their functionality which the participants would need to address, while the video game acted as a distracting secondary activity for diverting user attention from the robots.

The flaws exhibited by the robots were refusing to start, requiring a reset, returning to the charging dock prematurely, and requesting a dustbin cleaning early. These behaviors were selected for their plausibility to occur in a real scenario, as well as the amount of effort required to correct them. For example, the "dead" robot would never recover, while the one returning prematurely would immediately return to its duty once prompted via the app. Our objective was to gain insight into how participants would manage the two tasks and the robots when provided with access to the different communication styles (i.e., push and pull). Since we wanted to understand the impact on novices, participants were not provided with any training on the use of the app; we also limited training on the robots themselves, as described below.

We used a within-subjects design where each participant experienced two runs – one with only the manufacturer's on-robot interfaces (i.e., smartphone communication *disabled*) and one with the manufacturer's robot interfaces supplemented with our communication app (i.e., smartphone communication *enabled*). In both cases, the participants had access to the on-robot manufacturer's button interfaces.

Our participants were divided into four equally sized groups that were counterbalanced. Groups 1 and 3 shared the same group of robots between runs 1 and 2 with the run the the App was enabled alternated between these groups. Groups 2 and 4 were similar in that they also shared the same group of robots but the group was different from 1 and 3. Groups 1 and 2 shared the same order in which the App was enabled, as did 3 and 4 (Table 14.1).

Participants were not briefed or trained on how symbols on the robot's hardware or beeps the robots communicated to the user would correspond to what was causing the robots to be unable to operate, nor were participants showed how to start the robots. They were shown all the robots that would be used during the experiment including the robots that were to be swapped in on the second run and how to empty the dust bins on each robot. Similarly, participants were not briefed or trained on the use of the smartphone app for robot notifications.

Participants were given 6.5 minutes to both play the balloons game and use the robots to collect beads. A digital clock in the game-playing zone and a timer in the balloons game provided awareness of this period. Participants were responsible for keeping track of the time remaining by using either a digital clock showing the game time that was positioned inside the game-playing zone, or by using the game timer inside the balloons game (which also showed the experiment time).

While participants earned money based on points scored in the game, the total was adjusted by the fraction corresponding to the percentage of beads collected by the robots (i.e. if they collected 70% of the beads they would get to keep 70% of the points scored in the game). This percentage was determined by measuring the weight of the collected beads, as compared to the known amount scattered on the floor at the start of the run (roughly 100 grams). At the end of each run, the beads collected by the robots were measured by weight to calculate the percentage collected, while the remaining beads were removed from the cleaning area by the experimenter. Participants would lose an additional 100 points (approximately 45 seconds worth of balloons game playtime) for each robot that was not on a charging station when time ran out. Participants received compensation based on the higher of their two final scores. Each person received \$5 for simply completing the study, and could earn up to an additional \$10 based on their performance, using a score based on the time that the two working robots spent running and the score of the balloons game.

In addition to an informed consent form before the experiment, participants

were asked to fill out a pre-experiment questionnaire, two post-run questionnaires, and a post-experiment questionnaires to get their opinions and perceptions of the experiment.

14.4.3 Hypotheses

Our hypotheses in this study were as follows:

Hypothesis 1 (H1): Participants would be able to determine which robot they were communicating with using our system, despite similarities between robots.

Hypothesis 2 (H2): Participants would be able to retrieve information about the robots they were working with, identify solutions to problems faster, and allocate their time more appropriately when using the smartphone-based interface compared to only having the default manufacturer interfaces on the robots.

Hypothesis 3 (H3): Participants would prefer having access to the additional information provided by our smartphone-based interface over using the default manufacturer interface on the robot alone.

14.4.4 Independent Variables

This experiment used a within-subjects (repeated measures) design in which each participant performed two runs. While both runs included the manufacturers' default interfaces on the robots themselves, our smartphone-based communication was only enabled in one. Between runs, the experimenter replaced two of the three robots used in the previous run in full view of the participant as part of "resetting the task," while the participant filled out the post-run questionnaire. (The two robots not being used in a particular run were on charging stations behind the participant in the game playing zone, as shown in Figure 14.6.)

Conditions were counterbalanced by assigning participants into one of four experimental conditions corresponding to the two independent variables in the experiment: the order in which the smartphone app support was used in the runs and the order in which two different starting configurations were used.

The experiment used 5 robot vacuum cleaners: four iRobot Roombas and one Neato XV11. The four Roombas were two working Roomba 500 models $(R_A \text{ and } R_D)$, a working Roomba 600 model (R_C) , and a non-working Roomba Discovery model (R_B) . The Neato XV11 (R_E) was intentionally programmed to not work. Three of these vacuums were used during the first run, after which two of the three were replaced before starting the next run. There were two combinations of robots that were switched between, each of which consists of two "working" robots and one "non-working" robot: Group 1 consisted of R_A , R_B and R_C while Group 2 consisted of R_A , R_D , and R_E .

\mathbf{App}			Robot Starting Conditions				
Cond	Support	Run	R_A	R_B	R_C	R_D	R_E
1	Enabled	1	Easy	Dead	Help		
1	Disabled	2	Help			Easy	Dead
2	Enabled	1	Easy			Help	Dead
2	Disabled	2	Help	Dead	Easy		
3	Disabled	1	Easy	Dead	Help		
3	Enabled	2	Help			Easy	Dead
4	Disabled	1	Easy			Help	Dead
4	Enabled	2	Help	Dead	Easy		

Table 14.1: Experiment conditions, as described in Section 14.4.4

Each of the three robots exhibited a different level of functionality: one robot was "easy" to start, simply requiring the push of a button; one required "help" from the participant before it would start running (it needed to be reset); and the last one played "dead" and would never start working (Table 14.1). The two "working" robots performed their default cleaning behaviors except that the length of time they ran for was shortened and some "problems" were introduced. Two minutes after a robot started cleaning, it would automatically start returning to its dock. After the "easy" robot returned, it would require the participant to come to empty its dustbin before it would be able to start cleaning again. In contrast, the "help" robot could immediately be told to resume cleaning (even before it finished returning to its dock) when its two minutes were up, and never required the dustbin to be emptied.

The robot that needed to be reset flashed a red LED ring around the power button, illuminated a red error symbol in the shape of a circle with an exclamation mark in the center of it, and would periodically play a distinct error tone until the participant pressed and held down on the power button. When the smartphone app was enabled, this robot sent a push-style interaction, causing a message to appear on the smartphone. The reset event occurred at the beginning of each run.

The robot that needed its dustbin emptied illuminated a yellow LED ring around the power button, flashed a blue LED labeled "dirt detect" and would occasionally play a different error tone until the dustbin was removed. When the smartphone app was enabled, participants could view information about the full dustbin on the robot's status page (a pull-style interaction).

Time spent trying to make the third robot work was wasted. One of the dead robots had no lights on and showed no indication it even had power; the other dead robot had a single lit LED light, but showed no other signs of being functional. The dead robot could be viewed using pull-style interactions in the smartphone app, which would reveal the robot's broken status.

14.4.5 Dependent Variables

Aside from the three questionnaires, data was logged on the robots and smartphone, in manual notes, and video recordings. After each run, we documented the resting positions of the robots, the percentage of beads collected, and the game score. The post-run questionnaire included the NASA's Task Load Index (TLX) Questionnaire [7], as well as questions about participant confidence in the app and where participants thought app information originated.

Logged data included the time in game, the number of times the game was started and stopped, when and how often they left the game zone, and the game score over time. App specific data was also captured, including the number of times each robot's app page was accessed via a pull-style interaction, the time spent viewing each robot's app page, the number of times background notifications were observed, and whether the participant explicitly agreed or declined (via the app dialog buttons) to help the robot.

14.4.6 Participants

Twenty people (14 men, 6 women) between the ages of 18 and 33 participated, all of whom had previous experience with smartphones. Two participants had prior experience using Roombas and one person had used a Philips robot vacuum, while the remaining seventeen participants had no prior experience with robot vacuum cleaners.

14.4.7 Analysis Methods

The majority of the analyses were 2x4 mixed-groups factorial ANOVAs on app status (enabled, disabled) and one of the given conditions (1 through 4) counter balancing the starting conditions of the experiment (Table 14.1). These statistics were computed in R, using the 'psych' and 'ez' statistics packages. An alpha of 0.05 was used on analyses.

14.5 Results

14.5.1 Robot Usage

The app availability had a significant main effect for the combined time the robots spent cleaning [F(1, 16) = 17.052, p < 0.001] (Figure 14.7). There was a significant interaction between the experiment condition and the presence of the app [F(3, 16) = 3.49, p = 0.04], however there was no main effect for condition by itself [F(3, 16) = 1.064, p = 0.39] (Figure 14.8). A post-hoc two-tailed t-test showed that the time the robots spent cleaning was significantly higher during the run in which participants had access to the app $(M_1 = 406, SD_1 = 131)$ compared to the run in which it was not provided $(M_2 = 307, SD_2 = 98)$; t(19) = 3.49, p = 0.002. SD pooled is $SD_p = 91.7$ and Effect Size was $(M_1 - M_2)/SD_p = 1.08$. In other words, participants were able to keep the

robots cleaning for more of the session when using the app. While this was especially true for participants who used the app during the second run, it was also generally true for participants who used the app in the first run.

14.5.2 Game-Playing Zone

Participants spent less time outside the game-playing zone, and thus less time watching and physically interacting with the robots, when they used the app. The app was a significant main effect on the amount of time participants spent outside the game-playing zone [F(1,16) = 4.589, p = 0.048] (Figure 14.9). A post-hoc two-tailed t-test showed that the time participants spent outside the game-playing zone observing and interacting with robots was significantly less during the run in which they had access to the app $(M_1 = 151, SD = 53)$ compared to the run in which it was not provided $(M_2 = 176, SD = 44)$; t(19) = -2.23, p = 0.038, SD pooled is $SD_p = 37.3$ and Effect Size was $(M_1 - M_2)/SD_p = -0.670$.

Participants context switched between tasks less often when provided the app, presumably because it allowed them to determine if the robots needed attention without exiting the game-playing zone. The app was a significant main effect on the number of times participants switched between being inside and outside the game-playing zone [F(1, 16) = 4.47, p = 0.05] (Figure 14.10). A post-hoc two-tailed t-test showed that the number of switches was significantly less during the run when they had access to the app $(M_1 = 4.1, SD = 1.92)$ compared to the default $(M_2 = 5.1, SD = 2.57)$; t(19) = -2.078, p = 0.05, SD pooled is $SD_p = 1.75$ and Effect Size was $(M_1 - M_2)/SD_p = -0.573$.

14.5.3 Balloons Game

On average, participants had lower scores on the balloons game when they were also using the app and presence of the app had a significant main effect on the balloons game score [F(1, 16) = 7.567, p = 0.01]. A post-hoc two-tailed t-test showed that participants scored significantly fewer points during the run in which they had access to the app $(M_1 = 382, SD = 137)$ compared to the default $(M_2 = 463, SD = 131); t(19) = -2.66, p = 0.01$, SD pooled is $SD_p = 96.0$ and Effect Size was $(M_1 - M_2)/SD_p = 2.62$ (Figure 14.11).

The lower scores are probably due to a decrease in time spent playing the balloons game, since the rate at which players scored points was similar between conditions, and there was no discernible difference in the number of penalties incurred. This result aligns with the significant main effect for the app on the amount of time spent playing the balloons game [F(1, 16) = 4.52, p = 0.05]. A post-hoc two-tailed t-test showed that participants spent significantly less time playing the balloons game during the run in which they had access to the app $(M_1 = 180, SD = 33)$, compared to the default $(M_2 = 206, SD = 56)$; t(19) = -2.21, p = 0.04, SD pooled is $SD_p = 23.6$ and Effect Size was $(M_1 - M_2)/SD_p = -1.10$ (Figure 14.12).

14.5.4 Robot Interactions

There was a weak main effect of having access to the app on the number of times participants pressed buttons on the robots [F(1,16) = 3.70, p = 0.07]. There was no main effect of the experiment condition [F(3,16) = 1.79, p = 0.18] and there was no significant interaction between the app and the experiment condition [F(3,16) = 1.04, p = 0.4] on the number of times participants pressed buttons on the robots. A post-hoc two-tailed t-test showed (weak significance) that the number of times participants pressed buttons on the robots was fewer during the run in which they had access to the app $(M_1 = 10.85, SD = 9.6)$ compared to the run in which it was disabled $(M_2 = 15.7, SD = 6.56); t(19) = -1.92, p = 0.07, SD$ pooled is $SD_p = 6.94$ and Effect Size was $(M_1 - M_2)/SD_p = -0.699$ (see Figure 14.14). Simply put, participants spent less time using the robots' physical interfaces when they also had access to the app. Two potential explanations for this are that participants were using the app controls instead of the physical controls, and that they potentially had a better understanding of why a robot might not be responding to their actions.

As implied above, app availability led to significant differences in participants' physical interactions with robots (Figure 14.15). Time with robots was calculated by coding each interaction's start and stop times for every robot up until the time the robot that needed help was reset. Timing started whenever the participant knelt or leaned over a robot while reaching towards, touching, or looking at the robot or the phone. Timing was stopped whenever the user stood up, moved their hand away from the robot, or looked away from the robot or the phone. Time spent working with the dustbins was excluded, with time stopping when the dustbin was removed and re-starting once it was replaced. Inter-rater reliability of two coders (one experimenter and one researcher not involved with data collection, both included on the IRB protocol) was computed using Cohen's Kappa and showed significant agreement ($\kappa = 0.87, \alpha = 0.05$).

A pairwise comparison using a two-tailed paired t-test on the time spent with robots during the default run without the app showed a significant difference (p = 0.03), SD pooled is $SD_p = 19.1$ and Effect Size was $(M_1 - M_2)/SD_p =$ 0.691 between the time spent with the robot that needed to be reset ($M_1 =$ 28, SD = 27) and with the robot that was playing dead ($M_2 = 14.8, SD = 19$). The difference between the robot which needed its dustbin emptied $(M_3 =$ 24.3, SD = 20) and the robot that needed to be reset was not significant (p =0.6), SD pooled is $SD_p = 14.4$ and Effect Size was $(M_3 - M_1)/SD_p = -0.257$, nor was there a significant difference between the dustbin robot and the robot playing dead (p = 0.1), SD pooled is $SD_p = 14.3$ and Effect Size was $(M_3 - 1)$ $M_2)/SD_p = 0.665$. In comparison, a pairwise comparison using a two-tailed paired t-test of the time spent with robots during the run with the app showed a larger significant difference (p = 0.008), SD pooled is $SD_p = 15.5$ and Effect Size was $(M_4 - M_5)/SD_p = 1.09$ between the time spent with the robot which needed to be reset $(M_4 = 21.5, SD = 22)$ and the robot that was playing dead $(M_5 = 4.6, SD = 10.5)$. There was also a significant difference between the robot which needed to be reset and the robot that needed its dustbin emptied

14.5. RESULTS

 $(M_6 = 6, SD = 13.5, p = 0.02)$, SD pooled is $SD_p = 15.6$ and Effect Size was $(M_4 - M_6)/SD_p = 0.996$, but not between the robot which needed its dustbin emptied and the robot that needed to be reset (p = 0.5), SD pooled is $SD_p = 10.0$ and Effect Size was $(M_6 - M_4)/SD_p = -1.55$. This suggests the app helped participants more efficiently understand which robot they needed to reset.

Participants are required to hold down the clean button for 7 seconds in order to reset a robot. This would enable a robot that needed to be reset to continue cleaning. 7 seconds was chosen to reduce the chance of a user from accidentally resetting the robot. No participant who started without the app held down the clean button to reset the robot that needed help in the first run for longer than 5 seconds (Figure 14.16). Participants with the app in their first run appeared to apply their experiences to their second run without the app, as shown by the steadily increasing number of times participants held down the button for over 5 seconds in the second run. This same group may have also either been uncertain about which robot they needed to reset, or guessed the same technique would work on multiple robots since they also appear to have tried to reset the "dustbin" robot more than in any other situation. Everyone who had the app in their second run successfully reset the robot that needed help, and very few attempts were made at trying to reset a robot that had not requested help.

Half of the participants (10/20) viewed all three robots' status pages using the app. The robot that needed to be reset used a push-style interaction which caused a popup message to appear on the smartphone screen, interrupting what the user was doing. As a result, each of the participants ended up viewing this robot's page at least once. The robot that needed its dustbin to be emptied was viewed by 12 of the 20 participants, and the robot that played dead was viewed by 10 of the 20 participants.

Half of the participants were able to use the app to view information about 2 or more robots. Using the data from the 10 participants who viewed all three robots' status pages, a pairwise comparison using a two-tailed paired t-test showed significant differences on the time participants spent on those app pages between the robot that needed to be reset $(M_1 = 88.7, SD = 43.7)$, needed to have its dustbin emptied $(M_2 = 45.3, SD = 17.1)$, or was non-functioning $(M_3 = 9, SD = 9)$ (Figure 14.17). Participants spent significantly more time looking at the page of the robot that needed to be reset than the robot that needed its dustbin emptied (p = 0.02), SD pooled is $SD_p = 30.2$ and Effect Size was $(M_1 - M_2)/SD_p = 1.43$ or the robot that had been disabled (p < 0.001), SD pooled is $SD_p = 30.1$ and Effect Size was $(M_1 - M_3)/SD_p = 2.64$. They also spent significantly more time looking at the page of the robot whose dustbin needed to be emptied than that of the robot which had been disabled (p <0.001), SD pooled is $SD_p = 11.9$ and Effect Size was $(M_2 - M_3)/SD_p = 3.04$. In other words, the amount of time participants spent viewing information about the different robots was associated with the appropriate amount of attention needed to get and keep each robot working.

14.5.5 Post-Run Questionnaires

After each run, participants were asked to fill out a questionnaire asking about their experience with the robots, their perception of what help (if any) they needed to provide to robots, and their workload.

In each of the two runs, the three robots each engaged in one of three distinct behaviors. One robot immediately required being "reset" before it could begin cleaning, but once this had been completed could continue running without needing any other help. A second robot was immediately available to begin cleaning upon request, but thereafter needed to periodically have its dustbin emptied before it could continue cleaning. The last robot played dead, and simply refused to work for the entire duration of the run. During both runs, the two robots that needed help would emit visual (flashing lights) and auditory indicators (beeping sounds) to signal there was a problem until the issue was resolved. Following each run, participants were asked "What kind of help did the robot(s) require or request? Select all that apply.". With a single exception, participants' responses were limited to the two actions which they actually needed to take. More participants correctly identified the two solicited actions during the second run (24) than during the first run (19).

The number of participants who understood a robot needed to be reset was significantly higher (p = 0.04 using McNemar's test) during runs when the app was present (12/20) compared to the default (5/20). The run number did not significantly effect participants' understanding of whether or not a robot needed to be reset (p = 0.5). However, the order in which the phone was used did seem to have an effect; 5/10 participants who had access to the app during the first run understood they needed to reset one of the robots, compared to 7/10 who had the app during the second run. In comparison, without the app, only 2/10 participants during the first run and 3/10 from the second run understood that one of the robots needed to be reset.

Participants generally felt that the robots were predictable (Figure 14.18a) regardless of whether they had access to the app $(M_1 = 3.45, SD = 1.19)$ or not $(M_2 = 3.35, SD = 1.46)$ (no significant difference was found using a two-tailed paired t-test; t(19) = -0.31, p = 0.7), SD pooled is $SD_p = 1.32$ and Effect Size was $(M_1 - M_2)/SD_p = 0.076$.

Participants were more confident that they understood what robots were doing while using the app. A two-tailed paired t-test showed that participants reported significantly more confidence in their understanding of the robots' behaviors during runs in which they had the app $(M_1 = 4.65, SD = 1.57)$ compared to runs in which they did not $(M_2 = 3.7, SD = 1.42)$; t(19) = 3.13, p = 0.005, SD pooled is $SD_p = 1.38$ and Effect Size was $(M_1 - M_2)/SD_p = 0.691$ (Figure 14.18b).

Participants' satisfaction with the robots was higher (weak significance) during the runs when the app was enabled ($M_1 = 4.6, SD = 1.19$) than during runs in which it was not ($M_2 = 4.05, SD = 1.23$) according to a two-tailed paired t-test; t(19) = 1.93, p = 0.07, SD pooled is $SD_p = 1.14$ and Effect Size was ($M_1 - M_2$)/ $SD_p = 0.484$ (Figure 14.18c).

14.5. RESULTS

A two-tailed paired t-test also showed that participants found it significantly easier to determine what was needed to keep each robot working during when they had the app $(M_1 = 5, SD = 1.26)$, compared to the default $(M_2 = 3.7, SD = 1.75)$; t(19) = 3.21, p = 0.004, SD pooled is $SD_p = 1.28$ and Effect Size was $(M_1 - M_2)/SD_p = 1.02$ (Figure 14.18d).

As part of each post-run questionnaire, participants completed the NASA TLX questionnaire. Six 2x4 mixed-groups factorial ANOVA were performed to examine the effects of using the app and the experiment conditions on mental demand, physical demand, perceived performance, success, how rushed and how discouraged they felt. This analysis was done using the raw TLX data.

There was a weak significant main effect of having access to the app on mental demand [F(1, 16) = 4.26, p = 0.055]. Also, there was a significant main effect of the experiment condition on mental demand [F(3, 16) = 3.3, p = 0.04], and a weakly significant interaction between the app and the experiment condition on mental demand [F(3, 16) = 3.02, p = 0.06]. A post-hoc two-tailed paired t-test showed (with weak significance) that participants tended to have a lower mental demand during the run in which they had the app $(M_1 = 3.7, SD = 1.13)$ than the run without it $(M_2 = 4.1, SD = 1.33); t(19) = -1.8, p = 0.088, SD$ pooled is $SD_p = 1.13$ and Effect Size was $(M_1 - M_2)/SD_p = -0.355$. A pairwise twotailed t-test with Holm correction that was used to compare differences in mental demand between experiment conditions found significant differences between condition 1 $(M_1 = 4.8, SD = 1.14)$ and conditions 3 $(M_2 = 3.4, SD = 0.97, p =$ 0.048), SD pooled is $SD_p = 1.05$ and Effect Size was $(M_1 - M_2)/SD_p = 1.33$ and 4 (M = 3.1, SD = 0.074, p = 0.045), SD pooled is $SD_p = 0.818$ and Effect Size was $(M_1 - M_3)/SD_p = 2.078$, but not between any of the other experiment conditions.

These findings suggest that using the app lowered mental demand, but that perception of workload was influenced by prior experience. We found that participants who were placed in Conditions 3 and 4 had lower mental demands than participants in Condition 1. This result is interesting because participants used the app during their second run for both Conditions 3 and 4, while participants in Condition 1 used the app in the first run (the same was true of Condition 2, but the difference was not as pronounced as in Condition 1.)

There was a significant main effect of having access to the app on the physical demand [F(1, 16) = 4.65, p = 0.046]; however, there was no main effect for experiment condition [F(3, 16) = 1.80, p = 0.18] on physical demand, nor was there a significant interaction between the variables [F(3, 16) = 1.06, p = 0.39]. A post-hoc two-tailed paired t-test showed that participants tended to have a lower physical demand during with the app $(M_1 = 2.15, SD = 0.81)$ compared to the default $(M_2 = 2.7, SD = 1.34)$; t(19) = -2.15, p = 0.04, SD pooled is $SD_p = 0.986$ and Effect Size was $(M_1 - M_2)/SD_p = -0.558$. The app allowed participants to substitute walking between different locations in the room to accessing robots with an interaction in the app.

There was a significant main effect of having access to the app on how discouraged participants reported feeling [F(1, 16) = 5.04, p = 0.04]; however there was no main effect of the experiment condition [F(3, 16) = 1.54, p =

0.24], nor was there a significant interaction between the variables [F(3, 16) = 0.34, p = 0.79]. A post-hoc two-tailed paired t-test showed that participants reported feeling less discouraged during when they had the app $(M_1 = 3.4, SD = 1.67)$ than without it $(M_2 = 4.05, SD = 1.62)$; t(19) = -2.37, p = 0.03, SD pooled is $SD_p = 1.47$ and Effect Size was $(M_1 - M_2)/SD_p = -0.444$. This result is consistent with our findings that participants were more confident about understanding the robots' actions and found it easier to determine what to do to make robots work while using the app.

14.5.6 Post-Experiment Questionnaires

After the second run, participants were asked to complete a final questionnaire. Nearly all of the participants (18/20) reported that learning how to use the app was easy (Figure 14.19a). Most also reported that it was easier to figure out what they needed to do with the robots (14/20) and to control them (15/20) by using the app rather than by physically looking at the robots themselves (Figures 14.19b and 14.19c). The majority of participants (14/20) felt that they understood which robot the information in the app was referring to more than half the time, while only 3/20 felt that they "rarely" or "sometimes" understood. A two-tailed t-test did not report significant difference (t(14.6) = -0.5, p = 0.6) between ordering conditions.

All 20 participants said they preferred having access to the app (Figure 14.20a). Nineteen out of the 20 participants thought the information they received from the app was helpful and informative, and 16/20 felt more in control of the robots when they had access to the app (Figures 14.20b and 14.20c). There were no significant differences between run ordering conditions (p = 0.3, p = 0.3, and p = 0.4, respectively, using two-tailed t-tests).

14.5.7 Experiment Scores

There was no main effect for the app on experiment score [F(1, 16) = 1.16, p = 0.3]. However, there was a significant main effect for experiment condition on experiment score [F(3, 16) = 3.18, p = 0.05], and there was a weak interaction between the app and experiment [F(3, 16) = 3.06, p = 0.06]. A post-hoc pairwise comparison using paired t-tests showed significant differences between Condition 2 $(M_1 = 196, SD = 75)$ and Conditions 1 $(M_2 = 184, SD = 92, p = 0.03)$, SD pooled is $SD_p = 52.8$ and Effect Size was $(M_1 - M_2)/SD_p = 0.227$ and 3 $(M_3 = 183, SD = 78, p = 0.04)$, SD pooled is $SD_p = 52.7$ and Effect Size was $(M_1 - M_2)/SD_p = 0.247$.

We have been unable to produce a suitable explanation for why participants in Condition 2 did not perform as well as those in the other three conditions. The low experiment scores were a combination of both low balloon game scores (although the differences were not significant) and low combined time robots spent cleaning. The later was the result of just a single person (out of five) in Condition 2 successfully using more than one robot, compared to 4 out of 5 in Condition 1, 5 out of 5 in Condition 3, and 4 out of 5 in Condition 4. There were no significant differences between Condition 2 and the other conditions with respect to time spent outside the game-playing zone, app usage, time spent playing the balloons game, penalties incurred in the balloons game, or time spent physically interacting with the robots. That said, despite the lack of difference in experiment scores based on the presence of the app, our results still largely supported our hypotheses.

14.6 Discussion

14.6.1 H1 is Supported by the Results

A majority of participants indicated they could tell which robot was the source of information in the app (70%) and that it was easier to figure out what they needed to do to make the robots work (70%) with the app than by looking at the robots. Six of the ten participants who had the app during their first run were able to "reset" the robot that needed help, and five out of ten were able to transfer that knowledge by successfully identifying and resetting a different robot that needed help in their second run without using the app.

In contrast, no one who did not have the app in the first run was able to reset the robot that needed help. However, all of those same participants (10/10) were able to successfully reset another robot which needed help when provided the app in their second run.

14.6.2 H2 is Partially Supported by the Results

Participants were able to use the app to retrieve information about the robots and to identify solutions to problems. All of the participants used the app to view the robot that employed a push-style interaction (popup message), and half of the participants viewed all three robots using the app. There are a few possible reasons why more of participants did not view all three robots.

First, while basic use of the Android phone was demonstrated, participants were not provided with any training on how to use the robot smartphone app or its capabilities. The popup message appeared shortly after the run began, as one of the first experiences with the app. Therefore, participants may have believed that any future information would also come in the from of popup messages.

Another related possibility is that some participants may have simply not been interested in viewing the information of some robots, as with two participants who only viewed two of the three robots (both did not view the "dead" robot).

The third possibility is that some participants may not have been able to figure out how to access the information about the other robots. Of the eight participants who only viewed one robot, five of them reported having no prior experience using Android devices. Participants were able to use the app to identify solutions to problems and allocate their time more appropriately. Participants were much more likely to get two robots working when equipped with the app, with a combined time for robot cleaning that was over a minute and a half longer (on average) than without the app. During app runs, participants tended to focus most of their time and attention on the robot which needed to be reset first.

Participants using the app were also able to get the robots to clean for longer while also spending less time watching and physically interacting with them. We had predicted that this kind of behavior would lead to better performance since participants would have more time to score points in the balloons game and retain a higher percentage of their score. However, despite gaining an average of 30 additional seconds to play the balloons game when the app was enabled, participants using the app actually spent significantly less time playing the balloons game, causing their overall performance to be about the same. Much of the lost time was spent using the app. One potential explanation for this behavior is that the app's ability to communicate with and control the robots had a strong novelty effect on participants, leading them to spend more time with it than with the balloons game. This explanation is supported by the fact that only 3 participants had previously used a robot vacuum cleaner.

14.6.3 H3 is Supported by the Results

Participants liked having access to the app, the additional information it provided, and its controls over the robots' built-in interfaces. All of the participants reported that they liked being able to communicate with the robots through the app, and all but one (who was neutral) thought the app was helpful. The majority of participants said they felt more in control using the app. The majority also felt much more confident about understanding what the robots were doing and what needed to be done to make the robots work. According to the NASA TLX questionnaires, participants felt less discouraged while using the app. Unsurprisingly, participants also reported higher levels of satisfaction with using the robots during the run with the app. According to workload data from the NASA TLX questionnaires, the app reduced participants' mental and physical demands.

14.6.4 Effects of Run Ordering

Some of our results show evidence of an ordering effect. For example, participants who used the app in the first run were able to apply knowledge from their first experience during their second run (e.g., resetting the robot which needed help). Additionally, the 100% increase in participants who could reset the robot in the second run with the app who had previously been unable to without the app suggests that they had a better understanding of the information the app was providing due to their prior experience.

NASA TLX responses from after each run also show ordering effects. Some participants found the app to be distracting (4/20) or its messages confusing

(7/20), with most of the complaints coming from participants who used the app during the first run. Participants who used the app during their second run had the additional context on the task. The results from a similar question asked at the end of the experiment support this theory; in that question, all but one participant said they thought the information from the app was helpful.

14.7 Future Work

Perhaps the most important limitation of this study is that the participants were not representative of bystanders, who we consider to be an important target audience for this work, ultimately. Instead, participants were acting as operators or supervisors, with the robots' goals being aligned with their own goals. That said, given the lack of training provided, participants did have several characteristics of typical bystanders: a lack of familiarity or prior experience with the robot platform and a lack of training with its user interfaces. In order to measure how the feedback could assist people who are true by tanders to robots, we would test situations in which people must interact with a robot without previously having been informed that a robot would even be involved. but such experiments are difficult to construct. An alternative, less deception oriented experiment might ask people to find and help an autonomous robot carry out a task it has been assigned without prior knowledge of exactly where the robot is, what it looks like, what it is doing, or how to communicate with the robot. We believe that the results of our study are promising for these future bystander experiments, as we have shown that the app is able to convey status information about the context and methods for assisting the robots, an example of shared context.

Another limitation of this work is the lack of explicit comparisons of the effectiveness of *push* vs *pull* interactions. Both interaction styles were used during the experiment; however, they were directly paired with a single type of problem, and always occurred in the same order with participants receiving a *push* interaction popup message shortly after the beginning of the run. Testing the difference in effectiveness of these interaction styles is warranted. Finally, one of the most powerful applications of this work is its potential to be relevant for communicating with a wide variety of different kinds of robotic platforms. Further testing with different kinds of robots, including drones and self-driving cars, is necessary to determine if this technology would be suitable as a ubiquitous method for communicating with publicly deployed autonomous robots.

14.8 Conclusions

The results of this work support the use of smartphones as a ubiquitous interaction method to allow untrained users to communicate with autonomous robot systems. Participants were able to use the app without training, and reported that it was easy to learn and use. All of the participants preferred having access to the app, and all but one said the app was was helpful. Participants were able to retrieve information about nearby robots, and could distinguish the source of the information despite similarities in the robots' appearances. With the app, participants felt more confident they understood what the robots were doing and were more satisfied with the robots' performance. Although participants' experiment scores did not improve with the use of the app, their behavior (specifically, spending less time watching and physically interacting with the robots, and getting more robots working for longer periods of time) created the potential for improved performance. While additional experimentation is needed to better understand the differences between the *push* and *pull* interaction methods and how bystanders might use the system, these results are a promising first step towards building communication between people and the increasing number of autonomous robots in our society.

14.9 Acknowledgements

The research in this chapter was supported in part by the National Science Foundation under awards IIS-1552228 and IIS-1552256. Thanks to Chuta Sano and Christopher Munroe for their assistance with aspects of the experiment. Daniel Brooks was a doctoral candidate at the University of Massachusetts Lowell when this work was conducted.



Figure 14.7: Time robots spent cleaning without (left) and with (right) the app. Robots spent significantly more time (p = 0.002) cleaning when participants had access to the app.



Figure 14.8: Time robots spent cleaning by the app condition and then the run number in which participants experienced that condition (first or second), demonstrating that there is some learning effect between the first and second runs, regardless of whether the participant started with the app or without it.



Figure 14.9: Time spent outside of the game-playing zone without and with the robot status app.



Figure 14.10: Number of switches in and out of the game-playing zone without and with the robot status app.



Figure 14.11: Score in the balloons game, without and with the robot status app.



Figure 14.12: Balloons game playing time, without and with the robot status app.



Figure 14.13: Penalties incurred in the balloons game, without and with the robot status app.



Figure 14.14: Number of button presses on the robots, without and with the robot status app



Figure 14.15: Time physically spent with robots until robot reset, without and with the robot status app.



Figure 14.16: Clean button click/hold times (in seconds). The two left bar graphs show one set of participants, who experienced the no app condition first, followed by the app condition second. The two right bar graphs show the other set of participants, who experienced the app condition first, followed by the no app condition second.



Figure 14.17: Time spent in the robot status app



36

The robots actions were surprising or unpredictable.

(a) Robot predictablity







My confidence level with regard to understanding what each robot was doing was _____

(b) Confidence





Figure 14.19: Robot smartphone app characteristics



Figure 14.20: Robot smartphone app preferences

Bibliography

- D. J. Barber, I. Julian Abich, E. Phillips, A. B. Talone, F. Jentsch, and S. G. Hill. Field assessment of multimodal communication for dismounted human-robot teams. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 59(1):921–925, 2015.
- [2] D. J. Brooks, A. Shultz, M. Desai, P. Kovac, and H. A. Yanco. Towards state summarization for autonomous robots. In AAAI Fall Symposium Series, 2010.
- [3] J. Carlson and R. R. Murphy. How UGVs physically fail in the field. *IEEE Transactions on Robotics*, 21(3), 2005.
- [4] E. Cha, M. Mataric, and T. Fong. Nonverbal signaling for non-humanoid robots during human-robot collaboration. In *Proceedings of the 11th* ACM/IEEE International Conference on Human-Robot Interaction, pages 601–602, March 2016.
- [5] M. Desai, P. Kaniarasu, M. Medvedev, A. Steinfeld, and H. Yanco. Impact of robot failures and feedback on real-time trust. In *Proceedings of the 8th* ACM/IEEE International Conference on Human-Robot Interaction, Tokyo, Japan, 2013.
- [6] V. Groom, J. Chen, T. Johnson, F. A. Kara, and C. Nass. Critic, compatriot, or chump?: Responses to robot blame attribution. In *Proceedings of* the 5th ACM/IEEE International Conference on Human-robot Interaction, 2010.
- [7] S. G. Hart and L. E. Staveland. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. Advances in Psychology, 52:139–183, 1988.
- [8] Y. Hiroi and A. Ito. ASAHI: OK for failure: A robot for supporting daily life, equipped with a robot avatar. In *Proceedings of the 8th ACM/IEEE In*ternational Conference on Human-Robot Interaction, pages 141–142, 2013.
- [9] H. Hüttenrauch and K. Severinson Eklundh. To help or not to help a service robot: Bystander intervention as a resource in human-robot collaboration. *Interaction Studies*, 7(3):455–477, 2006.

- [10] iRobot. Features of the iRobot HOME App, 2017 (accessed December 27, 2017).
- [11] iRobot. iRobot Robot Vacuums, 2017 (accessed December 27, 2017).
- [12] iRobot. Resetting the iRobot HOME App "Care" status, 2017 (accessed December 27, 2017).
- [13] D. B. Kaber and M. R. Endsley. Out-of-the-loop performance problems and the use of intermediate levels of automation for improved control system functioning and safety. *Process Safety Progress*, 16(3):126–131, 1997.
- [14] P. Kaniarasu and A. Steinfeld. Effects of blame on trust in human robot interaction. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, Aug 2014.
- [15] A. B. Karami, K. Sehaba, and B. Encelle. Adaptive artificial companions learning from users' feedback. *Adaptive Behavior*, 24(2):69–86, 2016.
- [16] T. Kim and P. Hinds. Who should I blame? Effects of autonomy and transparency on attributions in human-robot interaction. In *Proceedings of* the 15th IEEE International Symposium on Robot and Human Interactive Communication, 2006.
- [17] R. A. Knepper, S. Tellex, A. Li, N. Roy, and D. Rus. Recovering from failure by asking for help. Autonomous Robots, 39(3):347–362, 2015.
- [18] M. K. Lee, S. Kiesler, J. Forlizzi, S. Srinivasa, and P. Rybski. Gracefully mitigating breakdowns in robotic services. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction*, March 2010.
- [19] B. Lussier, R. Chatila, F. Ingrand, M.-O. Killijian, and D. Powell. On fault tolerance and robustness in autonomous systems. In 3rd IARP-IEEE and RAS-EURON Joint Workshop on Technical Challenges for Dependable Robots in Human Environments, pages 351–358, Sept 2004.
- [20] Neato Robotics. Neato Botvac Connected, 2015 (accessed December 27, 2017).
- [21] Neato Robotics. Neato Robot Vacuums, 2017 (accessed December 27, 2017).
- [22] S. Rosenthal, M. Veloso, and A. K. Dey. Is someone in this office available to help me? *Journal of Intelligent & Robotic Systems*, 66(1):205–221, 2012.
- [23] V. Srinivasan and L. Takayama. Help me please: Robot politeness strategies for soliciting help from humans. In *Proceedings of the 2016 CHI conference* on human factors in computing systems, pages 4945–4955. ACM, 2016.

- [24] G. Steinbauer. A survey about faults of robots used in RoboCup. In RoboCup 2012: Robot Soccer World Cup XVI, pages 344–355. Springer, 2013.
- [25] D. Szafir, B. Mutlu, and T. Fong. Communicating directionality in flying robots. In Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, pages 19–26. ACM, 2015.
- [26] H. Yasuda and M. Matsumoto. Psychological impact on human when a robot makes mistakes. In *IEEE/SICE International Symposium on System Integration (SII)*, pages 335–339, Dec 2013.