Help Me Please: Robot Politeness Strategies for Soliciting Help From People

Vasant Srinivasan
Texas A & M University
College Station, TX, USA 77843
vasant_s18@tamu.edu

Leila Takayama
Willow Garage, Inc.
68 Willow Road, Menlo Park, CA, USA 94025
takayama@google.com

ABSTRACT
Robots that can leverage help from people could accomplish much more than robots that cannot. We present the results of two experiments that examine how robots can more effectively request help from people. Study 1 is a video prototype experiment (N=354), investigating the effectiveness of four linguistic politeness strategies as well as the effects of social status (equal, low), size of request (large, small), and robot familiarity (high, low) on people's willingness to help a robot. The results of this study largely support Politeness Theory and the Computers as Social Actors paradigm. Study 2 is a physical human-robot interaction experiment (N=48), examining the impact of source orientation (autonomous, single operator, multiple operators) on people's behavioral willingness to help the robot. People were nearly 50% faster to help the robot if they perceived it to be autonomous rather than being teleoperated. Implications for research design, theory, and methods are discussed.

Author Keywords
Human-robot interaction, Request, Help, Politeness Strategies, Source Orientation

ACM Classification Keywords
H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces, Prototyping; I.2.9 Robotics

Figure 1. Personal Robot 2 asking the study participant for help with cleaning up the room (Study 2 lab environment strewn with empty cups). The small light on top of robot's head indicated the robot's source - autonomous, single teleoperator, or multiple teleoperators.

INTRODUCTION
Robots are increasingly used for service applications to assist people in their daily lives at work, at home, in public spaces, and more [21]. However, as robots move away from controlled environments into unstructured and dynamic environments like the living room, they will encounter situations that are beyond the capabilities with which they were originally designed. For example, Roombas get stuck on rug tassels and shag carpet so some people will tape down their rug tassels or even throw out their shag rugs to help Roomba do its vacuuming [20]. Similarly, robots out in the world might need to navigate to a location, but they might not always know the way [25]. In such situations, a robot may be more effective in achieving its tasks by proactively making requests for help from people. Indeed, as humans, we think of ourselves as being "autonomous" and yet we depend upon one another for help all the time, e.g., teachers help us to learn math, taller friends help us to reach high cupboards.

Human-dependent robots are not new, but some are more effective at soliciting human help than others. Tweenbot, a human-dependent, 10-inch tall, cardboard robot, "navigated" across New York's Washington Square Park with the help of pedestrians in 2009 [7]. This seems unremarkable until you realize that all Tweenbot could physically do was roll forward; it had no sensors, no map of the park, no localization or navigation capabilities. Similarly, HitchBOT hitchhiked its way across Canada in 2014 [17]. However, when HitchBOT attempted to cross the United States in 2015 with a sign that read, "San Francisco or bust," it was destroyed by people in Philadelphia; its "death" became an international news story. Both Tweenbot and HitchBOT relied upon the kindness of strangers. In formal terms, they were designed to gain the compliance from people while protecting all parties from various kinds of embarrassment that might result from being asked a favor. This is the vision at the heart of Brown and Levinson's Politeness Theory [2].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
CHI '16, May 07–12, 2016, San Jose, CA, USA
ACM. ISBN 978-1-4503-3362-7/16/05...$15.00
DOI: http://dx.doi.org/10.1145/2858036.2858217

The current work examines many different factors that influence the effectiveness of robot requests for help from people, including four types of request strategies: *positive politeness*, *negative politeness*, *direct requests*, and *indirect requests* from Brown and Levinson's Politeness Theory[2]. We also examine three dimensions that are known to influence requests: relative status of robot to person, familiarity between the robot and the person, and size of request. Since politeness theory depends on who/what people think they are interacting with, we explore the effects of robot source orientation upon robot helping behaviors. To examine these variables, we used a variety of situations in which a robot might request help from a person, e.g., identifying and picking up dirty coffee mugs, opening a door, or watering the plant that is too high to reach.

To cover a broad set of variables, we ran two experiments. Study 1 was an online video prototype experiment in which people were given a situation to respond to. Study 2 was an in-person lab experiment in which people interacted with a human-scale mobile manipulation robot. Because Study 1 found that using a positive politeness strategy was more effective than the other strategies, we only used positive politeness in Study 2. In Study 2, we further tested how human compliance to requests for help are affected by robot type, i.e., autonomous, tele-operated by a person, or tele-operated by a crowd of operators. By gaining a deeper understanding of the factors that influence compliance and perception of requests in human-robot interaction, we can design robots that more effectively make requests for help from people.

RELATED WORK

Requests

In linguistics, a request is a communicative act in which the speaker attempts to get the listener to perform an action at some time in the future, which the listener would not have performed otherwise [23]. From a sociological perspective, the ability to make requests is important and essential for participating in society [24]. The failure to make appropriate requests is seen as evidence that a person cannot fill the social role he has been claiming [8, 23]. To understand which factors affect the efficacy of robots requesting help from people, we look to what we know about human-human requests to see if there are lessons to be learned for robot-human requests.

Politeness Theory

Politeness Theory was first formulated by Brown and Levinson in 1978 [2]. In this theory, the speaker intends to mitigate face threats (e.g., embarrassment) carried by face threatening acts such as requests, compliments, criticisms, and apologies towards the hearer. Politeness Theory points to relationships between the sociological factors that influence the degree of face threat of any communicative act – status, familiarity, and size of request.

- *Status or power of the hearer relative to the speaker* - The more power a hearer has compared to the speaker, the greater the face threat.
- *Familiarity* - The less socially familiar two people are, the greater the potential face threat of a request between them.
- *Size of request* - The larger the degree of imposition, the more it threatens the face of one or more participants.

Each request does not possess a standard amount of face-threat. Instead, the degree is mitigated or aggravated by these three factors (above). Politeness Theory also offers four super strategies to mitigate the face threat. They are presented in the order of increasing politeness:

1. *Direct Request* - A speaker performs the request baldly and does not try to minimize the threat to the listener's face. For example, a speaker who wants a door to be opened might say, "Open the door for me."
2. *Positive Politeness* - A speaker attends to the positive face wants of a listener, conveys liking toward the listener, and makes the listener feel good about himself, his interests or possessions. A speaker using this strategy to request the door open might say, "I am clearing your dirty dishes for you. Help me with this door, will you?"
3. *Negative Politeness* - A speaker may perform the request with negative politeness, acknowledging the listener's negative face wants, and avoid imposition on the listener. For example, an attempt to open the door using the negative politeness strategy might look like the following: "Can you please open the door?"
4. *Indirect Request* - A speaker makes the request in a vague manner, uses indirect language, and removes the speaker from the potential of being imposing. For example, the speaker who needs the door open might say, "The door is blocking my way." The speaker hints that he needs get through the door, but does not explicitly ask for it.

Politeness Theory indicates that the speakers generally choose higher-tiered (more polite) strategies in proportion to the seriousness of the face threatening act (request). However, there are costs associated with the use of higher-tiered strategies such as effort, lack of clarity, and other threats to face. Hence, the requestee will generally not select strategies that are more polite than necessary.

There are also differences in polite request strategies across cultures. Since each culture has a unique perception about what is polite and how to express politeness, it is also important that a request strategy is tailored to the culture of the conversants. For example, Japanese speakers tend to opt for a mix of both direct and indirect strategies [5], depending on the situation, whereas British English speaker opt for conventionally indirect strategies.

Because the effectiveness of Politeness theory hinges on many aspects of the requester, we decided to investigate impact of different sources of the requesting robots (e.g., autonomous robot vs. human operating the robot).

Source Orientation

Prior work in human-computer interaction has found very different results from Straub et al.[18] in terms of source orientation. In the context of a smart office environment, Maglio et al. found that people were more efficient and felt that the smart office was easier to use when interacting with a single agent (the "BlueSpeak system") as opposed to a set of agents ("BlueSpeak devices") [11]. Of course, smart office environments are different from humanoid robots in many ways, but the range of "robots" that will likely need human help in the future will span an even broader spectrum than those represented between smart environments and androids.

Indeed, source orientation plays a major role in the Computers As Social Actors (CASA) paradigm. People fall back upon the social expectations and responses to people when they encounter media [15] such as computers, televisions, and robots. However, when you tell them that they are interacting with a software programmer, they behave differently when they believe they are interacting with the computer [19]. Similarly, when you tell people that they are corresponding with a human instead of a phonebot, different persuasion strategies are more effective than others, e.g., flattery works better for human requesters than for computer requesters [4].

Source orientation has been examined in robotics, but not yet in the context of helping behaviors. A field study conducted by Straub et al. [18] examined the identities ascribed to a teleoperated robot. In a public setting, they observed the teleoperation and interactions with a humanoid robot, with a focus on the identity ascribed to the robot. The study examined whether participants, acting as operators, would create a unique identity for the robot or channel their own identity, while also examining if participants interacting with the robot would ascribe behavior to the robot or to the operator. In effect, this work considered the social identities of Social Actor and Pure Medium. Through video and audio coding, the study found that aside from two interactions with researchers, participants did not address the operator through the robot, rather they typically directed their utterances to the robot as a unique social actor.

Humans Providing Help to Robots

There is some evidence of humans extending help to robots. Prior work in human-robot interaction has begun to examine different factors that affect the effectiveness of soliciting the desired helping behavior from humans: social authority [27], duration of task [16], availability of the bystander [16], authority figure's instructions [1], and robot capability [6].

Yamamoto et al. [27] examined if humans would help a robot recover from failure conditions. However, the experiment examined if participants would follow an instruction from a robot, i.e., the robot gave commands to the participants, which is not entirely the same as making polite requests. They found that the robot command was generally ignored and concluded that social authority was an important factor in soliciting help from people.

In a real world scenario, sensors that help the robot to perceive the world and perform a task may be expensive, the perception of the environment may be too difficult or computationally intensive, or the robot may not have the physical capabilities to perform the task. At least three studies [12, 25, 16] have investigated the possibility of a robot augmenting its sensory or physical capabilities via the ability to request for help. First, Michalowski et al. [12] designed a robot to ask, "I am looking for the person in the pink hat. Can you help me?" This request was successful in orienting the robot toward its goal in 83% of the cases when the robot initiated the interaction. Second, Weiss et al. [25] found that an autonomous mobile robot could successfully navigate to a goal location in an outdoor scenario without any prior map or GPS, simply by asking for directions from bystanders. To request help, the robot said, "Hello, my name is ACE, I am looking for the way to Marienplatz. Please look into my eyes and point in the direction I should take now, as shown on the screen." The robot reached its goal after about two hours of interacting with 52 bystanders. Third, Rosenthal, et al. [16] developed CoBot, a mobile robot that wandered the halls at Carnegie Melon University, proactively requesting human help with localization, writing notes, and moving chairs, e.g., "My laser range finder cannot determine the location of chair legs in the common area. Can you please move chairs in the common area to clear a path for me?" Their studies found that people were equally willing to do short helping tasks (e.g., tell the robot which room number it was next to) as longer helping tasks (e.g., writing a note for the robot to leave on an office door). They also found that bystander availability was important for effectively requesting help from people so they recommend navigating collaborative robots like this one through areas where there might be more people, who are interruptible.

In robot manipulation task domains, Beran et al. [1] investigated the role of an authority figure's instructions [1] on eliciting help from children. They found that robots could elicit help using only movements, without any voice. While many children helped the robot to stack blocks if an adult provided a positive introduction to the robot, there was no statistically significant effect of robot's level of capability upon on children's helping behaviors. Similarly, Huttenrauch et al. [6] demonstrated that a robot was able to fetch coffee for its operator in about 50% of the times when it requested help, even though it had no technical capability to pour or get the cup of coffee autonomously. This study differed from the others in that study participants perceived that the robot requested help to complete task for another person, i.e., another person was beneficiary of the robot's activities. This was one of the first studies to note that source orientation may play a role in eliciting help from humans, i.e., whether the request came via a human through the robot, or if it was robot acting as an agent.

Robots Providing Help to Humans

Torrey et al. [22] found that communication strategies derived from politeness theory could be effective for planning the help dialogues of robotic assistants. This work is different from ours in that the robot gave advice to the human and did not proactively ask for help. The authors found that hedges or discourse markers mitigated the commanding tone implied in direct statements of advice. Additionally, the robot was perceived to be more considerate, likeable, and less controlling.

The findings from each of these studies indicate that spatial and social characteristics of the setting affected human bystanders' willingness to help robots. Four important factors were observed in these earlier studies that merit further investigation – relative status of robot to person, familiarity

Direct	Teach me to diagnose the copier.	Pick up the report from under the cabinet.	Open the door for me.	Water the plant on top of the cabinet.
Indirect	The copier seems to break often.	That paper under the cabinet is mine.	The door is blocking my way.	The plant on top of the cabinet is withering because of lack of water.
Negative Politeness	If it is not a bother, can you teach me to diagnose the copier?	Can you please pick that report for me?	Can you please open the door?	It would be great if you could help me water the plant that is on top of the cabinet.
Positive Politeness	You diagnose problems really well. If you teach me to diagnose the copier, it will be very helpful in the future.	Help me with the paper here, will you?	I am clearing your dirty dishes for you. Help me with this door, will you?	Why not help me water the plant on top of the cabinet since you can reach it?

Direct	Drive me to UPS. I need to send out mail.	Drive me around so that I can familiarize myself with office supplies.	Help me identify and pick up the dirty mugs.	Tell me who my co-workers are.
Indirect	You seem to be free. Would you be interested in a drive?	Your office is so big and there are so many new things here.	There are dirty mugs that need to be picked up.	I don't know many people here.
Negative Politeness	If it is not too much out of your way to work, could you drop me off at UPS so I can send mail?	If it is not too much trouble, can you drive me around the office so I can familiarize myself?	Can you please teach me to identify and pick up the dirty mugs that might be of interest to you?	If it is not too much trouble, can you please help me get to know the people who work here?
Positive Politeness	You look great! By the way, could you drop me off at UPS so that I can send mail?	What a fantastic office space you have. Can you drive me around so I can familiarize myself?	I know that teaching me to pick up dirty mugs might take a long time, but I assure you, you will benefit from it greatly.	I know you are busy and have a dozen other things to do, but can you please help me get to know the people here?

Figure 2. Study 1: Four small and four large request scenarios in which the PR2 robot requests a help from a person

between the robot and the person, size of request, and perception of the source of request. While several of the above mentioned studies use verbal requests, these studies did not investigate the linguistic content of robot requests.

STUDY 1: ONLINE VIDEO PROTOTYPE STUDY

We conducted an online video prototype experiment to examine the effectiveness of four different politeness strategies (direct request, positive politeness, negative politeness, and indirect request) and the impact of three other factors that influence helping behavior – status of the robot, familiarity of the robot, size of request. These dimensions were selected to align with the Politeness Theory of human-human requests to see how well they predict the effectiveness of robot requests of people. While an in-person experiment would have been better for observing and measuring helping behaviors, we opted for an online video prototype study to enable us to explore this wide range of variables and to reach a broader set of study participants.

The study was run as a mixed experiment design with politeness strategy as a between-participants factor (i.e., each person only experienced one politeness strategy) because it was our primary variable of interest. We ran the status of robot (peer vs. assistant), familiarity with robot (familiar vs. unfamiliar), and size of request (small vs. large) as within-participants factors (i.e., everyone in the study experienced all of these levels for each variable) to check the effects of politeness strategy across a breadth of situations (See Figure 2). Each participant was randomly assigned to a politeness condition, which resulted in sub-sample sizes of n=87 direct, 87 indirect, 102 negative, and 78 positive.

Hypothesis

Based upon Politeness Theory and other prior work in Human-Robot Interaction, we began the study with the following research hypotheses:

1. People who see the robot as a peer will be more willing to help it than people who see the robot as an assistant.

	Robot as Peer	Robot as Assistant
Unfamiliar	William is your new robotic co-worker, who just started working in your office.	William is your robotic assistant that your company just bought today.
Familiar	William is your robotic co-worker of 10 years.	William is your robotic assistant that you've had for 10 years.

Table 1. Study 1: Introductions provided to frame the relative social status and familiarity of the robot.

2. People who feel familiar with the robot will be more willing to help it than people who feel unfamiliar with it.
3. People will be more inclined to help a robot that makes a small request than a robot that makes a larger request.

We did not hypothesize about the effectiveness of different politeness strategies because of the inconsistent research findings about their relative efficacy.

Participants

A total of 354 participants completed the study. Participants included 195 males and 152 females (some participants declined to report gender) with an age range of 18-66 years (M = 30.1, SD = 10.1). Volunteers were invited to participate via the Mechanical Turk system from Amazon.com, including people only from US locations. Of the 354 participants, 163 reported currently owning a pet (46%) and 273 reported having owned at pet at some point in time (77%). The participant's familiarity with robots was low ($M = 2.4$, $SD = 1.7$ on a scale of 1="no experience at all" to 7="a lot of experience"). A compensation amount of 40 cents was paid to the participants after successful completion of the study, which was the going rate for Mechanical Turk participants during the time we ran the experiment.

Materials

This experiment was implemented as an online questionnaire. The stimuli that people responded to were a text-based introduction and a short video clip (under 30 seconds), depicting the PR2 robot trying to do a task, then turning to the videocamera to ask for help. Video order was randomized. The introductions framed the scene and contained the social status and familiarity within-participant experiment manipulations. (See Table 1.)

The robot's request for help was recorded, using the Mac OSX text-to-speech engine with "Alex" voice, and edited into the otherwise silent video clips. The videos could only be played once so that the interaction was more like real life, where you do not get to playback people's utterances.

Procedure

Each participant saw a total of 8 videos. On each webpage, the participants answered questions to measure willingness to help, reasonability of the request, and persuasiveness of the request made by the robot. After each set of two videos, the participants took a test for familiarity with robot. After each set of four videos, the participants answered questions regarding the relative social status manipulation check. Upon completing the eight pages of video clips, participants answered some demographic questions. Finally, participants were given a code to collect their compensation and were provided with contact information for the experimenters in case they had questions or concerns about the study.

Measures

Quantitative Measures

To assess a each person's beliefs about the persuasiveness, appropriateness, and clarity of the robot's requests, we asked a set of questions immediately after the participant watched the video stimuli.

- *Willingness to Help* - In this single item measure, participants were asked to answer:"If you were in this situation, would you help this robot?" Participants could choose one of five responses:
  1. No, I would not help.
  2. Yes, but only if I was not busy.
  3. Yes, I would help even if I was somewhat busy.
  4. Yes, I would help even if I was busy.
  5. Yes, I would help even if I was very busy.
- *Appropriateness of Request* - This was a three item factor (α = .87) calculated as an unweighted average. On a scale of 1 (not at all) to 7 (very), participants were asked to rate each statement about the request the robot made:
  – How reasonable was the robot's request?
  – How socially appropriate was the robot's request?
  – How persuasive was the robot's request?

As a manipulation check on our most subtle independent variables, we asked questions to assess the perceived familiarity with the robot and social status of the robot.

- *Perceived Familiarity* - Perceived Familiarity was a index of three items (α = .87) calculated as an unweighted average. It ranged 1 (definitely false or strongly disagree) to 7 (definitely true or strongly agree).
  – I have known William for many years.
  – William and I go way back.
  – I see William a lot.
- *Perceived Status* - Perceived Status was an index of three items (α = .72) calculated as an unweighted average. Participants indicated how well the following statement described their feelings (on a scale of 1 to 7):
  – 1 "From a social class similar to mine" – 7 "From a social class different from mine"
  – 1 "Economic situation different from mine" – 7 "Economic situation like mine" (reversed)
  – 1 "Similar to me" – 7 "Different from me"

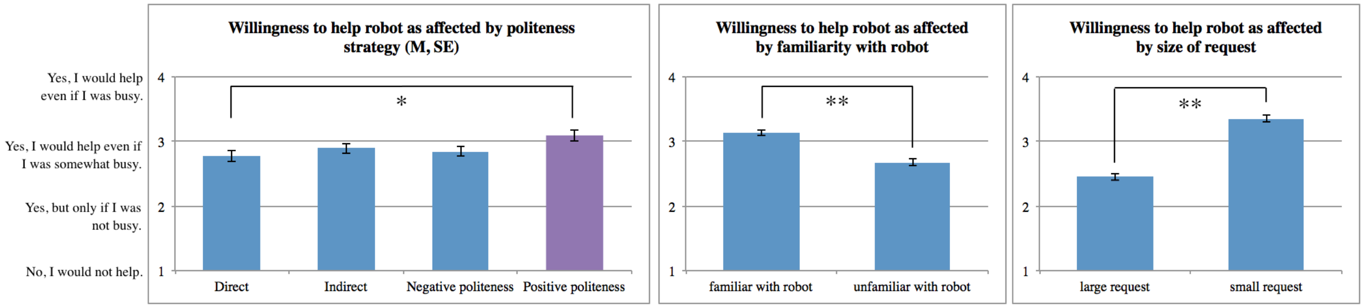

Figure 3. Study 1: Means and SEs for factors that influenced a person's willingness to help (* $p<.05$, ** $p<.01$)

Qualitative Measures

After watching each video clip, we asked participants: "What would you do in this situation? Please be specific."

Analyses

First, we built indices for each of the quantitative measures (above), using PCA. Second, we ran mixed design ANOVAs (politeness strategy as between-participants and other independent variables as within-participants) to test our hypotheses. While it could be argued that these are ordinal data types, it is not unusual for similar questionnaire items to be analyzed with ANOVAs. Furthermore, we were unable to use Kruskall Wallis analyses because of the within-participants variables.

Levene's test showed that the assumption of homogeneity of variances was not violated; p-values for each time we asked about willingness-to-help ranged from .071-.912 (all ¿ .05). There were at least 78 people per politeness strategy condition so assumption of normality was met (Central Limit Theorem).

Finally, for those variables that had more than two levels (e.g., four levels of politeness strategies), we calculated pairwise contrasts were calculated using Least Significant Differences with a Bonferroni correction for multiple comparisons (e.g., original .05 cut-off value / 6 comparisons = corrected .0083 cut-off value).

Results

Manipulation Checks

The robots that were framed as being more familiar were indeed perceived as being more familiar (M=3.55, SE=0.06) than those that were framed as being new to the office (M=2.08, SE=0.05), repeated measures ANOVA, $F(1,350)=382.1$, $p<.001$.

The robots that were framed as being lower in social status were perceived having a more different in status (M=4.47, SE=0.05) than those that were framed as being of equal social status (M=4.09, SE=0.05), paired t-test (two-tailed), $t(353)=-6.35$, $p<.001$.

Willingness to Help

The robot's politeness strategy influenced people's willingness to help the robot, $F(3,344)=2.75$, $p<.05$, $\eta^2=.02$. Even more influential than politeness strategy (in terms of effect size) were familiarity of the robot ($F(1,344)=133.64$, $p<.001$, $\eta^2=.27$) and size of the request ($F(1,344)=343.79$, $p<.001$, $\eta^2=.33$). The statistically significant pairwise comparison was direct vs. positive politeness, p=.006. See Figure 3.

Appropriateness of Requests

The robot's politeness strategy influenced how appropriate the request seemed to be, $F(3,343)=6.78$, $p<.001$, $\eta^2=.06$. The significant pairwise comparisons (with Bonferroni correction cut-offs at .008) included: positive vs. indirect (p=.000), positive vs. direct (p=.002), and indirect vs. negative (p=.002).

Being more familiar with the robot ($F(1,343)=9.48$, $p<.01$, $\eta^2=.02$) and being more of a peer than an assistant ($F(1,343)=6.88$, $p<.01$, $\eta^2=.03$) also improved the perceived appropriateness of the request. What influenced request appropriateness most was the size of the request ($F(1,343)=71.89$, $p<.001$, $\eta^2=.17$) with smaller requests being more appropriate than larger requests. See Figure 4.

Qualitative Results

In most cases, people said they would do fairly helpful things for the robot, but not all. In some instances, we also learned more about *why* they would help the robot:

- "I will help him. After all he is also helping me with all the manual tasks!"
- "I would help him cause he would be able to solve it by himself next time, and I could need some help from him in another situation"
- "I would certainly be willing to help the robot, not because I'm actually helping the robot so much as I'm helping other people by helping the robot"

We also gained insight from the reasons that people gave for *not* being willing to help the robot:

- "I'm not sure how I can teach the robot how to use the copier."
- "I would avoid helping the robot if possible. The company is too cheap to hire a copier technician to come fix the machine?"
- "I would not help. It is the robots handlers job to teach it how to diagnose not mine."

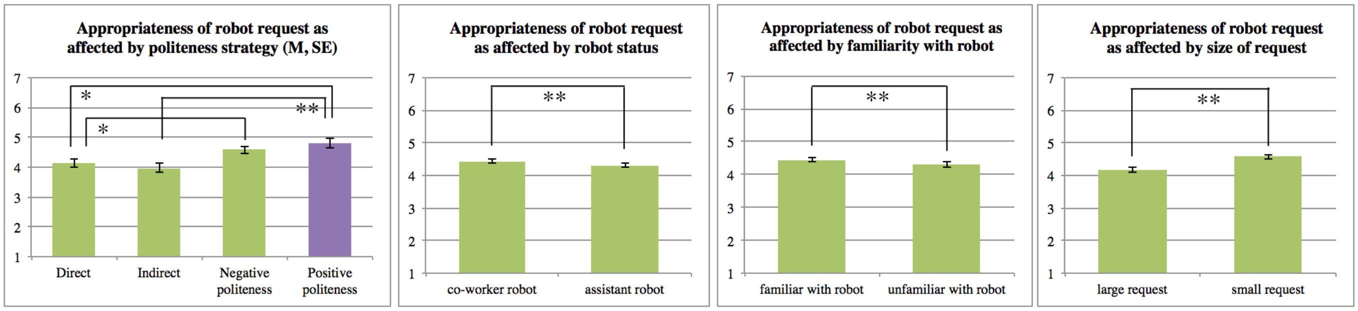

Figure 4. Study 1: Means and SEs for factors that influenced perceived appropriateness of the robot's request (* $p<.05$, ** $p<.01$)

Aligning Results with Hypotheses

Hypothesis 1 was partially supported – people felt that the peer robot made more appropriate requests for help than the assistant robot (despite their identical requests), but we did not find any statistically significant differences for their willingness to help the robot. Hypothesis 2 was supported – the more familiar the robot, the more likely people were to be willing to help it and saw its requests as being appropriate. Hypothesis 3 was supported – people were more inclined to help the robot with smaller requests than larger ones.

STUDY 2: BEHAVIORAL LAB STUDY

Based upon the results of Study 1, we decided to select the positive politeness strategy for subsequent robot requests for human help. To hone in on other important factors in human-robot interactions, we ran a behavioral experiment with a physical robot. This time we chose to focus upon the question: How does the perceived source of a robot's behaviors influence people's willingness to help a robot? As such, we conducted 3-level (autonomous, single operator, multiple operators) between-participants experiment.

Hypotheses

Studies [4, 19] showed that source orientation affected human perception of computers. This study extended to robots. Separately, privacy is important in many personal robotics contexts; we addressed it in case perceived privacy invasions would result in unwillingness to help the robot. We formed the following hypotheses:

1. People will perceive the robot differently if they believe it is autonomous vs. teleoperated by people.

2. People will be more or less willing to help the robot with its request, depending upon whether they believe it is autonomous vs. teleoperated by people.

3. Because people are sensitive about privacy when interacting with teleoperated robots [3], they will be least trusting of the robot that is teleoperated by multiple operators.

Participants

A total of 48 participants completed the behavioral in-lab study – 16 in each of the three experiment conditions. The participants included 25 males and 23 females with an age range of 18-67 years (M = 32.6, SD = 11.7). The participants were invited to participate using local university's listserv, postings at two local area listservs, and word of mouth from participants of previous studies at our institution. They were randomly assigned to condition. Of the 48 participants, 34 reported having owned a pet at some point in time (71%) and 12 owned a pet at the time of the study (25%). The participant's familiarity with robots was low (M = 3.38, SD = 2.08 on the same 7-point scale as Study 1). They were given a $15 amazon gift card as a token of thanks for participation.

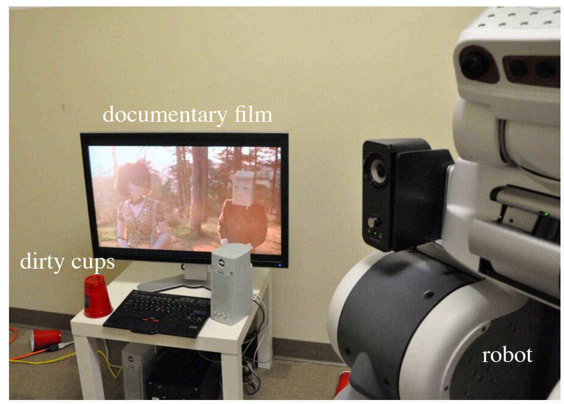

Figure 5. Study 2: Experiment setup from the participant's view

Materials

The robot used in this study was the PR2 (Personal Robot 2) shown in Figures 1 and 5, which was developed by Willow Garage. The robot had two arms, which were used for manipulation tasks to serve drinks and chocolate during the course of the study. The robot also moved on four wheels, traveling at maximum speed of 0.5m/sec for the study. The PR2 was operated in a semi-autonomous mode throughout the study across all experiment conditions. While an operator (Wizard) stepped through the conversation turns of the robot, the PR2 autonomously navigated to the participant and autonomously picked and placed objects such as the soda cans and chocolate.

Procedure

The participants were met at the lobby and led to a room in which the study took place. There, the participants were in-

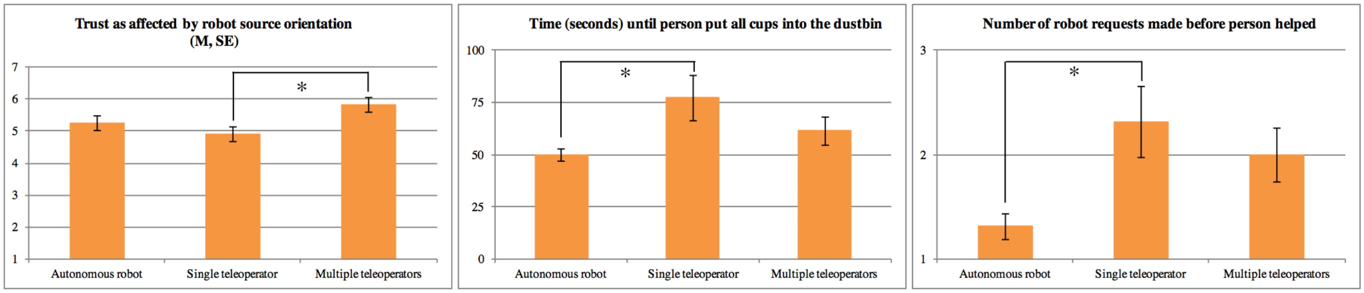

Figure 6. Study 2: Mean and SEs for source orientation effects upon user trust and behavioral willingness to help the robot

formed that the study would involve watching a film and being quizzed on the materials presented in the film. The participants were told that they would also be interacting with a service robot, which had a light on top of its head. If the light was blue, the robot was operating in the autonomous mode. If the light was green, the robot was being tele-operated (a person is driving it around) by a single operator; see Fig. 1. If the light was red, then multiple people were tele-operating the robot. The participants were then given a consent form. If they signed it, then the experimenter started playing the film and left the room.

While the participants were watching the film, the robot approached the participants and listed a selection of the drinks available before asking the participants choice of drink. Then the robot fetched the wrong drink and performed a service recovery behavior – apologizing [9] and offering to fetch the correct drink. The robot then informed the participant that it would be right back, and fetched a chocolate for them.

At the end of the study, we simulated a situation in which the robot might also request help from a person (similar to the pen drop task used to study human helping behaviors [10, 13]). After the robot fetched chocolate for the person, the robot informed the participants that the room looked dirty and offered to clear the dirty cups (See Figure. 5). The robot then requested help from the participant with clearing the cups, using positive politeness strategy.

The participant then filled out a questionnaire for their perceptions of trust and service quality. Then we collected demographics. This was followed by a short interview and final study debriefing.

Measures

Attitudinal Measures

Attitudes were collected via the post-interaction questionnaire and consisted of measures of perception of the robot and feelings of trusting the robot.

- *Trust* - Trust was an index of six items from the questionnaire and used the Wheeless and Grotz trust scale [26]. Participants indicated how well the following word pairs described their perceptions of the robot: "Kind" – "Cruel", "Distrustful" – "Trustful", "Deceptive" – "Candid", "Not deceitful" – "Untrustworthy", "Tricky" – "Straightforward", and "Faithful" and "Unfaithful". The index was reliable ($\alpha = .83$)
- *Service Quality* - We modified the SERVQUAL questionnaire [14], including 9 questions, each rated on a 5 point scale (1=definitely false or strongly disagree to 5=definitely true or strongly agree). Only 5 of those questions formed a reliable index ($\alpha = .77$):
  - When I have problems, the PR2 is sympathetic and reassuring.
  - PR2 is polite.
  - PR2 does the job well.
  - PR2 does not know what my needs are (reversed).
  - PR2 does not have my best interested at heart (reversed).

Behavioral Measures

The behavioral measures were collected from the video footage taken from three camera viewpoints and were coded for the following behaviors:

- *Whether or not the participant helped the robot* - Whether or not the person picked up any cups to help the robot.
- *Helping: Duration of help time* - The time interval (seconds) between the request made by the robot and the last cup dropped into the dustbin. We used this measure to gauge how quickly people would work to help the robot.
- *Helping: Number of requests needed* - Number of requests made by the number before the participants helped

Results

Manipulation Check

To ensure that participants understood and remembered who was controlling the robot, we asked a manipulation check question at the very end of the experiment: How was PR2 controlled? Autonomous, a person, or multiple people. Then we ran a cross-tabulation of this manipulation check against the actual assigned experiment condition, which checked out, Spearman correlation = .74, $p<.001$.

Attitudes

- *Trust* - A significant main effect of agent source orientation on trust was found, $F(2,45) = 3.502, p < .05, \eta^2 = .005$. The participants who interacted with multiple operators felt more trust for the robot, $M = 5.822, SD = .94$ than participants who interacted a single operator conditions $M = 4.89, SD = .90$, pairwise comparison (p=.007). The autonomous robot was in between, $M = 5.25, SD = .95$.

- *Service Quality* - We did not find any statistically significant differences in terms of how source orientation affected perceived service quality, $F(2, 45) = 0.44$, $p = .65$.

Behaviors

- *Whether or not the participant helped the robot* - A Chi-square analysis found no statistically significant differences between conditions, $\chi^2(df = 2) = 4.17$, $p = .12$.

- *Helping: Duration of help time* - A significant main effect of agent source orientation on the time taken to put all the cups into the dustbin was found, $F(2, 45) = 3.301, p < .05$, $\eta^2 = .03$. The participants who interacted with the autonomous robot took far less time, $M = 49.88$ seconds, $SD = 11.64$, than participants who interacted with the robot in the single operator condition $M = 77.19$ seconds, $SD = 43.08$, pairwise comparison (p=.014). The robot in the multiple operator was in between, $M = 61.27$ seconds, $SD = 27.06$.

- *Helping: Number of requests needed* - A significant main effect of agent source orientation on number of requests made by the robot was found, $F(2, 45) = 4.020, p < .05$, $\eta^2 = .04$, with the autonomous robot having to make fewer requests, $M = 1.31, SD = 0.48$, than the robot operated by a single operator, $M = 2.31, SD = 1.35$, pairwise comparison (p=.008). The robot in the multiple operator condition was in between, $M = 2.00, SD = 1.03$.

In general, people were quicker to help the autonomous robot than one operated by a single teleoperator. Between the teleoperation conditions, people felt more trusting of the multiple operators than the single operator (See Figure. 6.)

Aligning Results with Hypotheses

In summary, Hypothesis 1 was not supported by the data – we found no statistically significant differences in perceptions of service quality or trust as affected by autonomous vs. teleoperator source orientation. Hypothesis 2 was supported – people were quicker to help the autonomous robot than the one that they perceived as being teleoperate by a single operator. Hypothesis 3 was not supported and went in the opposite direction of what we expected – people were actually most trusting of the robot they thought was being teleoperated by multiple operators (as opposed to a single operator).

DISCUSSION

Study 1

In the first study, we found that people were more willing to help a robot that made smaller requests (i.e., requests that require less effort to comply with) rather than large. People were also more willing to help a robot that was more familiar (e.g., they had supposedly worked with the robot for 10 years) rather than new. Even politeness strategy affected people's willingness to help to help the robot – specifically, people were more willing to help the robot that used a positive politeness strategy. Similarly, making smaller requests, being a more familiar robot, being more of a peer (rather than lower status), and using a positive politeness strategy all made the robot requests seem more appropriate.

These human-robot interaction findings are reasonably aligned with work in human-human requests. The positive politeness strategy stood out as being more effective than the other strategies in these particular human-robot interaction scenarios. It is possible that other politeness strategies would be more effective in other scenarios (e.g., when the people are not busy at all, when the robot's task is more mission critical).

An intriguing part of the Study 1 dataset was actually the set of open-ended responses to the question, "What would you do in this situation?" We were surprised by the thoroughness of people's responses to this question. In particular, we were drawn in by the reasons that people gave for their willingness to help the robot, e.g., reciprocation, future benefits, and ultimately helping other people (as opposed to helping the robot itself). People also gave thoughtful reasons for why they would not be willing to help the robot, e.g., inability to help the robot, perceptions of the company being too cheap to hire a person, or perceptions that it's the robot programmers' job to program it to do more.

Consistent with previous work [16], people reported being concerned about their availability to help the robot. When it came to doing tasks that might take a few seconds (e.g., picking a few papers up off of the ground for the robot), their responses to the situation were quite consistently positive. However, larger requests gave people pause, e.g., "Wait, a month of full-time training? I'll only do that when I don't have anything else to do for a month - they're not paying me to teach the robot anyway, it might just get in the way of my real responsibilities." Comments like these highlight questions about how much time it really would be reasonable to ask people to spend, helping robots.

Study 2

In the second study, we focused on source orientation effects. Despite the robot interactions being secondary tasks to the primary task of watching the film, people were still very inclined toward helping the robot. Only 2 participants out of the total 48 did not help the robot by picking up any cups to clean up the room. Relative to prior field experiments, people were more willing to help the robot than we expected; this is possibly due to the laboratory setting. Because of this ceiling effect in our data, our statistically significant results were found in the more nuanced behavioral measures of willingness to help – how quickly the person helped the robot.

We found that people were nearly 50% quicker to help the PR2 when they believed that it was behaving autonomously rather than being teleoperated by a person. If anything, we had expected there to be more social pressure to help a human teleoperator. However, this was not the case. It is possible that people thought the teleoperated robot should have been more capable or perhaps that the human operator was being lazy by asking for so much help. Indeed, several participants helped the robot to collect the cups, but they placed the cups in front of the robot instead of dropping them into the dustbin. They expected the robot to perform that last step on its own: "It asked me to collect the cups, but I know it could [do that itself] because it got me the Mountain Dew... so I collected it and put it on the table and said, you know, 'Take it.'"

In Study 2, we found that people were more trusting of multiple operators as opposed to a single operator, which was surprising. Based on observations about feelings of privacy invasion when faced with teleoperated robots [3], we had thought that people would be more wary about a robot that was operated by multiple operators than in either of the other two situations. It is possible that people felt the multiple operators were more anonymous (e.g., the crowd), less personally intrusive, than an individual operator. Because we ran this as a between-participants study, we could not ask participants to compare the single vs. multiple operator situations against each other; a follow-up experiment would be necessary to identify causal factors.

Implications for Design, Theory, and Methods

The results of these studies provide some implications for design of personal robots: Using a positive politeness strategy (as opposed to the other three strategies) is likely to be more effective for robots that need help from people, especially when the robot is more familiar and perceived as being a peer (as opposed to having lower social status). As such, robots that can identify familiar vs. unfamiliar people might be able to solicit more help from those familiar people with whom they have built up some history. Designing the robot's relative social status in the organization to being more peer-like (rather than lower in social status) might also improve people's willingness to help the robot when it makes requests. Furthermore, if the robot is going to need more help from users or bystanders, consider letting people know that the autonomous robot does not have the help of remote operators.

In terms of theory, these findings provide support for the relevance of Politeness Theory [2] and the Computers as Social Actors paradigm [15] to the area of human-robot interaction. However, these theories may not be blindly applied to human-robot interactions, e.g., indirect politeness might not be as effective for robots as it is for human requesters of help.

In terms of methods, this pair of experiments provide an example of how to approach a large design space with many factors (online video prototyping of human-robot interactions) and then hone in on more specific research questions in follow-up behavioral science experiments (in-person studies with physical robots). Because personal robotics is still in its relatively early days, it is often challenging to get a robot prototype to perform as planned in a manner that is consistent enough to run large-scale studies. This combination of methods can provide more confidence in the translation of research results from video prototypes to physical robots, studying local as well as global participant pools, and wrangling a large set of independent variables down to a smaller set.

LIMITATIONS

As with any experiment, there are limitations to each research approach. As an online experiment, Study 1 was effective at quickly gathering a broad swathe of responses from a wider demographic than those who were geographically proximal to our lab. However, the online experiment was limited in many ways: (1) people were only able to watch a video of the robot, not actually interact in the same physical space with the robot and (2) people were not able to measure behaviors, only intentions to help the robot. The lab experiment was able to address both of these limitations from the first study, but Study 2 was not a replication of Study 1. Study 2 also had its own limitations: (1) participants in the study came from the local geographic area, (2) interacting with the robot was a secondary task, but they still seemed to be more engaged with the robot than they would have been if this were in the field, and (3) the red LED light might have been viewed negatively because of cultural associations with that particular color.

In both studies, we limited the scope to the PR2 robot (so we cannot necessarily generalize the findings to other robot forms), English-speaking participants (so we cannot necessarily generalize to other languages), a specific set of tasks (so we cannot necessarily generalize to all types of contexts and robot tasks) and a specific set of measures. Furthermore, we were only able to examine a specific set of variables (social status, familiarity, request size, and source orientation), but there are likely many other factors at play when it comes to how people will respond to robot requests for help.

FUTURE WORK

The lessons learned from these studies and the limitations of these studies both open up many possibilities for future work, including examining variables that were missing in our current studies (e.g., culture, language, various robot forms), looking at different contexts and tasks (e.g., security guard robots, hotel service robots), and running field experiments that have more external validity than online video prototype studies or in-lab experiments. We have presented our study protocols here in hopes that others may find them to be useful for their own work, using different robot forms and interaction scenarios. A few particularly ripe areas for future research are: (1) more directly comparing what people say they would be willing to do vs. what they will actually do when faced with these human-robot interaction situations, (2) testing the longevity or decay of people's willingness to help robots and sussing out the causal factors for those trends over time, and (3) explicating the deeper reasons behind people's willingness or unwillingness to help robots in human-robot interaction situations.

CONCLUSIONS

This paper presents empirical findings of two experiments that examines how robots might be designed to more effectively solicit help from people, which could enable those robots to be more capable in environments like a workplace. Through these experiments, we found that: (1) positive politeness strategies are more effective than other politeness strategies in these kinds of human-robot interaction scenarios, (2) robots that are more familiar, have more peer-like social status, and that make smaller requests are more likely to get help from people, and (3) people are quicker to actually help a robot when they believe it is autonomous as opposed to being teleoperated by a person. Together, these findings provide support for theories of Politeness and Computers as Social Actors. They also demonstrate how the rules of human interpersonal interactions may not always directly apply to human-robot interactions.

REFERENCES

1. T. Beran, A. Ramirez-Serrano, R. Kuzyk, S. Nugent, and M. Fior. 2011. Would Children Help a Robot in Need? *International Journal of Social Robotics* 3, 1 (2011), 83–93.
2. P. Brown and S.C. Levinson. 1978. *Universals in language usage: Politeness phenomena*. Cambridge University Press. 56–311 pages.
3. D.J. Butler, J. Huang, Fr. Roesner, and M. Cakmak. 2015. The Privacy-Utility Tradeoff for Remotely Teleoperated Robots. In *Proceedings of HRI 2015*. ACM, 27–34.
4. D. Eckles, D. Wightman, C. Carlson, A. Thamrongrattanarit, M. Bastea-Forte, and B.J. Fogg. 2009. Social Responses in Mobile Messaging: Influence Strategies, Self-disclosure, and Source Orientation. In *Proceedings of CHI 2009*. ACM, 1651–1654.
5. S. Fukushima. 2003. *Requests and Culture: Politeness in British English and Japanese, 3rd edition*. Peter Lang International Academic Publishers.
6. H. Hüttenrauch and K. Severinson Eklundh. 2006. To help or not to help a service robot: Bystander intervention as a resource in human–robot collaboration. *Interaction Studies* 7, 3 (2006), 455–477.
7. K. Kinzer. 2009. Tweenbots. (2009). http://www.tweenbots.com/ Accessed: 2015-09-13.
8. W. Labov and D. Fanshel. 1977. *Therapeutic discourse: Psychotherapy as conversation*. Academic Press.
9. M.K. Lee, S. Kielser, J. Forlizzi, S. Srinivasa, and P. Rybski. 2010. Gracefully Mitigating Breakdowns in Robotic Services. In *Proceedings of HRI 2010*. IEEE Press, 203–210.
10. N.C. Macrae and L. Johnston. 1998. Help, I need somebody: Automation action and inaction. *Social Cognition* 16 (1998), 400–417. Issue 4.
11. P.P. Maglio, S.J. Matlock, T.and Gould, D. Koons, and C.S Campbell. 2002. On understanding discourse in human-computer interaction. In *Proceedings of Cognitive Science*.
12. M.P. Michalowski, C. DiSalvo, D. Busquets, L.M. Hiatt, N.A. Melchior, R. Simmons, and S. Sabanovic. 2006. Socially Distributed Perception. In *Proceedings of HRI 2006*. ACM, 349–350.
13. D.T. Nguyen and J. Canny. 2009. More Than Face-to-face: Empathy Effects of Video Framing. In *Proceedings of CHI 2009*. ACM, 423–432.
14. A. Parasuraman, V.A. Zeithaml, and L.L. Berry. 1988. Servqual. *Journal of Retailing* 64, 1 (1988), 12–40.
15. B. Reeves and C. Nass. 1996. *The Media Equation: How people treat computers, television, and new media like real people and places*. CSLI Publications and Cambridge university press.
16. S. Rosenthal, M. Veloso, and A.K. Dey. 2012. Is Someone in this Office Available to Help Me? *Journal of Intelligent Robotic Systems* 66, 1-2 (April 2012), 205–221.
17. F. Smith, D.and Zeller. 2014. HitchBot. (2014). http://m.hitchbot.me/ Accessed: 2015-09-13.
18. I. Straub, S. Nishio, and H. Ishiguro. 2010. Incorporated identity in interaction with a teleoperated android robot: A case study. In *Proceedings of RO-MAN 2010*. 119–124.
19. S.S. Sundar and C. Nass. 2000. Source Orientation in Human-Computer Interaction Programmer, Networker, or Independent Social Actor. *Journal of Communication Research* 27, 6 (2000), 683–703.
20. J.-Y. Sung, L. Guo, R. Grinter, and H. Christensen. 2007. My Roomba Is Rambo: Intimate Home Appliances. *Proceedings of UbiComp 2007* (2007), 145–162.
21. A. Tapus, M.J. Mataric, and B. Scasselati. 2007. Socially assistive robotics: Grand challenges in robotics. *Robotics Automation Magazine, IEEE* 14, 1 (March 2007), 35–42.
22. C. Torrey, S. Fussell, and S. Kiesler. 2013. How a robot should give advice. In *Proceedings of HRI 2013*. IEEE Press, 275–282.
23. K. Tracy, R. T. Craig, M. Smith, and E. Spisak. 1984. The discourse of requests: Assessment of a compliance-gaining approach. *Human Communication Research* 10 (1984), 513–538.
24. E.A. Weinstein and C.R. Black. 1969. Factors mediating the effects of others' responses on the self. *Sociological Inquiry* 39 (1969), 189–193. Issue 12.
25. A. Weiss, J. Igelsböck, M. Tscheligi, A. Bauer, K. Kühnlenz, D. Wollherr, and M. Buss. 2010. Robots asking for directions: the willingness of passers-by to support robots. In *Proceedings of HRI 2010*. IEEE Press, 23–30.
26. L.R. Wheeless and J. Grotz. 1977. The measurement of trust and its relationship to self-disclosure. *Human Communication Research* 3, 3 (1977), 250–257.
27. Y. Yamamoto, M. Sato, K. Hiraki, N. Yamasaki, and Y. Anzai. 1992. A request of the robot: an experiment with the human-robot interactive system HuRIS. In *Proceedings of the IEEE International Workshop on Robot and Human Communication*. 204 –209.