How a Robot Should Give Advice

Cristen Torrey
Adobe Systems
San Francisco, CA
torrey@adobe.com

Susan R. Fussell
Dept. of Communication
Cornell University
Ithaca, NY, USA
sfussell@cornell.edu

Sara Kiesler
Human Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA, USA
kiesler@cs.cmu.edu

***Abstract*—With advances in robotics, robots can give advice and help using natural language. The field of HRI, however, has not yet developed a communication strategy for giving advice effectively. Drawing on literature in politeness and informal speech, we propose options for a robot's help-giving speech—using hedges or discourse markers, both of which can mitigate the commanding tone implied in direct statements of advice. To test these options, we experimentally compared two help-giving strategies depicted in videos of human and robot helpers. We found that when robot and human helpers used a hedge or discourse markers, they seemed more considerate and likeable, and less controlling. The robot that used discourse markers had even more impact than the human helper. The findings suggest that communication strategies derived from speech used when people help each other in natural settings can be effective for planning the help dialogues of robotic assistants.**

***Index Terms*— Communication, dialogue, assistive robots, social robot, human-robot interaction, politeness, mixed-method**

Fig. 1. A robot helper giving advice to a novice making cupcakes.

I. Introduction

In an imagined technical future, robots are intelligent and patient, able to assist people in a range of tasks. This vision of robots as helpers is being actively pursued in various academic and commercial research projects. Robots are guides in train stations [14], reference librarians [2], grocery shopping assistants for the blind [11], rehabilitation exercise coaches [17], personal assistants in the home [35], and weight loss coaches [18]. This vision of robots as helpers is beginning to be available commercially, although still limited in its realization.

Giving others advice about how to do things is a tricky business. Sometimes this help is appreciated and other times good intentions are not enough to avoid giving offense. As natural language interfaces improve and robots begin to take on the role of the advice giver, we want to know what kind of communication strategies a robot might use to navigate potentially sensitive help-giving situations. In this paper we explore social responses to a robot that offers its listeners help. We ask two related questions. First, how do people respond to different help-giving strategies? And, second, do responses to robot helpers differ from responses to human helpers?

Two decades ago, these questions would have paled next to the challenges of speech recognition, speech synthesis, and robot hardware [19]. Since then, research in human-robot dialogue for giving help has come a long way, although there continue to be important problems in speech misrecognition, detecting differences in social context and user intention (situation awareness), and adapting robot behavior to user intent [31]. Researchers have begun to tackle not only the transmission of help, but the style in which it is conveyed. The necessity of a robot's determining and establishing common ground with users is an active research area (e.g., [22]). Advances also have been made in determining socially appropriate verbal and nonverbal behavior such as how close a robot should stand to users [36] and whether it should apologize for errors [20]. Little attention has been paid, however, to how robots should speak to the people they are presuming to help. In many cases, the robot is polite (says "please" and "thank you") and encouraging ("Right!") but uses direct commands to convey advice or answers to questions. Here, for example, is a plan for a robot to help someone take a train:

Go to the station, buy a ticket, check the departure board for track information, go to the track, board the train. . Enjoy!

In this example (taken from a robotics talk), and in our own work [33][34], specific word choices and sentence construction are based on intuition on how best to offer information rather than on a careful examination of how people actually offer help. Do helpers simply enumerate directions or steps in a task? As we explain below, the literature in

978-1-4673-3101-2/13/$31.00

communication suggests that this direct approach is likely to come across as threatening or condescending.

To look more closely at how people give help, we conducted a field study of help-giving dialogues during a cupcake baking task [32]. This task represents a class of activity where the desired outcome is open-ended and the necessary steps are somewhat flexible. We identified several strategies used by human help givers to avoid making direct claims or giving direct orders to help recipients. Specifically, helpers often used hedges (words like “I think” and “probably”) and discourse markers (words like “I mean” and “so”) to lessen the forcefulness of their advice. These communication strategies distance helpers from making strong, direct claims on the listeners’ subsequent actions.

Our observations of human help givers in the cupcake task suggest strategies that may be useful for robot help givers, but they do not directly test people’s responses to these strategies. In the current study, we compared social responses to robot help givers with social responses to human help givers. We manipulated whether the robotic or human help-giver used or did not use hedges, and used or did not use discourse markers. The results inform the design of dialogue systems for help-giving robots. In the sections that follow, we describe the literature on which we drew, hypotheses, and experimental design.

A. Politeness theory

Many researchers have observed that speakers engage in indirect speech, particularly when navigating what Goffman calls a “face-threatening act” [12]. Offering someone help can be a face-threatening act, one that helpers navigate carefully to avoid insulting the recipient. Goffman refers to the delicacy observed in communication around face-threatening acts as “the language of hints” [13]. Fraser [10] and others describe this communicative work as “mitigation,” a strategy for reducing the force of a message.

Politeness theory [3] describes different ways speakers construct messages to support avoid threats to listeners’ autonomy. For example, instead of asking someone to “take the kettle off the stove,” they might use another approach, such as “I think the water might be boiling.” Politeness theory proposes that people use these mitigating communication forms in order to avoid giving orders to other people, a very threatening action. Although politeness theory specifically describes the communication features around making a request, the same principles may apply to help giving as well. Giving help or advice is not, strictly speaking, a request, but the act can similarly threaten the listener’s autonomy, so one might expect successful help givers to use politeness cues to mitigate the force of their language.

One common strategy for mitigating the force of a message is the use of hedges. Hedges literally express uncertainty; they include qualifying types of language such as “I guess,” “maybe,” “probably,” “I think,” and “sort of.” They are forms of speech that communication researchers call negative politeness, that is, words and phases that limit the universality of the statement, allowing the listener to disagree if necessary. Negative politeness protects listeners from threats to their autonomy whereas positive politeness (such as using “please” and “thank you”) encourages social connection and rapport [3]. Hedges are a form of negative politeness that speakers use mitigate the force of their communication. Indeed, the use of hedges in human conversation is so common that their removal makes speech seem aggressive.

From politeness theory, we hypothesized that speakers who use hedges in their help messages would be perceived as more considerate, less controlling, and more likeable than those who do not use hedges. Based on the prior work on social responses to technology [21][24][25], we would not anticipate significant differences between human helpers and robot helpers, especially because speech makes robots seem human-like. We thus predicted in H1 that social responses to robot helpers would mirror responses to human helpers:

H1: Human and robotic speakers using hedges will be rated more considerate, less controlling, and more likeable than speakers not using hedges.

B. Informal speech

When people speak, their speech often contains various types of disfluencies called “discourse markers.” Discourse markers include repeated words, false starts, and fillers such as “uhm.” In our observations of human help-giving communication [32], we noted that helpers often showed a pattern of consistent, repetitive use of discourse markers such as “like you know,” “I mean,” “well,” “just,” “like,” and “yeah”. These words operate at a pragmatic level; their meaning is derived not exclusively from their literal definition but from their use in context.

There are several different interpretations of the use of discourse markers (see Fischer [8] for a review). Some researchers think discourse markers are related to speech production difficulties. One interpretation of speakers who use “like” frequently is that they are having trouble finding the right words [29]. The phrase “you know” also may indicate the speaker is stalling for time [15].

Other researchers believe the use of discourse markers, particularly the phase “you know,” is associated with more casual speech or speech between young people [30]. “You know” invites addressees to make an unspoken inference [16]. It may be common within younger communities or in more casual situations because speakers in these contexts are more willing to engage with their addressees’ interpretations [9]. Between young people, discourse markers may appeal to positive politeness by reinforcing the similarity between the speaker and the listener.

Discourse markers can help mitigate the impact of help messages. By saying “you know” but remaining vague about the details, speakers may be engaging in a form of negative politeness. Using “you know” allows helpers to be less explicit about the direction they are giving and to give greater weight to the listeners’ interpretations. Andersen [1] describes talk involving the use of “like” as “loose,” meaning it reduces the commitment of the speaker to what follows. “I mean” warns of upcoming adjustments [26]. Helpers who frequently use “I mean” in their help messages may be more comfortable adjusting their messages as they are being produced, or they may be particularly sensitive to the way the message could be interpreted [9]. The use of "I mean" may signal a casual, flexible way of speaking that could impact negative politeness as well. If the way a message is given is in flux, then it may signify the speaker would modify if pressed.

There is no single theory that predicts perceptions of discourse marker use in help dialogues. Nevertheless, the prior work suggests that discourse markers can distance the speaker from making a strong, direct claim on the listener. Therefore, we hypothesized that the use of discourse markers would improve perceptions of the speaker as more considerate, less controlling, and more likeable. Discourse markers soften commands by sounding casual. Under the assumption that social responses to robot helpers will mirror responses to human helpers, we posited:

H2: Human and robotic speakers using a style of speech containing several discourse markers will be rated more considerate, less controlling, and more likeable than speakers not using discourse markers.

C. Combining hedges and discourse markers

The first two hypotheses anticipate a benefit for the use of both hedges and discourse markers. Although there is no specific prior work that observes the combination of these two communication strategies, it is possible that their combined effect would be additive. That is, the use of both strategies together could be better than either strategy alone. In the observational work [32] the use of frequent discourse markers was accompanied by the use of hedges, so we hypothesized that their combination would be particularly effective.

H3: Human and robotic speakers using both hedges and discourse markers will be rated more considerate, less controlling, and more likeable than speakers using either strategy in isolation.

II. Method

To compare listeners' impressions of robot and human speakers who give advice using hedges and discourse markers, we ran an experiment using videos of human and robot helpers giving advice to a person learning to make cupcakes. We used a within-subjects design in which participants saw videos of human and robot helpers. To create the four communication conditions of the experiment, a direct help message (with neither hedges or discourse markers) was created. Three additional help messages added hedges, discourse markers, and their combination, for a total of four communication, help-giving styles in four videos. We drew these help messages from the steps of the cupcake baking task during which participants struggled during the laboratory cupcake baking sessions.

If the condition included hedges, the help message used a single hedge. In our prior observations, hedges were used consistently by certain helpers when they offered information, but they were used one at a time. If the condition included discourse markers, the help message included three to four discourse markers. This condition specifically tests the effectiveness of frequently-observed discourse markers from the observational study. For each of the four communication conditions we created a script for a short, video vignette involving a baker and a helper. Table 1 illustrates the help messages in the scripts for each communication condition.

To compare impressions of human and robot helpers, we produced these video vignettes with a human helper and then digitally spliced a robot helper over the human helper in each video segment. Video vignettes were produced for all four communication conditions, one set showed a human speaker and a second set showed a robot speaker. Participants viewed all four communication conditions but they saw two communication conditions with human speakers, and two communication conditions with robot speakers. We counterbalanced order of condition.

TABLE I. Examples of help messages communicated in each communication condition, using the step "Cream butter and sugar."

	No Discourse Markers	Discourse Markers
No hedge	"The mixture should be smooth and fluffy. The color will get lighter too."	"Basically just keep going until it's like a smooth mixture. Lighter color and fluffier."
Hedge	"Until the batter looks smooth. It'll get kind of fluffier and the color will lighten up."	"And kind of mix it, until it's just like fluffy. Basically, a nice smooth consistency, a little bit lighter color."

A. Participants

Seventy-seven Carnegie Mellon University students and staff members, as well as members of the general Pittsburgh community, were recruited from Carnegie Mellon's experiment scheduling website. They were paid $10 for their participation.

B. Procedure

The experiment was conducted in a laboratory where the videos could be projected on a large screen with accompanying high-quality audio. When participants arrived at the laboratory, they were asked to answer several questions about baking, as a test of their experience with the task domain.

Participants were then informed they would be watching a series of short videos and would be asked to answer several questions following each video. Four video clips were shown. After each video clip, participants were asked to respond to several pages of items on a questionnaire. After viewing the final video clip, participants were paid for their participation and dismissed.

Because the experimental procedure required counterbalancing along three different dimensions (communication condition, human or robot helper, and male or female actor), it was infeasible to counterbalance completely. Instead, we conducted a total of sixteen sessions during which each communication condition was counterbalanced for the order of presentation, and the actor in each video clip was counterbalanced for order of appearance and for communication condition used.

C. Video Stimuli

Four scripts were created for use in the video vignettes, based on the four communication conditions previously described. Each script contained help messages for five steps in the cupcake baking task from measuring the flour to filling the cupcake tin with batter. The steps were identified as sources of frequent help giving in the prior observational study. Each video vignette was approximately three minutes in length and included a short introduction to the characters and the task. Following this introductory material, a statement of the current step and the baker's response to it was displayed before showing the help giving interaction. For example, prior to seeing the helper offer information about how long to mix the butter and sugar together with an electric mixer, participants saw the following message displayed on the screen: Evan reads from the recipe, '"Cream butter and sugar." In the subsequent frame, Evan is shown mixing the butter and sugar. Then, the helper is shown speaking the help message. Over the course of the three-minute video clip, five different baking steps are introduced. After each step is introduced, the help message is delivered.

Because different speakers using the exact same language (hedges and/or discourse markers) may generate different impressions based on paralinguistic qualities or physical appearance, we captured the dialogue from each script using four different actors. The actors were referred to in the videos with pseudonyms (Ada, Max, Rob, and Kit). To separate the influence of individual characteristics from the verbal aspects of the help messages, we hired four undergraduate drama majors (two men, two women) to act out all four scripts. By counterbalancing which actor was seen in which communication condition, it was possible to investigate the effect of the communication condition independent of the effect of the actor's appearance, voice and other characteristics.

To compare impressions of robot helpers, we duplicated the set of video clips showing human help givers and overlaid video of a robot helper directly on top of the human help giver. In this way, the baker's behavior was kept consistent and the image of the helper is the only visual difference between the sets of videos. The robot is designed with variable forehead and chin shapes. These modifications were used to create four robots with slightly different facial features, to stand in for the four human actors in the videos. The robots were identified in the videos with the same pseudonyms as the human actors (Ada, Max, Rob, and Kit).

Fig. 2. Left: A human helper giving advice to a novice making cupcakes. Right: the image of the human has been replaced by the robot.

Because of the acoustic variability in the way the human helpers spoke (every actor spoke their lines in a subtly different way), it was important to keep those paralinguistic features consistent in the videos containing a robot helper. Synthetic speech is typically used for a robot's communication, but it was difficult to recreate the subtleties of naturally occurring speech (particularly speech containing discourse markers) with speech synthesis. Using the recorded human speech as a base, we modified the audio track to create a more metallic, harmonic sound. To create the impression that the audio in the clip was the robot's voice, the human audio track was modified by duplicating the track, changing the pitch on the duplicate track, and lowering the volume on this second track.

D. Measures

We observed participants' reactions to the help-giving scenarios with both quantitative and qualitative measures. After viewing each video, participants rated their agreement with a series of statements about the helper, the help message, and the task itself. Participants responded to 15 questionnaire items about their perception of the helper and the helper's communicative behavior. Factor analysis with varimax rotation indicated three underlying factors that accounted for 68% of the variance.

Considerateness: Six questions pertained to the participants' impressions of the speaker as considerate and supportive (e.g., "The helper is tuned into the baker's needs," "The helper was not considering the baker's feelings."). The scale is reliable (Cronbach's $\alpha = 0.86$).

Controlling: Two items pertained to participants' impression of the speaker as controlling ("The helper took control of the baking activity," "The helper spoke strongly about how the cupcakes should be made."; Cronbach's $\alpha = 0.79$).

Liking: Three items pertained to the participants' perceptions of positive rapport with the helper (e.g., "I like the helper," "If I were making cupcakes, I would want this helper around.") As a scale, these three items were highly reliable (Cronbach's $\alpha = 0.86$).

After considering their agreement with each statement on the questionnaire, participants were asked to describe the video, in their own words, for someone who had not seen it. The open-ended responses were transcribed and coded according to the *Linguistic Category Model* [7]. The Linguistic Category Model attempts to measure the level of abstraction with which people describe individuals. We used this coding manual to explore the differences between the way that human helpers and robot helpers were described in open-ended responses.

The coding manual defines five different codes: Adjectives, Descriptive Action Verbs, Interpretative Action Verbs, State Action Verbs, and State Verbs. Instances of the various verb categories were rare in this sample, to describe either human or robot helpers. The Adjective category was more common. Descriptions in this category, for example "the robot was aggressive," suggest that participants are ascribing qualities or traits to the helper, such as being "honest" or "aggressive." By using the adjective form, participants are describing helpers as having these qualities indefinitely, as opposed to engaging in a one-time activity, for example "the robot offered information about baking cupcakes."

III. Results

We first present the results from our survey measures; then we examine the adjectives used by participants to describe the helpers in the video using Fiedler's linguistic category model[7].

A. Survey measures

First, we considered the effects of communication condition and the impact of the human/robot helper variable on the questionnaire scales. Our statistical ANOVA model includes the effect of hedges (hedges vs. no hedges), discourse markers (markers vs. no markers), human vs. robot helper, and all related two-way and three-way interaction effects of those three independent variables. The model also includes several control variables: order of video (1-4), helper gender, participant age, number of years in the U.S., and baking expertise (number correct from pre-test).

Considerateness. As hypothesized in H1 and H2, hedges and discourse markers increased participants' agreement that the helper was considerate. Hedges and discourse markers alone, as well as both of them together, improved ratings of considerateness over direct speech, interaction $F(1, 211) = 23.49$, $p < .001$. (See Figure 3.) H3 was not supported; that is, using both hedges and discourse markers did not add together to increase perceptions of considerateness over either of these strategies used alone. Hedges and discourse markers worked for the human and the robot helper equally well, although male helpers were rated as significantly more considerate than female helpers, main effect $F(1, 212) = 22.96$, $p < .001$.

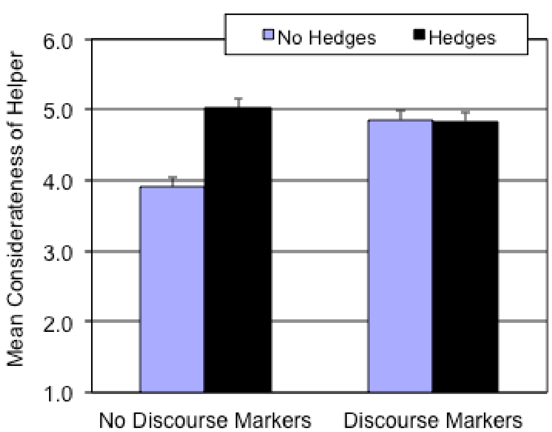

Fig. 3. Either hedges or discourse markers or both significantly increased perceptions that the helper was considerate, but the combination did not improve perceptions any further.

Controlling. Both the human and robot helper were rated equally as less controlling when they used either hedges, discourse markers, or both, as shown in the significant interaction, $F(1, 209) = 4.9$, $p < .05$. (See Figure 4.) Male helpers were perceived as less controlling than the female helpers even when the words they were using were identical, $F(1, 210) = 5.5$, $p < .05$.

The robot helpers were perceived to be less controlling than the human helpers, $F(1, 209) = 6.2$, $p < .05$. A significant interaction indicated that the impact of the robot helper on ratings of control was particularly evident when the robot used discourse markers, interaction $F(1, 277) = 5.28$, $p < .05$. (See Figure 5.) A robot helper using discourse markers was rated as significantly less controlling than a human helper using the identical discourse markers.

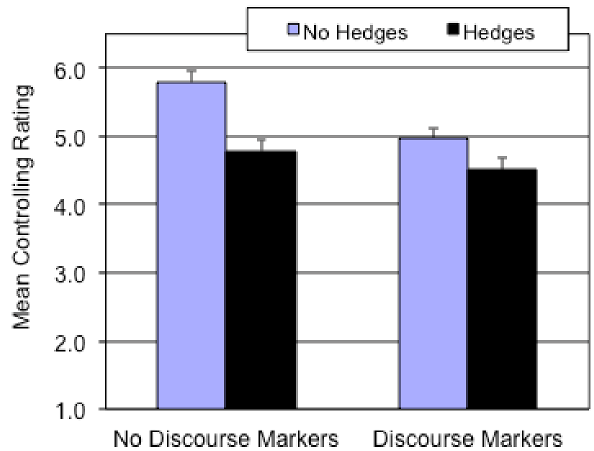

Fig. 4. Hedges and discourse markers significantly reduced ratings of control.

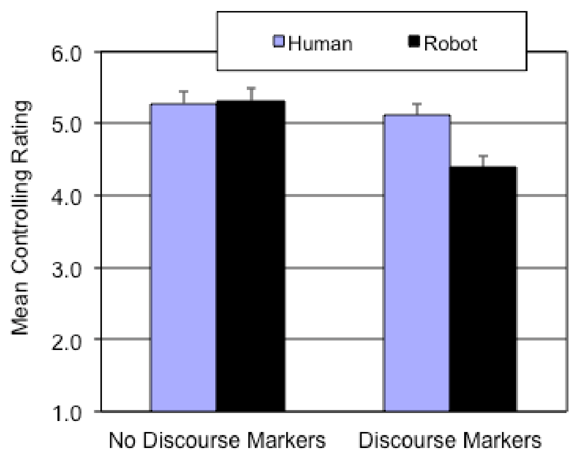

Fig. 5. Helpers using discourse markers were perceived to be less controlling, even more so when the robot used them.

Liking of the helper. Liking for the helper followed a similar pattern as judgments of whether the helper was perceived as considerate and controlling. When the helper used hedges, discourse markers, or both, participants liked the helper more than when the helper used direct speech, $F(1, 210) = 7.73$, $p < .01$. (See Figure 6.)

Several of the control variables influenced liking for the helper as well. Male helpers were better liked than female helpers, $F(1, 211) = 17.6$, $p < .001$. The older the participant, the more the participant liked the helper, $F(1, 76) = 5.97$, $p <$

.01. And, finally, there was a significant interaction between participant expertise and the use of the hedging strategy. The less expertise a participant had, the more the participant liked helpers that used the hedging strategy.

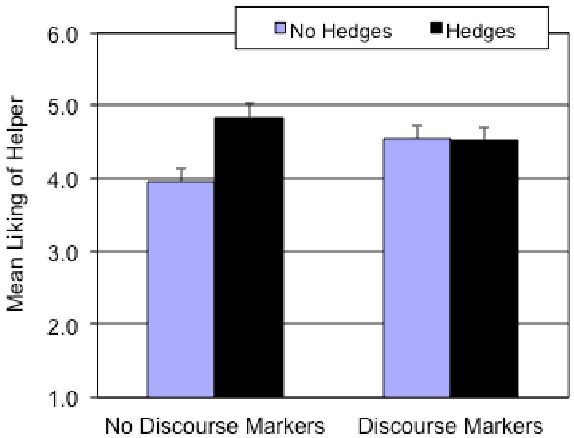

Fig. 6. The use of hedges, discourse markers, or their combination all increased measures of liking equally.

To summarize, consistent with H1 and H2, hedges and discourse markers were successful as a communication strategy for both human and robot helpers. But surprisingly, these two strategies were not more effective when they were combined, contrary to H3. Both strategies made relatively little impact when the other strategy was already present, and this was true for both robot and human helpers.

There were no differences in the less positive perceptions of human helpers and robot helpers when the help messages were direct (and used neither a hedging nor discourse marker strategy). Further, when robot helpers used discourse markers they were perceived as even less controlling than human helpers. So there is a little support for the idea that robot helpers cam be perceived differently than human helpers, and this difference can be more positive for the robot.

B. Open-ended descriptions of the videos

We also collected open-ended descriptions of the videos from the participants. These descriptions were written after participants saw each video, to explore whether participants would describe robot and human helpers differently. We measured this difference by coding each description for the use of adjectives describing the helper. The use of an adjective, such as "the aggressive helper," indicates the participant is ascribing a quality or trait to the helper, as opposed to saying something like, "the helper acted aggressively" which indicates the helper acted this way in a single instance.

We found more frequent use of adjectives when participants were describing human helpers than when they were describing robot helpers. (See Figure 7.) Using a log transformation to normalize the distribution of adjectives across participants and controlling for total number of words in each participants' description, adjectives were significantly more frequent in descriptions of human helpers than in descriptions of robot helpers, $F(1, 218) = 8.1$, $p < .01$. We did not find significant evidence that the use of adjectives to describe robot helpers affected participants' ratings of the helper as considerate, controlling, or likable.

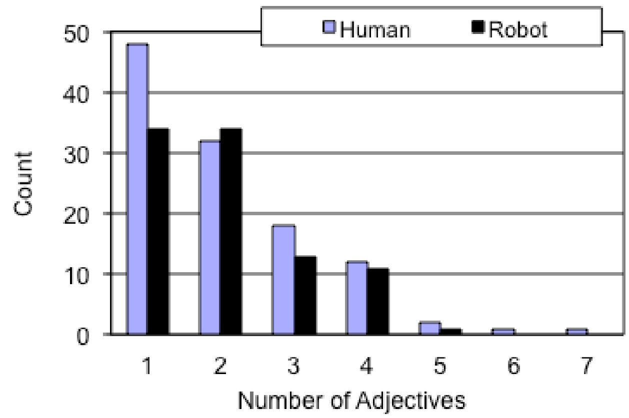

Fig. 7. More adjectives were used to describe human helpers than robot helpers per unit of speech.

IV. Discussion

The goal of this experiment was to explore social responses to robot and human helpers using different communication strategies. The two communication strategies, hedging and use of discourse markers, were drawn from the literature and modeled from observed human helpers' speech. We tested the impact of these strategies, individually and in combination, on impressions of human and robot helpers by asking participants for their reactions to a series of video vignettes.

Participants responded positively on several measures to the use of hedges and discourse markers, but the combination of the two strategies was not any more successful than the individual strategies alone. This is somewhat surprising, given the consistency with which we observed the combination of hedges and discourse markers in naturally occurring speech. We anticipated that the combination of the two strategies would be the most successful approach, but, in this experiment, we could not determine any benefit of using more than one strategy at a time.

Reactions to robots were, for the most part, similar to responses to humans. Direct speech, without either mitigation strategy, was judged to be the least desirable on several dimensions for both humans and robots. This means that robots will be held accountable for communication that might appear to be rude, condescending, or controlling, if the same speech were made by a human. Developers of robots that give directions, offer information, or provide other kinds of help cannot assume that a direct approach will be forgiven because it is used by a robot speaker. Robot dialogue planning will need to encompass the emotional context of speech beyond the positive politeness strategies that now well accepted [e.g., 23].

Surprisingly, the use of discourse markers improved perceptions of robots more than it improved perceptions of people. For example, when a robot says "I mean, yeah, just mix all that stuff in there," the robot was perceived to be significantly less controlling than when a human helper said the same thing. It is possible that people's lack of familiarity with this kind of speech in a robot is the source of this effect. Expectancy violation theory was first identified with respect to nonverbal behavior but has subsequently been explored for verbal behavior as well [4][5]. When something is unexpected, expectancy violation theory predicts greater levels of arousal and a subsequent stronger reaction to the stimuli. In this case, the use of discourse markers by a robot might have had a stronger effect when used by a robot, because it was unexpected.

Expectancies also could have influenced participants' overall perceptions of the robot versus human helpers. Regardless of communication condition, participants viewed robot helpers as being significantly less controlling than human helpers overall. Further, participants wrote fewer adjectives when describing the robot helper than when describing the human helper. These results deserves further investigation. They suggest that people have less complex cognitive conceptualizations of robots than they do of people (leading them to write fewer adjectives). They also may have less complex conceptualizations of robot "motives" and intents, leading to greater surprise when robots show complexity. Technology designers are typically discouraged from devising interactions that surprise people or are not what they expect, but this lack of similarity between humans and robots may be a positive opportunity to defy expectations, resulting in less domineering robot helpers.

A. Limitations and Future Directions

This experiment compared human helpers and robot helpers in as controlled a manner as possible, but there are a number of trade-offs associated with this approach. First, this experiment used a novel video vignette paradigm that did not require the participants to interact directly with the human or robot helper. This situation could have influenced the participant to take the helper's communication less seriously than if the communication had an instrumental impact on their own activity. Participants were mere bystanders. Ratings might have been milder than if the consequences of the help messages were more salient for participants.

On the other hand, when participants were viewing the scene as bystanders they had greater opportunity to notice and to dwell on the differences between the help givers, something they might not have the opportunity to do if they were engaged in baking the cupcakes themselves. This intense concentration on the help giver's communication could have accentuated participants' responses to the different strategies.

Another limitation of the design is that this experiment used actors to reenact scenes that were previously observed in naturally occurring speech. Although the actors did an excellent job with their lines, the dramatic characterization of the interaction may have had unforeseen effects.

Still another limitation is the particular instantiation of the help messages, with the use of a single hedge and/or three or more discourse markers in each help communication, may not allow us to generalize to other uses of hedges or other uses of discourse markers. We observed this pattern of communication behavior in naturally occurring speech, but we cannot claim that the impressions we found in this investigation would hold if the frequency and combination of hedges and discourse markers were varied in another way.

This experiment varied the use of hedges and the use of discourse markers but did not combine this language use with other potentially influential factors, such as nonverbal communication behavior, including the physical distance between the helper and the baker and prosodic cues. It is possible that the use of hedges or discourse markers could be quite different in a different tone of voice or a different speed of speech. This experiment used four speakers to vary some of these aspects of communication, but we cannot exclude the possibility that other nonverbal communicative elements, in combination with the communication manipulations, would alter impressions.

Impressions of helpers and of help messages are undoubtedly influenced by factors outside of the scenario depicted in these video vignettes. Help recipients care more about some activities than others, and they are likely to be more sensitive to help-giving interactions around activities with which they strongly identify. Various additional elements of the task, the help recipient, and the help provider could, undoubtedly, influence responses to receiving aid. For the sake of simplicity, this experiment looked closely at the influence of the help message, but this influence might weaken when other situational elements are factored into the equation.

Finally, this exploration of communication mitigation strategies was culturally limited to American English. The use of discourse markers, as a phenomenon associated with young, casual speech, might even be limited geographically to specific regions of the United States. In future research, HRI researchers might be better able to reflect on the generalizability of positive impressions of hedges and discourse markers, but for the present, we acknowledge the limitations of this exploration of politeness as culturally bound.

V. Conclusion

HRI researchers have recognized the huge potential of robots serving as assistants. In many cases, help will be delivered through speech because the recipients of help need to hear advice in language they can understand and accept. This research tackles the communication strategy that ought to be integrated into the delivery of advice by a robot. Our finding that robot helpers created a positive impression with hedges and discourse markers suggests an important step for future HRI research in assistive robotics. This step is to draw on the vast communication literature, and to use this literature to suggest key opportunities in modeling human interaction and testing its effectiveness. Maybe, uhm, it will help robots give help.

Acknowledgments

We thank Jeff Hancock, Anind Dey, and Jodi Forlizzi for their suggestions, all delivered with hedges. We thank Nina Demeter, Sara Trapnell, Arya Arabshahi, Evan Tyler Barron, Noah Plomgren, and Daniel Scott Weschler for their assistance creating the videos. This work was supported by National Science Foundation grants 0333420, 0121426, and 0624275.

References

[1] Andersen, G. (2000). The role of the pragmatic marker like in utterance interpretation. In G. Andersen & T. Fretheim (Eds.), *Pragmatic Markers and Propositional Attitude*. Philadelphia: John Benjamins.

[2] Behan, J., & O'Keefe, D. (2008). The development of an autonomous service robot. Implementation: "Lucas" the library assistant robot. *Intelligent Service Robots, 1*, 73-89.

[3] Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage*. New York: Cambridge University Press.

[4] Burgoon, J. (1978). A Communication Model of Personal Space Violations: Explication and an Initial Test. *Human Communication Research, 4*, 129-142.

[5] Burgoon, J., Le Poire, B. A., & Rosenthal, R. (1995). Effects of Preinteraction Expectancies and Target Communication on Perceiver Reciprocity and Compensation in Dyadic Interaction. *Journal of Experimental Social Psychology, 31*, 287-321.

[6] Cesta, A., Cortellessa, G., Pecora, F., & Rasconi, R. (2007). Supporting interaction in the RobotCare Intelligent Assistive Environment. *Proceedings of AAAI* (www.aaai.org).

[7] Coenen, L. H., Hedebouw, L., & Semin, G. R. (2006). *Measuring Language Abstraction: The Linguistic Category Model Manual*.

[8] Fischer, K. (2006a). *Approaches to Discourse Particles:* Emerald Group Publishing.

[9] Fox Tree, J. E., & Schrock, J. C. (2002). Basic meanings of you know and I mean. *Journal of Pragmatics, 34,* 727-747.

[10] Fraser, B. (1980). Conversational mitigation. *Journal of Pragmatics, 4*, 341-350.

[11] Gharpure, C., & Kulyukin, V. (2008). Robot-assisted shopping for the blind: issues in spatial cognition and product selection. *Intelligent Service Robots, 1*.

[12] Goffman, E. (1955). On face-work: An analysis of ritual elements in social interaction. *Psychiatry, 19*, 213-231.

[13] Goffman, E. (1967). *Interaction Ritual: Essays on Face-to-Face Behavior*. New York: Doubleday.

[14] Hayaski, K., Sakamoto, D., Kanda, T., Shiomi, M., Koizumi, S., Ishiguro, H., et al. (2007). Humanoid robots as a passive-social medium: a field experiment at a train station. *Proc. HRI 2007.*

[15] Holmes, J. (1986). Functions of you know in women's and men's speech. *Language in Society, 15*, 1-22.

[16] Jucker, A. H., & Smith, S. W. (1998). And people just you know like wow: Discourse markers as negotiating strategies. In A. H. Jucker & Y. Ziv (Eds.), *Discourse Markers: Descriptions and Theory* (pp. 171-201). Philadelphia: John Benjamins.

[17] Kang, K. I., & Mataric, M. (2005). A hands-off physical therapy assistance robot for cardiac patients. *Proc. IEEE International Conference on Rehabilitation Robotics, ICORR '05*.

[18] Kidd, C. D., & Breazeal, C. (2007). A robotic weight loss coach. *Proceedings of AAAI 2007*.

[19] Kulyukin, V. A. (2006). On natural language dialogue with assistive robots. *Proc. HRI '06.*

[20] Lee, M. K., Kiesler, S., Forlizzi, J., Srinivasa, S., Rybski, P. (2010). Gracefully mitigating breakdowns in robotic services. *Proceedings of HRI '10.*

[21] Lee, S., Kiesler, S., Lau, I. Y., & Chiu, C. (2005). Human mental models of humanoid robots. *Proc. IEEE International Conference on Robotics and Automation, ICRA 2005*.

[22] Lemaignan, S., Ros, R., & Alami, R. (2011). A symbolic approach to perspective-aware grounding, clarification and reasoning for robot. *Robotics, Science, and Systems, Grounding Human-Robot Dialog for Spatial Tasks workshop*. Downloaded from: http://hal.archives-ouvertes.fr/hal-00667030/

[23] Petrick, R.P.A., & Foster, M. E. (July 2012). What would you like to drink? Recognising and planning with social states in a robot bartender domain. *Proceedings of the International Cognitive Robotics Workshop (CogRob 2012)*, 69-76,

[24] Powers, A., Kramer, A., Lim, S., Kuo, J., Lee, S., & Kiesler, S. (2005). Eliciting information from people with a gendered humanoid robot. *Proceedings of RO-MAN 2005*.

[25] Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. New York: Cambridge University Press.

[26] Schiffrin, D. (1987). *Discourse Markers*. Cambridge: Cambridge University Press.

[27] Schober, M., & Brennan, S. E. (2003). Processes of interactive spoken discourse: The role of the partner. In A. Graesser, M. Gernsbacher & S. Goldman (Eds.), *The Handbook of Discourse Processes* (pp. 123-164). Mahwah, NJ: Lawrence Erlbaum.

[28] Schourup, L. (1999). Discourse markers: tutorial overview. *Lingua, 107,* 227-265.

[29] Siegel, M. E. A. (2002). Like: The discourse particle and semantics. *Journal of Semantics, 19,* 35-71.

[30] Stubbe, M., & Holmes, J. (1995). You know, eh and other exasperating expressions: an analysis of social and stylistic variation in the use of pragmatic devices in a sample of New Zealand English. *Language & Communication, 15*, 63-88.

[31] Tapus, A., Mataric, M. J., & Scassellati, B. (March 2007). Socially assistive robotics. *IEEE Robotics and Automation Magazine*, 35-42.

[32] Torrey, C. (2009). *How Robots Can Help: Communication Strategies that Improve Social Outcomes*. Unpublished doctoral dissertation, Carnegie Mellon University.

[33] Torrey, C., Powers, A., Fussell, S. R., & Kiesler, S. (2007). Exploring adaptive dialogue based on a robot's awareness of human gaze and task progress. *Proc. HRI 2007* (pp. 247-254).

[34] Torrey, C., Powers, A., Marge, M., Fussell, S. R., & Kiesler, S. (2006). Effects of adaptive robot dialogue in information exchange and social relations. *Proc. HRI 2006* (pp. 126-133).

[35] Walters, M. L., Dautenhahn, K., Woods, S. N., & Koay, K. L. (2007). Robotic etiquette: Results from user studies involving a fetch and carry task. *Proceedings of HRI 2007.*

[36] Yamaoka, F., Kanda, T., Ishiguro, Hagita, N. (2008). How close?: Model of proximity control for information-presenting robots. *Proceedings of HRI'08.*