Rhetorical Robots: Making Robots More Effective Speakers Using Linguistic Cues of Expertise

Sean Andrist, Erin Spannan, Bilge Mutlu
Department of Computer Sciences, University of Wisconsin–Madison
1210 West Dayton Street, Madison, WI 53706 USA
E-mail: {sandrist, espannan, bilge}@cs.wisc.edu

***Abstract*—Robots hold great promise as informational assistants such as museum guides, information booth attendants, concierges, shopkeepers, and more. In such positions, people will expect them to be experts on their area of specialty. Not only will robots need to be experts, but they will also need to communicate their expertise effectively in order to raise trust and compliance with the information that they provide. This paper draws upon literature in psychology and linguistics to examine cues in speech that would help robots not only to provide expert knowledge, but also to deliver this knowledge effectively. To test the effectiveness of these cues, we conducted an experiment in which participants created a plan to tour a fictional city based on suggestions by two robots. We manipulated the landmark descriptions along two dimensions of expertise: practical knowledge and rhetorical ability. We then measured which locations the participants chose to include in the tour based on their descriptions. Our results showed that participants were strongly influenced by both practical knowledge and rhetorical ability; they included more landmarks described using expert linguistic cues than those described using simple facts. Even when the overall level of practical knowledge was high, an increase in rhetorical ability resulted in significant improvements. These results have implications for the development of effective dialogue strategies for informational robots.**

I. Introduction

Social robots are envisioned to use human language to assist people with informational needs in places such as museums, shops, and information booths. Robots are well suited to such roles, because they afford more intuitive interactions due to their embodiment and ability to use human communicative mechanisms. As embodied social agents, they have the potential to go beyond simply providing information to actually persuading listeners to make certain choices—much like an effective human speaker would do. In order to achieve this goal, they must convince their users that they are *experts* in the subject matter under consideration. For example, most people are unlikely to be persuaded to visit a specific art exhibit suggested by a robot museum guide who does not express sufficient understanding of the museum. Thus, interaction designers must ensure that the behavior of informational robots, specifically their speech, appropriately conveys expertise.

Expertise is commonly thought of as being synonymous with simply having a high degree of knowledge and experience in a particular area. However, there is another dimension of expertise: rhetorical ability, defined as the ability to communicate effectively and persuasively [10]. Someone who is knowledgeable may be considered an expert, but their expertise might not benefit others without the ability to effectively communicate it to them. Such a person might be more correctly referred to as *knowledgeable* rather than *expert*. To make the leap from knowledgeable to expert requires rhetorical effort [10]. Expert human communicators display certain linguistic cues which serve to convince listeners that they are fully qualified experts. In addition to the practical knowledge of the information being presented, many speakers will use expressions of goodwill, references to past instances of sharing their expertise, intuitive organization of information, and appropriate metaphors to convince listeners of their expertise [10].

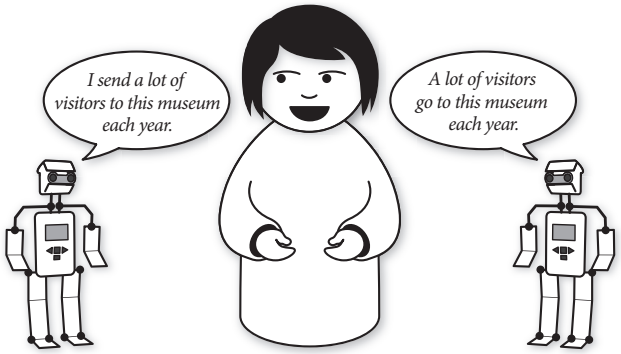

Fig. 1. A piece of information presented by robots with high (left) and low (right) rhetorical ability. They are both conveying the same information, but the robot on the left is using an expression of prior experience to increase its credibility.

If robots are to be perceived as experts, they will not only need to possess a sufficient amount of practical knowledge, they will also need the rhetorical ability to be able to effectively express this knowledge. In this paper, we present a model of expertise, including both dimensions of practical knowledge and rhetorical ability, that can guide the design of effective and persuasive dialogue for robots. We draw upon literature in psychology and linguistics to present speech cues that can increase a speaker's rhetorical ability. We also conducted an experiment in which robots made suggestions to participants on which landmarks to visit on a tour of a virtual city (Figure 1) to evaluate whether these cues help robots achieve the rhetorical ability that the model suggests.

In the next section, we review the literature from which we drew the expert rhetorical strategies most applicable to

robots. In the remainder of the paper, we present our model of linguistic expertise, describe the experiment we designed to evaluate this model, present our results, and discuss the implications of our findings for the design of expert speech for robots.

II. Background

In order to design expert speech cues for robots, we first examined expert speech in humans. In this section, we first present an overview of the relevant literature from psychology and linguistics on the rhetorical strategies used by humans with high rhetorical ability that inform our model of expert speech. Next, we review related work with robots, both as experts and in informational roles where they may have the opportunity to influence people.

A. Humans as Experts

Research in psychology suggests that, in situations that require a certain degree of problem-solving, experts and novices display a number of easily identifiable differences; novices are more dependent on surface features, whereas experts focus more on the underlying principles [4]. For example, when a novice looks at the map of a city they have never visited, they might only know what the map tells them: street names, landmark names, and anything else that is visible on the map itself. A resident of the city (i.e., an "expert") is able to draw from information beyond the map, such as landmarks that are not listed [14]. Research also suggests that experts can find solutions to problems significantly faster than novices do [17].

However, there is an important distinction between *having* knowledge and being able to *share* it. Communicating expertise presents an entirely different challenge from only utilizing prior knowledge. Research suggests that knowledge should be presented in a useful way to someone who may have no prior experience with the topic at hand, using linguistic cues to support the perception of expertise [16]. This rhetorical effort helps speakers move from being highly knowledgeable to being received as an expert; by using rhetoric to increase perceptions of expertise, the speaker—going beyond simply giving information—convinces the listener that the information is important and useful and that they should accept it as true [10].

Research in linguistics has argued that the process of communicating expertise is a careful combination of imparting *knowledge*, proving *credibility*, and building *trust* [10]. Credibility, in particular, can be a very important factor when a speaker is attempting to persuade a listener. Fogg and Tseng [6] argue that credibility is purely a perceived quality and not inherent in any person or piece of information. When listeners regards a speaker as credible, they not only believe that the speaker knows important information, but also that the speaker is being honest. A lack of the perception of credibility, on the other hand, might undermine the entire interaction.

Ultimately, one of the goals of conveying expertise is persuasion. By demonstrating expertise using rhetorical ability, speakers attempt to establish power in the situation, justifying why a listener should value their words. This use of rhetoric can then lead to the listener's compliance, both in thought and behavior [19]. Alternative forms of rhetoric exist, such as invitational rhetoric, which focuses more on offering perspectives rather than a direct attempt to effect change in a listener through persuasion [7]. However, this paper focuses on designing dialogue strategies for robots to utilize the more traditional forms of persuasive rhetoric.

B. Robots as Experts

Research in human-robot interaction (HRI) has also explored how robots might communicate more effectively in informational roles. Torrey et al. [25] demonstrated that the amount of information that a robot might present to its user depends on the level of expertise of the user and the context of the information exchange. Including too many details could be insulting to someone with past experience on a subject, while not including sufficient detail could be unhelpful for someone who is entirely unfamiliar with it. Some information, while helpful, may be restricted only to certain groups, such as health information at a clinic. Another study by Torrey [24] on linguistic cues for help-giving robots showed that not only should direct commands be avoided, as they generally created a poor impression of the robot, but that the use of hedges ("I guess", "I think", "maybe", etc.) softened the intensity of commands and was found to be polite by many of the participants. These findings have strong implications for designing effective dialogue strategies for robots.

Researchers have explored the use of robots as persuasive entities [15], [18], [21]. For instance, Kidd and Breazeal [15] showed that having a physical "expert" entity present during a long-term weight-loss program is beneficial, as people felt more accountable for following the program. Roubroeks et al. [21] explored participants' reactions to robots either opposing or agreeing with them when identifying the most important setting in a washing machine and concluded that messages should be carefully worded to avoid creating a poor impression of the robot, as this could result in people no longer wishing to interact with it.

Studies of robots as persuasive entities have also explored the role of a robot's nonverbal behaviors on its persuasiveness. A recent study showed that participants complied with a robot's suggestions significantly more when it used nonverbal cues associated with persuasiveness such as gaze, gestures, and proximity than when it did not use these cues [5]. Another study demonstrated that the combination of persuasive gaze and gesture cues were more effective than the use of only one of these cues [9]. Finally, Siegel et al. demonstrated that the perceived gender of the robot played an important role in its persuasiveness [22]; in a study that involved museum guide robots attempting to solicit museum donations, they showed that men were more likely to donate money to a female robot, while women did not show a preference. Participants also rated opposite-sex robots as being more engaging, credible, and trustworthy.

Although previous work has examined dialogue strategies for informational robots and robots as persuasive entities, little work has been done on identifying specific verbal strategies to increase perceptions of expertise for robots. In this paper, we seek to address this knowledge gap by developing a model of expert speech for robots, including specific linguistic cues, based on a recent analysis of expert rhetoric for human speakers [10]. The next section introduces this model and the specific cues that robots might use to more effectively communicate their expertise.

III. A Model of Expert Speech

To guide the design of robots that can be perceived as experts, we developed a model of expert speech. The model draws from recent work that defines the overall concept of expertise in two dimensions: *practical knowledge*, which captures prior knowledge and experience, and *rhetorical ability*, which refers mainly to speaking prowess [10]. As represented in Figure 2, we propose that these two dimensions can be considered as dividing the entire space of expertise into four quadrants. A speaker with low practical knowledge and low rhetorical ability is what we refer to as a *true novice*. On the other hand, an individual with a novice-level understanding of the subject matter who uses expert rhetoric in conversation may be falsely perceived as an expert. This situation corresponds to the *perceived expert*. A *perceived novice* describes a speaker who has a high amount of practical knowledge but does not possess the rhetorical skills to communicate that knowledge effectively. Finally, the *true expert* is an individual who possesses both the knowledge and the rhetorical prowess to be effectively persuasive when conveying information.

In this work, we explore how robot speech might be improved in the dimension of rhetorical ability. Research literature in psychology and linguistics points to a number of linguistic cues that can contribute to the speaker's rhetorical ability, in turn increasing his or her effectiveness in persuading listeners [3], [8], [10], [13], [20], [23]. We have identified five of these linguistic cues as being important and useful for robots to convey expertise. A speaker with high rhetorical ability, either a *perceived expert* or a *true expert*, will use many of these linguistic cues throughout his or her speech. An individual with low rhetorical ability, either a *perceived novice* or *true novice*, will use few to none of these cues. Informational robots as social entities will need to use these cues effectively to be perceived as experts. These cues are summarized with examples in Table I and are discussed below.

Goodwill: Effective and persuasive speech involves more than providing information; research suggests that an effective speaker also conveys to the listener that considering this information is in his or her best interests [13]. To achieve this effect, the speaker displays expressions of *goodwill*, which indicate that the speaker wants what is best for the listener [10]. For example, when describing a museum of history, the speaker might not only point out that the museum is a popular landmark, but also suggest that the listeners might enjoy themselves and learn about history if they visit the museum.

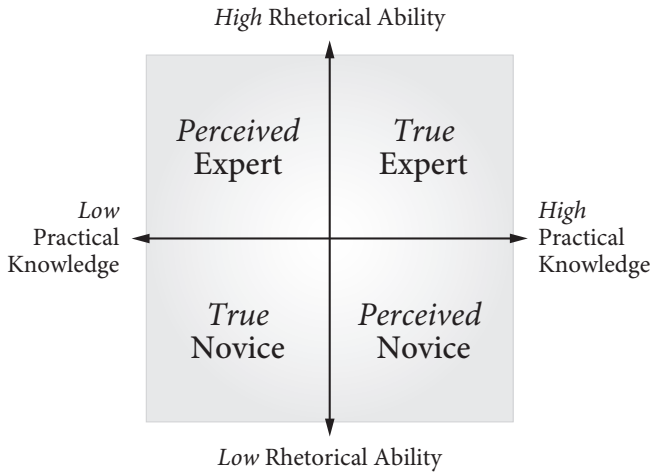

Fig. 2. The overall space of expertise can be split into two dimensions: that of *practical knowledge*, comprised of prior knowledge and experience, and *rhetorical ability*, which refers mainly to speaking prowess. Our model uses these axes to divide the space into four quadrants of expertise.

Prior Experience: Another linguistic cue that research on expertise has identified as key to expert speech is in expressions of *prior experience* [23]. This research suggests that effective speakers also seek to gain their listeners' trust by building multiple favorable experiences with the listener or conveying that the speaker has had prior experience giving to good information to previous listeners. When listeners trust the speaker, they consider the information to be more valuable and rely more on it in making decisions. For example, the speaker might point to a previous listener who visited the history museum and his or her favorable experience at the museum.

Metaphors: Another rhetorical strategy for effective speech is the use of *metaphors*, which helps establish common ground between the speaker and the listener and indicates that the speaker is making an effort to connect with the listener and share his or her experience and expertise [10]. For example, the speaker might depict visiting a particular collection in the museum of history as stepping into the lives of people who lived in the past century. Metaphors are particularly effective linguistic elements to make concepts understandable for listeners with little experience with the subject.

Organization: Rhetorical ability is also shaped by the *organization* of utterances [8], [20]. Research on how novice tutors—individuals who are topic experts but novices in tutoring—teach found that they often took tangents or looped back to an earlier point in the conversation to bring up a detail that they had missed the first time [8]. These tangents and disruptions are examples of poor organization and can damage the credibility of the speaker by creating the impression that he or she is not well versed in the subject [20]. For example, when describing the museum of history, the speaker must display a logical progression of information with sufficient detail and good flow.

TABLE I
THE FIVE LINGUISTIC CUES THAT WE IDENTIFIED FROM THE LITERATURE TO BE KEY FOR INCREASING RHETORICAL ABILITY.

Expertise Cues	Example Utterances
Goodwill *Wanting the best for the listener.*	Expert: *"This cafe is a great place to go for lunch to get out of the hot sun."* Novice: *"This cafe is a great place to go for lunch."*
Prior expertise *References to past helping experience.*	Expert: *"I send a lot of visitors to this museum each year."* Novice: *"A lot of visitors go to this museum each year."*
Organization *More natural organization of information.*	Expert: *"At 1000 years old, the castle is the the oldest landmark in the city. It has Gothic architecture."* Novice: *"The castle is 1000 years old. It has Gothic architecture. It's the oldest landmark in the city."*
Metaphors *Making descriptions more accessible.*	Expert: *"Stepping onto the sunny beach is like wrapping yourself in a towel from the dryer."* Novice: *"The sunny beach is quite hot."*
Fluency *Reduced pauses and confidence in speech.*	Expert: *"The statue is 200 years old. [A 300 ms. pause] It was built to honor the King."* Novice: *"The statue is 200 years old. [A 1200 ms. pause] It was built to honor the King."*

Fluency: Finally, the timing or *fluency* of speech is a key para-verbal cue for rhetorical ability. Research suggests that speakers with both practical knowledge and rhetorical ability can draw upon their knowledge and speak more quickly than speakers who are not comfortable speaking or are not experts on the topic of the conversation [4]. While a certain amount of pausing is expected in any spontaneous speech at any level of the speaker's expertise, pauses that are beyond a "normal" length are perceived as cues for an inarticulate speaker [3] and as disruptive to the flow of the conversation [1]. Brief (less than 200 ms) and medium (200-1000 ms) pauses are frequent and expected in spontaneous speech, while the frequent occurrence of long pauses (more than 1000 ms) might become disruptive.

IV. HYPOTHESES

The model and the linguistic cues described in the previous section will guide the design of robots for roles in which they are expected to provide suggestions and information to people with the goal of persuading them to take a particular course of action. We have developed the following hypotheses on the impact that a robot following the proposed model and set of linguistic cues to form its speech might have in human-robot interaction.

Hypothesis 1: Participants will follow suggestions made by a robot that possesses a high amount of practical knowledge more frequently than suggestions made by a robot with a low amount of knowledge.

This prediction is based on research that suggests that people are more likely to accept advice from advisors with high practical knowledge than from those with low practical knowledge [11].

Hypothesis 2: Participants will follow suggestions made by a robot with high rhetorical ability (using more linguistic cues of expertise) more frequently than suggestions made by a robot with low rhetorical ability.

This prediction is based on findings from studies that show that the use of rhetoric to convey expertise elicits listener compliance [19]. Previous work has also shown that verbal cues of expertise help build trust and that advice-takers tend to follow the suggestions of those they trust [23]. Research in HRI also provides basis for this hypothesis; the use of expert language might shape how helpful a robot is perceived [24].

Hypothesis 3: Within any level of practical knowledge, displaying high rhetorical ability will improve the participants' compliance with the robot's suggestions.

In addition to Hypotheses 1 and 2, we also posit that rhetorical ability will increase the credibility of the speaker and thus the listener's compliance with the speaker's suggestions at any level of practical knowledge.

V. EVALUATION

To test our model of expertise, we designed an experiment in which two Lego Mindstorm robots made suggestions to participants on which landmarks to visit on a tour of a virtual city (Figure 3). The robots differed either in terms of their amount of practical knowledge, their rhetorical ability, or both.

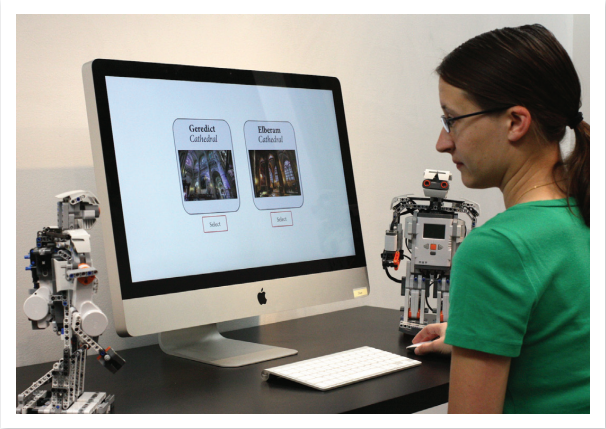

Fig. 3. An experimenter demonstrating the experimental task. The robots made suggestions to participants on which landmarks to visit on a tour of a virtual city. At each decision point, the participant listens to each of the robot's descriptions, then chooses which landmark she would like to include in a virtual tour of the city.

A. Participants

Forty-eight participants (24 females and 24 males) were recruited through postings in physical and online bulletin boards from the University of Wisconsin–Madison campus. All participants were native English speakers, and their ages ranged from 18 to 53 ($M = 23.69, SD = 7.83$). On average, participants rated their familiarity with robots as 2.77 ($SD = 1.55$) on a seven-point rating scale.

B. Study Design

We manipulated the level of practical knowledge (low vs. high) and the level of rhetorical ability (low vs. high) for each robot to create four types of speech strategies that correspond to the quadrants of the model described in Section III and illustrated in Figure 2: perceived expert, true expert, perceived novice, and true novice. To compare these strategies, each trial involved pairs of robots that followed a different speech strategy. We tested every possible pairing of quadrants and randomly assigned eight participants to each of the following pairing conditions: (1) perceived expert vs. true expert, (2) true novice vs. perceived novice, (3) perceived expert vs. true novice, (4) true expert vs. perceived novice, (5) perceived expert vs. perceived novice, and (6) true expert vs. true novice.

For each location, we developed four scripts for the four quadrants of the model. Practical knowledge was manipulated by varying the number of discrete facts that the descriptions of the landmarks presented; high knowledge scripts contained four discrete pieces of information, while low knowledge scripts contained two. Rhetorical ability was manipulated by varying the number of linguistic cues of expertise present in the robot's speech; high rhetorical ability scripts contained three of the expert speech strategies from Table I, while low rhetorical ability scripts contained none. Because human speakers employ only a limited number of cues in a given message, we used only three cues in each script but used all five cues in approximately equal proportions across all scripts in the study.

The examples below illustrate the description of the landmark Elberam Cathedral in all four linguistic strategies. Four facts are associated with this landmark: (1) it is visible from anywhere in the city; (2) it has high towers; (3) its towers have spiral stairways; and (4) the stairways have more than 350 steps.

> *True Expert*: "Elberam Cathedral is visible from anywhere in the city. Its towers are as tall as the clouds on some days, and include spiral stairways, which have more than 350 steps."

The script presents all four facts, uses good fluency and organization, and includes a metaphor to describe the height of the towers.

> *Perceived Novice*: "Elberam Cathedral is visible from anywhere in the city. [*1500 ms pause*] The cathedral has high towers. [*1500 ms pause*] The towers include spiral stairways which have more than 350 steps."

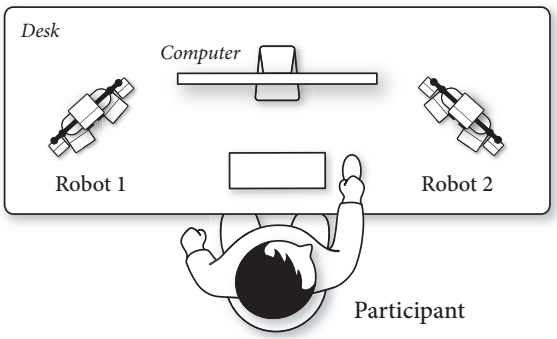

Fig. 4. A diagram of the physical setup of the experiment. One at a time, each robot would turn its head towards the participant to provide information about a landmark on the screen, then turn its head away.

The script presents all four facts but displays low rhetorical ability with no linguistic cues of expertise.

> *Perceived Expert*: "Visible from anywhere in the city, Elberam Cathedral has towers that are as tall as the clouds on some days."

The script presents only two of the facts, but uses good fluency and organization, and includes a metaphor to describe the height of the towers.

> *True Novice*: "Elberam Cathedral is visible from anywhere in the city. [*1500 ms pause*] The cathedral has high towers."

The script presents only two of the facts and displays low rhetorical ability with no linguistic cues of expertise.

To assess whether these manipulations achieved different levels of practical knowledge and rhetorical ability, we asked two independent coders blind to the manipulations to rate each script for information content and rhetorical expertise and analyze their ratings using a repeated measures analysis of variance (ANOVA), using coder and script IDs as covariates. Practical knowledge significantly affected the ratings of information content, $F(1,4) = 12.50$, $p = .024$, but only marginally affected ratings of rhetorical expertise, $F(1,4) = 7.21$, $p = .055$. Rhetorical ability significantly affected ratings of rhetorical expertise, $F(1,4) = 23.34$, $p = .009$, but did not affect ratings of information content, $F(1,4) = 0.03$, $p = .87$. These results suggest that the scripts match our target levels of practical knowledge and rhetorical ability.

C. Procedure

Following informed consent, the experimenter introduced the participants to the two Lego Mindstorm robots, Sam and Ky (chosen as gender-neutral names), and told them that the robots were being trained as tour guides. The robots were placed on the two sides of a computer monitor as shown in Figure 4. In the experimental task, participants planned a virtual tour of a fictional city with options presented by the two robots. They were shown ten pairs of similar locations, such as two art museums or two amusement parks. Depending on the pairing condition, one robot described locations using

TABLE II
THE SUBJECTIVE MEASURES USED TO CAPTURE THE PARTICIPANTS' IMPRESSIONS OF THE ROBOTS.

Measure *(Reliability)*	**Items**
Competency *(Cronbach's α = 0.91)*	*The robot was informative.* *The robot made travel recommendations like an expert.* *The robot was knowledgeable on the city landmarks.* *The robot had a lot of expertise on the travel destination.* *The robot seemed to have deep knowledge of the city and its landmarks.*
Persuasiveness *(Cronbach's α = 0.88)*	*The robot's suggestions convinced me to choose its locations.* *The robot was like a good salesperson.* *I would use the robot's suggestions in the future.*
Sociability *(Cronbach's α = 0.86)*	*The robot was likeable.* *The robot was sociable.* *The robot was friendly.*
Trustworthiness *(Cronbach's α = 0.74)*	*The robot was believable.* *The robot was trustworthy.* *I trust the information that the robot gave me.*

the script from one quadrant of the model, while the other robot used a different quadrant. Each robot, in random order, turned its head to the participant, provided the information, and then turned away. After each robot gave its description, the participant chose a location to visit using an experimental interface as illustrated in Figure 5. All scripts were recorded using two modulated voices designed to be gender-neutral, which were assigned randomly to each robot at the beginning of the tour. The order of the landmark pairs and which landmark appeared to the left or right of the screen were also randomized. After finishing the tour of ten locations, participants filled out a questionnaire that measured their perceptions of the two robots. The study took approximately thirty minutes, and participants were paid $5.

D. Measurement & Analysis

The study involved two manipulated independent variables: practical knowledge and rhetorical ability. The primary dependent variable was participant compliance, which was measured by labeling each choice of landmark with the speech strategy employed by the robot that provided the description for that landmark. We also collected data on a number of subjective measures of participants' impressions of the robots as secondary dependent variables. Seven-point rating scales were used for all items in these measures. Item reliabilities, measured by Cronbach's α, for all measures were sufficiently high. Table II summarizes these measures.

Our data analysis included mixed-model analysis of variance (ANOVA) tests to assess how the robots' level of practical knowledge and rhetorical ability affected participants' decisions of which landmarks to visit and their subjective evaluations of the robots. Practical knowledge, rhetorical ability, and the interaction between the two variables were modeled as fixed effects. Additionally, the analysis of data from objective measures considered Trial ID as a random effect. All contrast tests used Scheffé's method to control for family-wise Type 1 error rate.

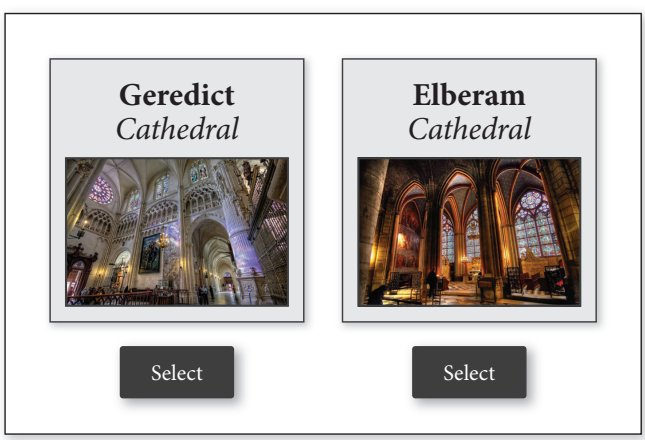

Fig. 5. A schematic of the interface that the participants used to choose which landmark to include on a virtual tour. Each landmark was described by a different robot, differing in the amount of practical knowledge and/or rhetorical ability displayed.

VI. RESULTS

A. Results on Objective Measures

Hypothesis 1 predicted that participants would more frequently comply with suggestions by the robot with higher practical knowledge. Our analysis also supported this hypothesis; participants chose landmarks that were described with all possible information more frequently than they chose landmarks that were described with only some information, 66% vs. 34%, $F(1,956) = 109.06$, $p < .001$. There was no significant interaction between levels of practical knowledge and rhetorical ability, $F(1,956) = 0.68$, $p = .41$.

Hypothesis 2 predicted that participants would chose more landmarks described by the robot with higher rhetorical ability. Our analysis provides support for this hypothesis, as the robot's level of rhetorical ability had a significant effect on which locations participants chose to visit, $F(1,956) = 21.83$, $p < .001$. The participants chose to visit landmarks described by robots using high rhetorical ability 57% of the time, compared to 43% for landmarks described by the robots using low rhetorical ability.

Our data also provided support for Hypothesis 3, which predicted that high rhetorical ability would improve participant compliance at any level of practical knowledge. Contrast tests provided support for this prediction; high rhetorical ability significantly improved participant compliance when practical knowledge was low, $F(1,956) = 15.11$, $p < .001$, and when practical knowledge was high, $F(1,956) = 7.40$, $p = .007$.

B. Results on Subjective Measures

The analysis of data from our subjective measures showed that participants rated robots with high practical knowledge to be more competent, $F(1,91) = 49.68$, $p < .001$, more persuasive, $F(1,91) = 28.65$, $p < .001$, more trustworthy, $F(1,91) = 13.07$, $p < .001$, and marginally more sociable, $F(1,91) = 3.61$, $p = .061$, than they rated robots with low practical knowledge. Robots with high rhetorical ability were rated more competent, $F(1,91) = 4.66$, $p = .034$, more persuasive, $F(1,91) = 19.55$, $p < .001$, and more sociable, $F(1,91) = 27.56$, $p < .001$, than

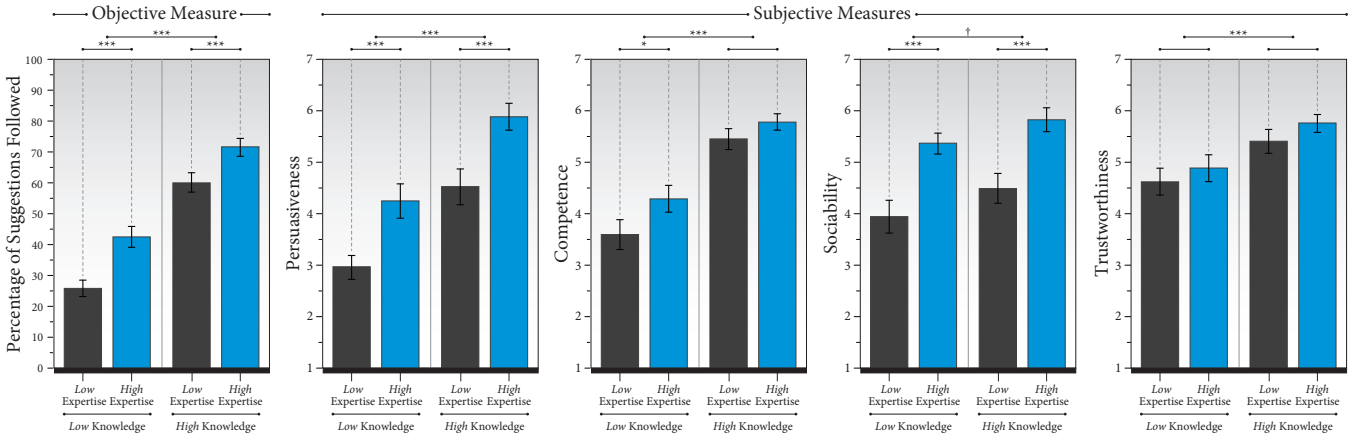

Fig. 6. The results from our objective and subjective measures. The use of rhetorical cues of expertise affected participant compliance, as well as their perceptions of the robot's competence, persuasiveness, and sociability. The robot's level of knowledge also significantly affected compliance, as well as perceptions of competence, persuasiveness, and trustworthiness. ***, **, *, and † indicates p-values smaller than .001, .01, .05, and .10, respectively.

robots with low rhetorical ability were. Rhetorical ability had no effect on participants' ratings of the robot's trustworthiness, $F(1,91) = 1.83$, $p = .18$. The analysis showed no interaction effects between practical knowledge and rhetorical ability in measures of competency, $F(1,91) = .60$, $p = .44$, persuasiveness, $F(1,91) = 0.017$, $p = .90$, sociability, $F(1,91) = 0.003$, $p = .95$, or trustworthiness, $F(1,91) = 0.012$, $p = .91$.

Contrast tests revealed interesting findings in measures of competency and persuasiveness. Rhetorical ability improved perceptions of competency when practical knowledge was low, $F(1,91) = 4.30$, $p = .041$, and not when it was high, $F(1,91) = .96$, $p = .33$. Also, high practical knowledge with low rhetorical ability was not significantly different from low knowledge with high rhetorical ability in ratings of persuasiveness, $F(1,91) = 0.43$, $p = .51$, indicating that rhetorical ability can make up for a lack of knowledge in the perceived persuasiveness of a robot. All of our results are visualized in Figure 6.

VII. Discussion

Our results showed that the rhetorical cues of expertise we presented in our model significantly influenced participants' decisions. As predicted by Hypothesis 1, participants more frequently followed suggestions made by a robot with high practical knowledge than those made by a robot with low practical knowledge. Participants were also more likely to follow suggestions made by a robot that showed high rhetorical ability than suggestions made by a robot that did not show this ability, providing support for Hypothesis 2. Regardless of the robot's level of practical knowledge, rhetorical ability significantly improved participant compliance, providing support for Hypothesis 3. Participants also rated the robot with high rhetorical ability to be more competent, persuasive, and sociable. We found that rhetorical ability could make up for a lack of practical knowledge in ratings of persuasiveness. Interestingly, while practical knowledge improved ratings of trustworthiness, rhetorical ability did not.

A. Design Implications

Robot designers may wish to design robots that maximize both dimensions of expertise, both in practical knowledge and rhetorical ability. However, we have shown that even when robots might not possess a high level of practical knowledge, rhetorical ability can be used to improve compliance and perceptions of persuasiveness, competence, sociability, and trustworthiness of the robot. The linguistic cues we identified can serve as a starting point for those seeking to design effective speech behaviors for robots. However, designers should also exercise some caution in the use of these cues, as rhetorical ability might not be appropriate to use in some situations. For example, when discussing which prescription drug a user should take, sounding like an expert without having a high degree of knowledge about the drug or the user's specific situation might have harmful consequences.

B. Limitations and Future Work

The context in which we evaluated our model of expertise—two robots giving competing advice in a tour-planning setting—might seem unrealistic or non-generalizable, as there would likely not be two competing robots in real-world applications. In a preliminary version of the study, we used a single robot to describe landmarks using different linguistic strategies. The results were similar with participants more frequently choosing landmarks that were described with high rhetorical ability. The current study addresses an important limitation of this preliminary study; the presence of two robots allows participants to make subjective evaluations of robots who use different rhetorical techniques. Our preliminary study also showed that participant gender had no impact on the results, therefore we did not include gender in the analysis of the current study.

While the Lego Mindstorm robots provided all of the affordances we needed for our study, particularly embodiment and shifting attention between the screen and the participant by moving their gaze, whether our results would generalize to other robot designs is unknown. On the other hand, to have

shown that linguistic cues of expertise have strong effects even when they are presented by toy-like robots could be considered a strength of our model.

Because this study was meant to serve as a first step toward a better understanding of how robots might express expertise using linguistic cues, we used multiple cues of expertise without trying to establish their relative effectiveness in achieving expert speech. Future work could further explore which cues are most effective while potentially targeting and removing cues which might be unnatural for robots.

The most obvious and important next step with this work would be to integrate the model into a natural language generator. The model could take the form of a module near the end of a natural language generation pipeline. Given a set of knowledge to be expressed in language, the "expert rhetoric module" could reorganize the language to improve organization and fluency, and automatically generate metaphors and expressions of past experience and goodwill. While some work has been done in developing computational frameworks for handling metaphors in natural language processing [2], [12], automatic generation of metaphors may remain unfeasible in the near future. For now, designers of dialogue systems for robots could handcraft a stable of metaphors in a narrow application domain for the robot to draw from, which would simplify the robot's task to selecting the appropriate metaphor at the appropriate time.

VIII. Conclusion

Social robots hold great promise as informational assistants, providing services in museums, information kiosks, shops, and so on. These positions will require robots to provide their users with credible information that their users value and trust. To do so, robots not only need to be knowledgeable on the topic in which they provide information, but they also need to show the rhetorical ability to communicate their knowledge effectively. In this paper, we drew from literature in psychology and linguistics to construct a model of speech cues that shape the rhetorical ability of a speaker. These cues include expressions of goodwill, references to prior expertise, organization of information, the use of metaphors, and fluency in speech. In a laboratory experiment, participants planned a tour of a fictional city with the advice of two robots who differed in their level of practical knowledge or rhetorical ability. Our results showed that cues for rhetorical ability effectively shaped people's perceptions of the robot's expertise; people were more likely to follow suggestions offered by a robot using linguistic cues of expertise than they were from a robot who simply listed facts. They also found the robot with high rhetorical ability to be more competent, persuasive, and sociable. While future work is needed to further identify which cues are most effective in making robots sound like experts as well as how such cues might be integrated into natural language generation for robots, we have shown that it is possible to improve compliance with the robot's suggestions, as well as how it is perceived, using linguistic cues of expertise. These cues can make the interaction between humans and robots more effective.

Acknowledgments

We would like to thank Flickr users *cuellar* and *marcp_dmoz* for the use of their cathedral photos under a Creative Commons license. We would also like to thank Brandi Hefty for her help with conducting the experiment.

References

[1] E. Campione and J. Véronis. A large-scale multilingual study of silent pause duration. In *Proc. Speech Prosody '02*, pages 199–202, 2002.

[2] J. G. Carbonell. Metaphor: an inescapable phenomenon in natural language comprehension. Technical report, Carnegie Mellon University, Pittsburgh, PA, 1981.

[3] R. Carlson, K. Gustafson, and E. Strangert. Cues for hesitation in speech synthesis. In *Proc. Interspeech '06*, volume 6, 2006.

[4] M. Chi, R. Glaser, and E. Rees. *Expertise in problem solving*. Learning Research and Development Center, University of Pittsburgh, Pittsburgh, PA, 1981.

[5] V. Chidambaram, Y.-H. Chiang, and B. Mutlu. Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues. In *Proc. HRI '12*, HRI '12, pages 293–300, 2012.

[6] B. Fogg and H. Tseng. The elements of computer credibility. In *Proc. CHI '99*, pages 80–87, 1999.

[7] S. K. Foss and C. L. Griffin. Beyond persuasion: A proposal for an invitational rhetoric. *Communications Monographs*, 62(1):2–18, 1995.

[8] M. Glass, J. Kim, M. Evens, J. Michael, and A. Rovick. Novice vs. expert tutors: A comparison of style. In *Proc. MAICS '99*, pages 43–49, 1999.

[9] J. Ham, R. Bokhorst, R. Cuijpers, D. van der Pol, and J.-J. Cabibihan. Making robots persuasive: The influence of combining persuasive strategies (gazing and gestures) by a storytelling robot on its persuasive power. In *Social Robotics*, Lecture Notes in Computer Science, pages 71–83. Springer-Verlag, 2011.

[10] E. J. Hartelius. *The rhetoric of expertise*. PhD thesis, The University of Texas at Austin, 2008.

[11] N. Harvey and I. Fischer. Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, 70(2):117–133, 1997.

[12] J. R. Hobbs. Metaphor, metaphor schemata, and selective inferencing. Technical Report 204, SRI International, Menlo Park, CA, 1979.

[13] M. Hyde. *The ethos of rhetoric*. Univ South Carolina Press, 2004.

[14] E. Isaacs and H. Clark. References in conversation between experts and novices. *J Experimental Psychology: General*, 116(1):26, 1987.

[15] C. Kidd and C. Breazeal. Robots at home: Understanding long-term human-robot interaction. In *Proc. IROS '08*, pages 3230–3235, 2008.

[16] I. Kinchin. Effective teacher student dialogue: a model from biological education. *J Biological Education*, 37(3):110–113, 2003.

[17] J. Larkin, J. McDermott, D. Simon, and H. Simon. Expert and novice performance in solving physics problems. *Science*, 208(4450):1335, 1980.

[18] R. Looije, M. A. Neerincx, and F. Cnossen. Persuasive robotic assistant for health self-management of older adults: Design and evaluation of social behaviors. *Intl J Human-Computer Studies*, 68(6):386–397, 2010.

[19] G. Marwell and D. Schmitt. Dimensions of compliance-gaining behavior: An empirical analysis. *Sociometry*, pages 350–364, 1967.

[20] J. Noel. On the varieties of phronesis. *Educational Philosophy and Theory*, 31(3):273–289, 1999.

[21] M. Roubroeks, J. Ham, and C. Midden. The dominant robot: Threatening robots cause psychological reactance, especially when they have incongruent goals. *Persuasive Technology*, pages 174–184, 2010.

[22] M. Siegel, C. Breazeal, and M. Norton. Persuasive robotics: The influence of robot gender on human behavior. In *Proc. IROS '09*, pages 2563–2568, 2009.

[23] J. Sniezek and L. Van Swol. Trust, confidence, and expertise in a judge-advisor system. *Organizational Behavior and Human Decision Processes*, 84(2):288–307, 2001.

[24] C. Torrey. *How robots can help: Communication strategies that improve social outcomes*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, 2009.

[25] C. Torrey, A. Powers, M. Marge, S. Fussell, and S. Kiesler. Effects of adaptive robot dialogue on information exchange and social relations. In *Proc. HRI '06*, pages 126–133, 2006.