

Eliciting Information from People with a Gendered Humanoid Robot*

Aaron Powers¹, Adam D.I. Kramer², Shirlene Lim¹, Jean Kuo¹,
Sau-lai Lee¹, Sara Kiesler¹

¹ *Carnegie Mellon University*

Human-Computer Interaction Institute

5000 Forbes, Pittsburgh, PA 15232, USA

apowers@cmu.edu, adik@uoregon.edu, sllim@andrew.cmu.edu, cjkuo@andrew.cmu.edu,

sllieh@hkusua.hku.hk, kiesler@cs.cmu.edu

² *University of Oregon*

Department of Psychology

1227 University of Oregon, Eugene, OR 97403, USA

Abstract – A conversational robot can take on different personas that have more or less common ground with users. With more common ground, communication is more efficient. We studied this process experimentally. A “male” or “female” robot queried users about romantic dating norms. We expected users to assume a female robot knows more about dating norms than a male robot. If so, users should describe dating norms efficiently to a female robot but elaborate on these norms to a male robot. Users, especially women discussing norms for women, used more words explaining dating norms to the male robot than to a female robot. We suggest that through simple changes in a robot’s persona, we can elicit different levels of information from users—less if the robot’s goal is efficient speech, more, if the robot’s goal is redundancy, description, explanation, and elaboration.

Index Terms – human-robot interaction, social robots, humanoids, communication, dialogue, common ground, knowledge estimation, mental models, gender

I. INTRODUCTION

When we interact with another person, we are efficient much of the time, using as few words as we need to communicate our meaning [1]. For instance, if I am discussing football with another fan, I do not need to explain the rules of the game. I use more jargon and partial sentences when talking to an expert than I use when I am talking with a novice. Likewise, if I am interacting with a robot I can be efficient if the robot knows a lot about my topic. The more I estimate the robot knows about my topic, the less I need to be redundant, to elaborate, describe, and explain—the more efficient I can be in my communication.

In human-robot interaction, we can use this principle to elicit efficient communication from users. For example, if a mechanic’s helper-robot conveys that it has knowledge of tools, users should feel they have an overlapping store of knowledge with the robot. Because of this common ground, users will be able to query the robot for a wrench without explaining what they mean by “wrench.” More common ground with the robot is associated with the elicitation of less information and more efficiency.

Suppose we do not have an efficiency goal for this interaction. If there is noise in the communication channel, we may want users interacting with the robot to be redundant, to explain themselves, and to elaborate. For example, we may want the user to describe the desired wrench in detail to the robot, perhaps to distinguish different types of wrenches or other tools that the robot may confuse it with. In this case, we will not want the robot to exhibit common ground with the user, but instead to seem more ignorant. Equivalent situations can be found in human conversational settings. For instance, teachers, police interrogators and research interviewers often display confusion or uncertainty to elicit more detailed information from students, suspects, and subjects. In human-robot interaction, a robot might elicit more detailed information from users by conveying less overlapping expertise and less common ground with users. In this paper we examine this idea, showing how a robot’s gendered persona can increase or decrease common ground, thereby influencing the information people convey to the robot.

II. ESTABLISHING COMMON GROUND THROUGH PERSONAS

To perceive common ground with a robot, users need to have a reasonably accurate idea about what the robot knows [2]. An obvious starting point is what users themselves know, or think they know. People’s own knowledge acts as a default or anchor for estimating the knowledge of others [3]. People also use social context cues such as a robot’s appearance or demeanor to build a mental model of what the robot knows. Social cues point to social groups such as gender, age, profession, and nationality, and these social groups convey a persona, that is, a personality with social and intellectual attributes [2]. Personas help people estimate others’ knowledge [4]. For example, if a robot is a humanoid from New York or Hong Kong, people conclude that it has knowledge about these localities [5].

Common ground with a robot is likely to affect how people communicate with the robot and what they expect in return from it. Speakers design messages to be appropriate to what they assume to be the knowledge of the recipients

* This work is supported by NSF Grant #IIS-0121426.

[6]. People represent information sparsely when they are communicating with others with whom they share much knowledge; by contrast, people represent information more elaborately if they have to communicate it to others who know nothing about the subject matter [7, 8]. In figure 1, we summarize this argument. In the figure, social cues influence user's perception of the robot's persona. The persona, and the user's own knowledge, influence the user's estimate of the knowledge of the robot. Comparing the robot's estimated knowledge with the user's own knowledge results in the degree to which the user and the robot share common ground. More common ground results in the user employing more efficient speech; less common ground results in the user employing more elaborated speech.

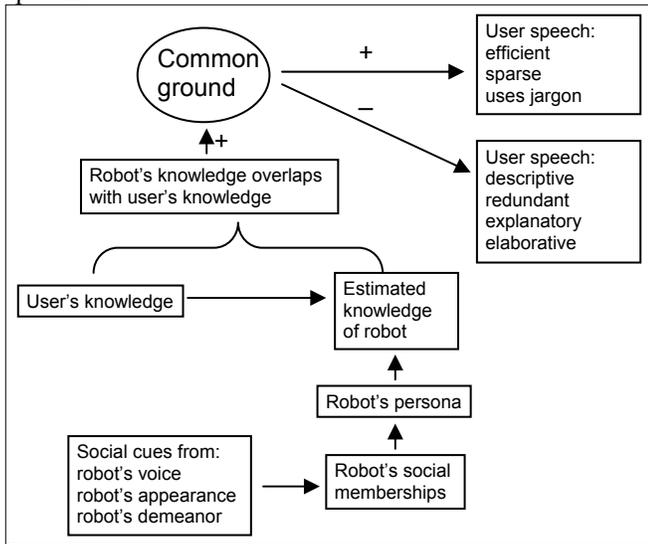


Fig. 1. How common ground predicts user speech with a robot.

To examine the process outlined above, we studied the way a gendered robot conveys common ground and how that common ground influences users' speech. Starting with the social cues depicted at the bottom of Figure 1, we argue that a robot that speaks with a feminine (high frequency) voice or that has feminine appearance will be associated with the female gender, and will take on the persona of a female. As such, the robot will be estimated to have knowledge that many women have, such as knowledge of women's clothing sizes and women's sports celebrities. By contrast, a robot that speaks with a masculine (low frequency) voice and looks masculine will be estimated to have knowledge that many men have, such as knowledge of men's clothing sizes and men's sports celebrities.

III. GENDERED PERSONA PREDICTIONS

In our experiment, we chose to focus on the topic of romantic dating practices. In human populations, women are more knowledgeable about dating norms and social practices, and they have more social skill than men do [9]. Therefore, hypothesis 1 is that subjects (not necessarily consciously) will assume a "female" robot knows more about dating practices and norms than a "male" robot does.

In the common ground model, people use their own knowledge to estimate others' knowledge. (Dawes [10] showed that people do well, statistically speaking, to take their own opinions or knowledge as representative of that of the group to which they belong.) People then compare their knowledge to their estimates of others' knowledge to arrive at an estimate of common ground. Therefore, we argue, women should assume more overlapping knowledge with a female robot and men should assume more overlapping knowledge with a male robot. According to this logic, the most overlapping knowledge and common ground should be between women interacting with a female robot about dating norms that pertain particularly to women.

What does this process predict about how subjects will speak to the robot if the robot asks them about dating norms? In the model, subjects will describe and explain dating norms efficiently to a female robot because the female robot already shares some of this dating knowledge, i.e., has more common ground with the subject. They will explain themselves more to a male robot because of the comparative lack of common ground. Therefore, hypothesis 2 is that if women assume more overlapping knowledge with a female robot and men assume more overlapping knowledge with a male robot, then women should explain dating norms less to the female robot than men do. Likewise, men should explain dating norms less to the male robot than women do. Finally, we have argued that women have the most common ground with a female robot who asks them about dating norms that apply to women. Hence, the most efficient communication should be found in the case where women are speaking with a female robot about dating norms for women.

We tested these predictions in the belief that, if valid, they have significant implications for understanding and designing human-robot social interaction. The theory implies that people who interact with a gendered humanoid robot do not approach the robot *tabula rasa*, but rather develop a default mental model of the robot's knowledge. That knowledge estimate influences their assumed common ground with the robot and their discourse with the robot. Designers can affect these models and the consequent speech of subjects in appropriate directions, eliciting more efficiency or more elaboration depending on their goals for the human-robot interaction.

IV. RELATED WORK

The last decade has seen a number of projects involved in the construction of social robots [11 - 15]. More recently, Valerie, a receptionist-robot, engages in a dialogue with people [16]. Robovie [17] is a child-sized robot in Japan who speaks English with school children. The Nursebot robot, Pearl, the same robot we used in the study described in this paper, was initially developed to interact with older people [18].

In a previous study, we showed that the language and appearance of a robot could change people's estimations of

its knowledge [5]. In that study subjects assumed that a robot made in Hong Kong who spoke Chinese knew more about landmarks in China than a robot made in New York who spoke English. Here we follow up on this previous study, to demonstrate that a robot's persona can be gendered easily using simple cues, and that the robot's gender will influence how much people say to the robot.

V. METHOD

We tested the predictions in an experiment in which young adults of both genders engaged in a one-on-one dialogue with a humanoid robot. We chose the topic of "first dates" because almost all young adults have personal knowledge of dating practices and because there are well-established schemas for behavior of women and men on first dates [19]. Indeed, norms for first dates have changed little since the 1950s [20].

In the experiment, a subject, alone with the robot, engaged in a dialogue with the robot. The robot was presented as either female (feminine voice, pink lips) or male (male voice, grey lips). We used only these two cues on purpose, to demonstrate that differences in robot persona and subject behavior can be accomplished through minimal variation of a robot's appearance and voice.

The male or female robot told the subject that it was training to be a dating counselor, and that it needed advice about what typically happens on dates. The robot then asked the subject various questions about events that transpire on a first date, and the robot responded to what the subject typed. The dialogue was scripted to begin with general questions about dating, such as where people meet others and about the appropriateness of conduct such as dating a boss or co-worker. As the dialogue progressed, the robot talked about a hypothetical couple, "Jill" and "John," who were about to go on a first date. The robot asked the subject a series of questions about Jill and John, such as whether John should call Jill back if she was busy the first time he called, or if Jill should bring John flowers. Subjects' answers to these questions about Jill and John allow us to evaluate how female and male subjects talked with a male or female robot differently depending on whether they were talking about a woman (Jill) or a man (John).

A. Experimental Design

There were two factors that were varied between subjects, robot gender ("male" or "female"), and subject gender (male or female). In addition, each subject spoke with the robot both about Jill and about John.

B. Subjects

Thirty-three native American-English speakers from Carnegie Mellon participated for US\$10 cash as payment (17 males, 16 females; average age 21 years). Half the study was run by a male experimenter and half by a female; there were no differences due to experimenter gender.

C. Procedure

When subjects arrived at the experimental lab, the experimenter told them he/she was creating a dating service for Carnegie Mellon students, and that their conversation would help train the robot's AI system to give people better advice.

Subjects conversed with the robot through an interface like that of Instant Messaging. The IM interface was on the screen on the robot's chest (see Figures 2-3). We used a robot without speech recognition, because we wanted to be able to generalize to the current generation of robots, which have limited speech recognition of highly descriptive dialogue. The robot used Cepstral's Theta for speech synthesis, and its lips moved as it spoke [21]. The text also showed on the screen, as in IM interfaces.



Fig. 2. The robot talking with a subject.

We adapted the robot's questions from Laner & Ventrone's studies of dating norms [19, 20]. They asked students what events typically transpire on a first date. We chose twelve of the events with the largest gender difference, and created a scenario about "John and Jill," two hypothetical individuals who were interested in each other. Six of the items were most commonly thought to be initiated by men (e.g. "decide on plans by yourself", 61% say men, 6% say women), and six were most commonly performed by women (e.g. "buy new clothes for date," 2% say men, 75% say women).

The robot asked some questions about what Jill should do and some questions about what John should do. In each case, the robot asked half of these in a way that supported a gender stereotype (e.g. "Do you think that John should make the plans for the date?"), and half in a way that reversed the stereotype (e.g. "Do you think it's appropriate for John to buy new clothes for a first date?"). The questions about Jill and John were embedded in other questions about dating conduct and norms (such as the wisdom of Internet dating), so as to disguise our interest in gender.

After chatting with the robot, the subject completed a survey, which included ratings of the masculinity and

femininity of the robot, taken from Bem's Sex-Role Inventory [22], as well as other questions about the robot's personality, knowledge, and humanlikeness.

D. Dialogue

The robot interpreted and responded to the subject using a customized variant of the Alice chat-bot [23], a publicly available pattern-matching text processor. We developed and refined the dialogue through pretesting, and carried out two pilot studies (one with an animation of a robot and one with the actual robot).

The first version of the dialogue resulted in very large variability across subjects both in the amount and content of their speech. Some subjects did respond to the robot's questions but asked for clarifying information; others did not answer the questions and had to be prompted to answer. One reason for this variability across people is that, initially, the chat-bot often did not respond well to subjects' questions. For example, when the robot asked "Should [person's name] go to a club?" some subjects asked, "Can he/she dance?" In each successive test, we tailored the chat-bot's responses towards the questions and comments that subjects made, and dropped dialogue that subjects did not understand.

Most questions the robot asked could be answered with "yes," "no," or a number. Hence, it was possible for the subject to be very efficient. However, as we expected from the model in Figure 1, many subjects did not say a simple "yes" or "no" (e.g., see Figure 3). For the robot to understand most replies, we had to compile hundreds of variants of common responses from subjects' responses in the pilot tests, such as "that would be nice," and "of course not." We added "smart" exchanges. For instance, if the subject said, "Absolutely!" the robot remarked on the subject's certainty. We also added an ability to respond to question-specific answers, such as "they should split the check," to allow more comprehension by the robot. When the subject said something vague like "maybe," or "only if ...," or when the subject otherwise failed to answer the question in a manner the robot could understand, the robot prompted the subject, "Please rephrase that," or "Please be more specific," and "tell me whether it would be OK for them to date if John was Jill's boss."

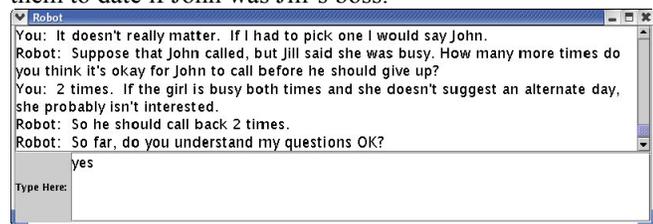


Fig. 3. IM-like chat interface, with responses from a pilot test.

Subjects in the pilot studies commonly made spelling and grammatical errors, and did not correct themselves, such as typing "shoudl," or "No, is Jill and John ike each other and Jil is comfortable with asking him on a date." To fix this problem, we added the Linux Aspell spell checker

to find many spelling errors and automatically correct them in the robot's interpreter. Thus when the subject spelled something wrong, the robot was still able to interpret it, and if the robot repeated the subject's words, the robot spelled the words correctly. As a result of these improvements in the robot's script and interpreter, the number of nonresponses by subjects declined precipitously.

Although branching on the subjects' responses made the interaction feel more fluid, some of the branches were boring or redundant and branches tend to complicate statistical analysis. Therefore, for the main experiment, we shortened the robot's dating question script from our original 1026 words in 65 sentences to 876 words in 50 sentences. We reduced 11 branches to only 3. We made the questions clearer, increasing the average number of words in each question from 16.3 to 17.5.

E. Analyses of Main Experiment

The main dependent variable was how much the subjects said to the robot about what Jill's and John's conduct should be before, on, and after their first date. We used the Text Analysis and Word Counts (TAWC) program [24] to count the number of words the subjects used to answer the robot's questions about Jill and John. To normalize the counts, which were skewed and left censored (a person can't say fewer than zero words to any question), logs of the totals were computed and the data were centered. The result is a standardized measure of the log of total words spoken about Jill and about John.

The data were analyzed using analysis of variance with two between factors (gender of subject and gender of robot) and one within factor (words about Jill compared to words about John).

VI. RESULTS

A. Subjects' Perception of Robot Persona

The first analysis was a check on the manipulation. That is, did the subjects perceive the robot to be gendered? We asked subjects a write-in question whose result was significant (chi square = 40, $p < .0001$). Sixteen of 17 subjects in the female robot condition said the robot was female and one said "female?" In the male robot condition 14 of 16 subjects said the robot was male, one said female, and one said "male?" We also asked subjects to respond to a pair of 5-point rating scales (1 = low, 5 = high) asking how masculine and how feminine the robot was. Subjects rated the female robot an average of 3 on the feminine scale and 2.2 on the masculine scale, and they rated the male robot an average of 2.1 on the feminine scale and 3.6 on the masculine scale (interaction $F [1, 29] = 25, p < .001$).

Then, we tested whether there were differences by robot gender in perceptions of the robot's speech skills. Three rating scales (1 - 5) addressed this question: the robot's speech quality, the robot's response time, and the robot's conversation skill. There were no differences due to robot or subject gender. On average, subjects rated the

robot's speech quality 3.3, response time 2.7, and conversation skill 2.8. These scores are lower than the ratings (approximately 3.5 – 4) that people give to other people or to themselves, but higher than in the previous version of our dialogue.

We used a scale measuring extraversion of the robot (cheerful, attractive, happy, friendly, optimistic, warm) because extraverts tend to elicit more talk from other people than do introverts. We found no differences due to robot gender. Across conditions, the robot was rated as moderately extraverted.

Other items measured the robot's dominance, compassion, and likeability. In these items, most ratings were the same across robot gender and subject gender, and in moderate ranges of the scale. However, men's ratings of the female robot were significantly lower than women's ratings of either gendered robot and lower than men's ratings of the male robot. Thus, men rated the female robot as lower in leadership and higher in dominance ($p < .001$), as somewhat less tender and compassionate ($p < .07$), and as marginally less likeable ($p > .10$). Because of these differences, we examined whether subjects' ratings influenced how much they talked with the robot. We found that their ratings of the robot's assertiveness, compassion, and likeableness were correlated with amount of talking, so we used these ratings as control variables in the subsequent analyses. Use of these control variables did not change the direction of results.

B. Hypothesis #1: Robot's Perceived Knowledge

We next examined data from items about the robot's knowledge and personality. We did not find an overall difference in the subjects' estimates of the robot's knowledge of dating (after their interaction with the robot) but, instead, women subjects tended to rate the female robot as having more knowledge about dating whereas men rated the male robot as having more knowledge about dating; these differences did not quite attain statistical significance ($p = .14$). Overall, those who felt the robot knew less spent

a greater amount of time talking to the robot ($r = -.31$) and answering Jill and John questions ($r = -.13$).

C. Hypothesis #2: Subjects' Talk

As noted above, we measured the number of words that men and women subjects used in communicating with the male or female robot about Jill's and John's appropriate conduct on a first date. We predicted that (a) subjects would use fewer words in talking to the female than the male robot, (b) that women would talk less to a female robot than men would, and men would talk less to a male robot than women would, and (c) that the least talk should occur in the condition where women conversed with a female robot about dating norms for women.

Overall we found a significant triple interaction of subject gender, robot gender, and Jill vs. John questions ($F [1, 25] = 4, p = .05$). Although these results are not very strong, due to the comparatively small sample size (fewer than 10 subjects per condition), the results reflect the predictions, shown in Figures 4a and 4b.

1. Subjects said more words to the male robot than to the female robot.
2. Men said more words to the female robot than women did, whereas women said more words to the male robot than men did.
3. The fewest words were said to the female robot by women about Jill.

VII. DISCUSSION

This work has a few design implications for human robot interaction. First, the theory predicts that people will make assumptions about the knowledge of a robot based on social cues attached to it. Hence we cannot assume that people approach a robot *tabla rasa* but instead have a mental model that, at the outset, is anchored by impressions of the robot's persona. Second, designers can manipulate a robot's appearance, conduct, and context to convey the robot's knowledge or they can design a robot whose cues adapt to different user models. Third, because people will

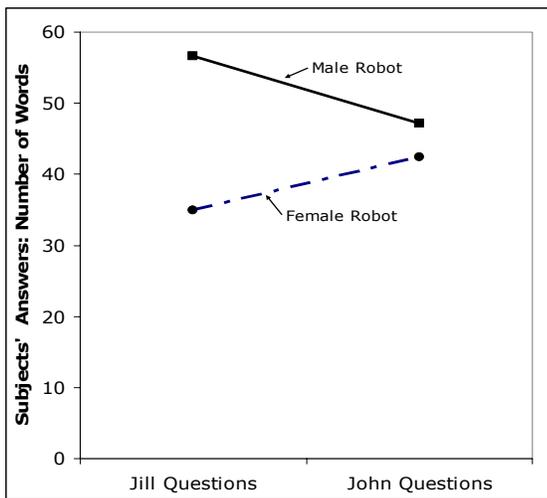


Fig. 4a. Women's responses.

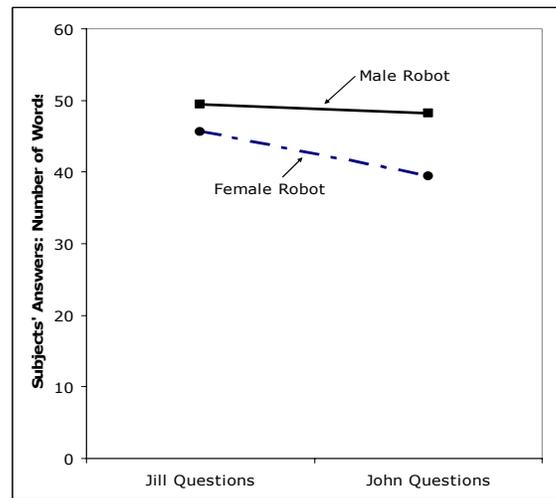


Fig. 4b Men's responses.

adjust their conversation with a robot depending on their perceived common ground with it, designers will need to make decisions about their goals for this conversation. Do they want the human-robot interaction to be as efficient as possible, or do they want it to be more discursive?

Our study results suggest that if a robot's task is stereotypically associated with different social groups, then we may want to design the robot's interface to fit or to violate the stereotype. For example, a mechanic's helper-robot, if stereotypical, would be male. If we wanted this robot to have minimal and efficient conversation with users about what tools they need, how to assist, and so forth, then the mechanic's helper should be male. Suppose instead that we wanted users to provide more information, to explain themselves, to be redundant. This might be a design goal if the robot were not just a mechanic's helper robot, but a general assistant, and the robot was not specially designed for the task. Then, the robotic assistant should be anti-stereotypic for the task – i.e. female for the mechanic's helper. This might be beneficial if the robot's dialog system were aided by redundancy. In that case, we speculate that people will be more redundant in their conversation with the robot if it does not fit the stereotypic persona for that topic (e.g., a female mechanic, a male nurse). For example, we can create a broader model in users about what the robot knows by making the robot look more feminine and youthful, characteristics associated with jobs like docent, teacher, or nurse, rather than mechanic's helper. Since to achieve common ground people adapt their speech to the perceived needs of the other, they should adjust their speech to the perceived needs of the robot. So, just as we speak more clearly to three year olds than we speak to our peers, we will speak more clearly to the robot we think is ignorant than we will to the smart one.

A. Future Work

Because this study represents a first demonstration of a common ground effect in human-robot interaction, we must regard it as preliminary. We believe there are many worthwhile domains to explore in seeking replication and extension of the theory to human-robot interaction, for example, whether people find common ground with a robot's emotional state, preferences, or decision biases. This work may also lead to some new ways that designers can adapt dialogue systems such that people and robots will communicate more clearly.

ACKNOWLEDGMENTS

We thank the People and Robots, Nursebot, and Social Robots project teams for their suggestions and help in designing the human-robot interactions used in this study.

REFERENCES

[1] Clark, H. H., & S. E. Brennan, "Grounding in communication." In L. B. Resnick, J. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127-149). Washington, DC: APA, 1991.

[2] R. S. Nickerson, "How we know—and sometimes misjudge—what others know: Imputing one's own knowledge to others," *Psychological Bulletin*, vol. 125, no. 6, pp. 737-759, 1999.

[3] A. Tversky and D. Kahneman, "Judgement under uncertainty: Heuristics and biases," *Science*, vol. 185, pp. 1124-1131, September 27, 1974.

[4] M. Ross and D. Holmberg, "Recounting the past: Gender differences in the recall of events in the history of a close relationship." In J. M. Olson and M. P. Zanna (Eds.), *Self-inferences processes: The Ontario Symposium*, Vol. 6, Hillsdale, NJ: Erlbaum, 1990, pp. 135-152.

[5] S.-L. Lee, S. Kiesler, I.Y. Lau, and C.Y. Chiu, "Human Mental Models of Humanoid Robots," *2005 IEEE International Conference on Robotics and Automation*, April 2005.

[6] H. H. Clark, "*Arenas of language use*," Chicago: University of Chicago Press, 1992.

[7] E. A. Issacs and H. H. Clark, "References in conversation between experts and novices," *Journal of Experimental Psychology: General*, vol. 116, no. 1, pp. 26-37, 1987.

[8] S. Fussell and R. Krauss, "Coordination of knowledge in communication: Effects of speakers' assumptions about what others know." *Journal of Personality and Social Psychology*, vol. 62, pp. 378-391, 1992.

[9] W. Wood, N. Rhodes, M. Whelan, M., "Sex differences in positive well-being: A consideration of emotional style and marital status." *Psychological Bulletin*, vol. 106, pp. 249-264, 1989.

[10] R. M. Dawes, "Statistical criteria for establishing a truly false consensus effect," *Journal of Experimental Social Psychology*, vol. 25, pp. 1-17, 1989.

[11] M. Scheeff, J. Pinto, K. Rahardja, S. Snibbe, and R. Tow, "Ex with Sparky: A social robot," *Proceedings of the Workshop on Interactive Robot Entertainment*, 2000.

[12] C. Breazeal and B. Scassellati, "How to build robots that make friends and influence people," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Kyoju, Japan, 1999.

[13] W. Bugard et al, "Experiences with the interactive museum tour-guide robot," *Artificial Intelligence*, vol. 114, nos. 1-2, pp. 3-55, 1999.

[14] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz, "Probabilistic algorithms and the interactive museum tour-guide robot minerva," *International Journal of Robotics Research*, 19(11):972-999, 2000.

[15] T. Willeke and C. Kunz and I. Nourbakhsh, "The History of the Mobot Museum Robot Series: An Evolutionary Study," *Proceedings of FLAIRS 2001*, 2001.

[16] R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. Shultz, and J. Wang, "Designing robots for long-term social interaction," *2005 IEEE International Conference on Robotics and Automation*, April 2005.

[17] T. Kanda, T. Hirano, and D. Eaton, "Interactive robots as social partners and peer tutors for children: A field trial," *Human Computer Interaction*, vol. 19, pp. 61-84, 2004.

[18] M. Montemerlo, J. Pineau, N. Roy, S. Thrun and V. Verma, "Experiences with a mobile robotic guide for the elderly," *18th National Conference on Artificial Intelligence*, pp. 587—592, 2002.

[19] M.R. Laner, and N.A. Ventrone, "Egalitarian daters/traditionalist dates," *Journal of Family Issues*, vol. 19, no. 4, pp. 468-477, July 1998.

[20] M.R. Laner, N.A. Ventrone, "Dating scripts revisited," *Journal of Family Issues*, vol. 21, no. 4, pp. 488-499, May 2000.

[21] K.A. Lenzo, and A.W. Black, *Theta*, Cepstral, <http://www.cepstral.com>

[22] S.L. Bem, *Bem Sex-Role Inventory*, Palo Alto: Consulting Psychologists Press, Inc., 1976.

[23] R. Wallace, *Alice*, ALICE Artificial Intelligence Foundation, <http://www.alicebot.org/>

[24] A.D.I. Kramer, S.R. Fussell, and L.D. Setlock, "Text analysis as a tool for analyzing conversation in online support groups," *Extended Abstracts of the 2004 Conference on Human Factors and Computing Systems*, pp. 1485-1488, April 2004.