How People Talk with Robots: Designing Dialogue to Reduce User Uncertainty

Kerstin Fischer

■ *If human-robot interaction is mainly shaped by users' strategies to deal with their unfamiliar artificial communication partner, as it is suggested here, robot dialogue design should orient toward reducing users' uncertainty about the affordances of the robot and the joint task. Two experiments are presented that investigate the impact of verbal robot utterances on users' behavior; results show that users react sensitively to subtle linguistic cues that may guide them into appropriate understandings of the robot. Furthermore, the role of user expectations and robot appearance are discussed in the light of the model presented.*

Verbal interaction with robots is not characterized by a set of well-defined properties; what users actually say when faced with a robot, especially if they are not engineers or computer scientists, can hardly be predicted. In this article I argue that the peculiarities of human-robot dialogue are best understood as users' strategies to deal with what they understand the challenges of the situation to consist in (Fischer 2006a). That is, users interact with artificial communication partners on the basis of what they consider to be potentially problematic, what the task comprises, what the robot can understand, and so on, that is, what they consider its affordances to be (Gibson 1977). The variation observable between users is then due to different concepts of what the affordances of the system really are.

Although this observation seems trivial, there are in fact several other suggestions for what human-robot interaction might be shaped by. The most prominent is certainly the idea that users transfer mindlessly from social interactions between humans to interactions with computers and robots (for example, Nass and Moon [2000], Nass and Brave [2005]). In particular, Nass and colleagues have investigated a broad range of social behaviors, with the result that people were found to react to a computer's behavior as they do to other humans' behavior (Fogg and Nass 1997), or that they transfer to the agent human characteristics, such as intentionality (Ju and Takayama 2009), ethnicity (Pratt et al. 2007) or gender (Nass and Brave 2005). Nass proposes that the reasons for this transfer, which has also been called *media equation* (Reeves and Nass 1996), is mindlessness, an error due to thousands of years of evolution in social environments (Nass 2004).

A second proposal, which often remains implicit, is that there is either a particular conventional way of speaking to computers and robots, or that users make use of other conventional varieties, for instance, the way of talking to children. The linguistic term for such a situation-specific variety is *register* (Krause and Hitzenberger 1992).

However, the interpersonal variation observable in the ways in which people talk to robots (see Fischer 2006a, 2006b, 2011)

 ISSN 0738-4602

cannot be accounted for in either of the two suggestions. Especially the finding that some users do not treat robots as social actors constitutes a challenge to the mindless transfer hypothesis (Fischer 2011; Fussel et al. 2008). Furthermore, in a recent study we demonstrate that the adaptations people make are informed by the kind of feedback the robot provides; variation occurs especially where no such feedback is available (Fischer et al. 2011). Thus, users' adjustments depend on the information the robot displays to them.

Further evidence for the model proposed here, in which users' behavior is guided by their understanding of the affordances of the human-robot interaction situation, is that people can be shown to be generally extremely cooperative (Fischer 2006a). The considerable amount of cooperative behavior people are willing to invest in the interactions with artificial systems is in line with pragmatic theories that consider communication to rest essentially on cooperation (for example, Grice 1975).

If users' behavior in verbal human-robot interaction depends on their sense-making processes, there are three relevant domains for dialogue design in which the designer can shape the interactions; on the side of the user, these are certain preferences and preconceptions, on the side of the robot these concern its functionalities as they are apparent from the robot's appearance and the robot's behavior in the interaction itself (Zoltan-Ford 1991).

This article provides an overview of current research in these areas and presents two empirical studies of the complex role the robot's verbal output may play in human-robot interaction. The results from this investigation will have considerable design consequences.

User Expectations

An important factor in human-robot interaction concerns the personal preferences and preconceptions of the human user. While it is self-evident and somewhat trivial that people differ from each other, user expectations are usually not taken into account in studies on human-robot interaction; however, several studies have demonstrated the great impact of expectations on interactional success.

For instance, Paepcke and Takayama (2010) find a priori expectations to have a considerable impact on users' experience with, and perception of, robotic pets. They manipulated participants' expectations using different advertising texts by means of which the robots were introduced; depending on whether expectations were high or low, users evaluated the robots' capabilities significantly differently.

Especially for personal companions, individual user expectations may play a crucial role. For instance, Turkle (2006) reports on interview data she elicited in nursery homes and elementary schools from elderly people and children who had kept either the Sony's robotic dog Aibo or My Real Baby, a doll-like robot, for several weeks. The interviews show that both children and elderly people vary considerable regarding their relationships with these relational artifacts. Turkle concludes that human-robot relationships are highly individual and similar to Rorschach test projections of the self.

That users' relationships with relational artifacts are influenced by individual personal needs, is supported by a study by Lee et al. (2006) who find participants' degrees of loneliness to correlate with their judgment of the social presence of the robot and its overall behavior; furthermore, they find lonely people to provide more positive responses to social agents.

While users' personal needs may influence the relationships they engage in with relational artifacts, user expectations may also be relevant for task-oriented interactions. For instance, Fischer and Moratz (2001) show that users subconsciously employ a hierarchy of robot behaviors according to which they produce their instructions. In the experiment, users considered moving to a goal to be more complex than moving along a path, which in turn was considered to be more difficult than moving at all. The endpoint of the hierarchy was regarded to be instrumental action, like turning the wheels. If this hierarchy was an appropriate model of robot functionality, an instruction like "turn your rear wheels" should be easier to process for the robot than "move left" or "move to the left object," which is of course only true if the robot designer implemented such states at all. In case of communication problems, users will only move down in their conceptual hierarchies and thus may not be able to identify the single kind of instruction the robot was designed to understand, in the case of the robot employed in Fischer and Moratz (2001), the most complex, goal-oriented instruction. Following is a (translated) example of the hierarchy instantiated in a user's instructions; after every instruction typed by the user, the robot printed "error" on the screen:

drive straight ahead to the right
turn 45 degrees to the right
turn to the right
drive 10 cm ahead
come on
stat driving
start driving
engine on

While robots can be implemented to fulfill sophisticated behaviors, users may expect them to have learned more elementary actions first. User

expectations thus play a crucial role in successful human-robot interaction, especially with naïve users.

Furthermore, user expectations may determine different interaction styles. From an engineering perspective those expectations are interesting that correspond to different user groups; even more interesting are styles that can be identified unobtrusively and online and that predict user behaviors throughout a dialogue. One such phenomenon is users' attention to the robot as a social actor or not (Fischer 2006a, 2006b, 2011; Lee, Kiesler, and Forlizzi 2010); thus, unlike other proposals which suggest the interpretation of a robot as a social actor to be involuntary and based on evolutionary psychology (Reeves and Nass 1996, Nass and Moon 2000), several studies have shown that while some people do treat the robot as a social actor, others do not, and that in both cases people are not mindless about their choices (Fischer 2011; Fischer forthcoming; Fischer et al. 2011). A simple indicator of the group to which a given user belongs concerning his or her inclination to treat the robot as a social actor is a user's spontaneous reaction to the robot's greeting, which predicts numerous linguistic behaviors even over time (Fischer 2011). Whether people react to the social aspects of the greeting "yes, hello, how do you do?" by means of, for instance, *thanks* or even *and how do you do?*, can be statistically correlated with certain linguistic behaviors; furthermore, Lee, Kiesler and Forlizzi (2010), taking up suggestions by Fischer (2006b), find that also numerous conversational strategies can be predicted from the greeting.

Since users' expectations thus play a crucial role, human-robot interaction research may profitably engage in determining how such expectations can be shaped. One possible site to do so is the robot's appearance.

Robot Appearance

Numerous studies have shown that the robot's appearance influences the interaction. Especially social interpretations have been found to be triggered by anthropomorphic features (for example, Nass 2004). For instance, Koda and Maes (1996) find differences in the amount of anthropomorphization depending on the amount of human characteristics of the robot. This finding is supported by numerous studies, for instance, Reeves and Nass (1996), Powers and Kiesler (2006) and DiSalvo et al. (2002).

Additional evidence comes from the study of avatars and figures in virtual environments; for instance, Baylor, Rosenberg-Kima, and Plant (2006) investigate different personae as online tutors and find considerable differences in the amount of influence these tutors have on the student depending on their appearance. Similarly, Forlizzi et al. (2007) find effects for agent appearance that reflect common gender stereotypes.

Moreover, there are effects depending on how similar the artificial agent is to the user (Lee and Nass 2003, Yee et al. 2007). Finally, Parise et al. (1996) compare the degrees of cooperativeness exhibited by participants in interactions with humanlike versus doglike agents. Although participants found the doglike agent cuter and more likable, they cooperated more with the humanlike agent and even more so with the human interlocutor. These findings suggest that the speech directed at robots that differ in appearance will differ considerably.

However, my own studies of robot appearance lead to slightly different conclusions (Fischer 2008, Fischer forthcoming); I compared three robots, a large metal spider (called *Scorpion*) built by Spenneberg and Kirchner (2002), the robotic dog Aibo, and a Pioneer robotic platform with a mounted camera, which is basically only a red box on wheels with a wooden construction on top. The robots' behavior was scripted and carried out by a human wizard behind a screen and thus identical for all participants in all three conditions. While this procedure may seem a little awkward or unnatural, for the participants themselves the situation appears to be completely interactional; occasional mismatches between the users' utterances and the robot's responses are easily categorized by the participants as recognition errors or as problems in lexical choice, for example, which in itself provides us with useful clues to the ways users make sense of the human-robot interaction situation (see Fischer [2010]).

The participants' task was joint attention spatial instruction; that is, users had to instruct the robot to move to particular places within the room. In these data, no significant differences with regard to the major linguistic measures (complexity, interactivity, perspective taking, anthropomorphization, and so on) could be found. My suggestion is that the lack of findings in these experiments is due to the fact that the robots' appearances did not provide any information as to the robots' functionalities relevant for the spatial instruction task. However, in the experiments with Aibo and Scorpion, yet not with the Pioneer whose camera was obvious, participants wondered for example about the robots' sensors (Fischer and Bateman 2006). Since it is of great influence for an interaction to know what the partner can perceive (for example, Hanna, Tanenhaus, and Trueswell 2003), designing robots in a way that their sensors are apparent to the user may facilitate the interaction by reducing users' uncertainty about their artificial communication partner by communicating the robot's functionalities.

Appearance may thus provide users with clues to the robot's capabilities if they are compatible with users' particular partner models (for instance, social characteristics displayed by robots may not influence those users into social behaviors whose preconceptions about robots do not allow this; Fischer [forthcoming]) and if these characteristics provide clues to the robot's functionalities, that is, if they are *interpretable* to the user.

To sum up, robots' appearance can be used to influence users' preconceptions if it provides consistent clues to particular robot functionalities.

To Talk or Not to Talk

Another possibility to influence users' behavior is through human-robot dialogue. Thus, in this section, we consider the role of verbal output — does the robot have to talk? We consider what purposes robot utterances serve and what their effects are on users' behavior. In order to address these questions, we compare interactions with a robotic wheelchair (Lankenau and Roefer 2001) in two conditions that differ regarding whether or not the robot produces verbal behavior; in both conditions, participants had to carry out the same four tasks with the same robotic wheelchair (see figure 1).

Tasks

Participants had to steer the robot around in order to "train" it on the environment and to provide it with verbal explanations, in particular (1) familiarize the robot with locations in a room for handicapped people by steering it around and labeling particularly interesting locations; (2) summarize the locations trained from a static position; (3) familiarize the robot with locations in the building; and (4) instruct the robot how to get to one of these locations.

Procedure

Participants were seated in the robotic wheelchair, whose steering mechanism was briefly explained. Then they were given the instructions for task 1 to read. Instructions, setup, and the robot were identical in both conditions. The conditions differ in that in condition 2, the robot engages in verbal dialogue with the user in tasks 1, 2, and 4. In task 3, the robot stated explicitly that it would remain silent during the training. Participants had to steer the robot themselves and were free to move to as many locations as they considered relevant. There were no behavioral instructions. The robot was supposed to move autonomously only at the end of the instructions when it was supposed to take the user to one of the locations it was previously trained on.

All verbal robot output in condition 2 was scripted and manipulated by a human "wizard" hidden behind a flexible wall. Thus, for each location the respective participant steered the robot to, there was a set of robot utterances to be played in a particular order (see table 1, Version 1). For some utterances, the wizard had different choices depending on the label the participant had used, for instance, *sofa* versus *couch*, *fridge* versus *refrigerator*, *stove* versus *hot plate*. While this procedure may seem unnatural, the resulting dialogues are in fact quite fluent:

R022: (1) we are now going to **turn** — and **go** to the (1) dining room table. (6)

Robot: Okay. (15) I'm sorry, which kind of table?

R022: (1) m dining room table. - dinner **table**,

Robot: (3) The what?

R022: (1) dinner table, (2)

Robot: Ah, the dinner table. (2) Just to get this right, I am in front of it, is that right?

R022: - yes that's right.

ROBOT: Okay, thank you.

R022: you're welcome.

Scripting the robot output does not only render all robot output identical and thus the dialogues comparable across persons and conditions, providing a unique methodological opportunity to study the influence of isolated variables, it is also computationally the cheapest method possible. Thus, it should be impossible to discard the results of this study on the basis of the assumption that the current dialogues necessitate unrealistically sophisticated speech technology.

Participants

Participants were 9 native speakers of English in condition 1 and 11 in condition 2. All were either staff or exchange students at the University of Bremen.

Data Analysis

All dialogues (or monologs in condition 1) were recorded and transcribed. The transcripts were analyzed semiautomatically concerning the amount of speaking, length and depth of interaction, linguistic variability, off-talk, and out-of-domain vocabulary.

Results

The analysis reveals significant differences in participants' linguistic behavior; we compare here participants' behavior in the third task since the robot is not speaking in either of the two settings and so the two conditions are completely comparable. The comparison shows that in condition 1 (without verbal robot output in the previous two tasks), users employ almost five times more different words (306) than in condition 2 (68). While in the second condition no out-of-domain vocabulary is used, at least 57 words are completely out-of-domain in the first condition, for instance:

Figure 1. The Robotic Wheelchair "Rolland."

alley, boom, bulletin, cigarette, earlier, extinguisher, extra, forms, fun, hurt, lady, language, loop, love, mind, multiply, nobody, nudge, packaging, people, person, pretty, rabbits, renovating, shelf, smoke, smoking, track, traffic, tricky, trouble, visit

Thus, in the nonverbal condition 1, users were not very focused on what the robot could possibly understand. In contrast, participants in condition 2 had a much more accurate understanding of the requirements of the task.

Furthermore, there is significantly more off-talk, $F(1, 17) = 17.79430$, $p < .001$, in condition 1 than in condition 2. That is, participants turned to the experimenter almost seven times more often in condition 1 than in condition 2 (the mean number of instances of off-talk in condition 1 is 0.13 [*sd* 0.08] and 0.02 [*sd* 0.02] in condition 2).

Another result concerns the use of personal pronouns; these pronouns are indicators for how participants understand the situation: as joint action (*we*) or as individual action (*I* or *you*). The analysis shows that there is much less identification with the robot as body extension in condition 1: There are significantly fewer instances of *we,* $F(1, 17) = 8.75$, $p < .01$, in condition 1 compared with condition 2 in the third task (the mean is 0.03 [*sd* 0.04] for condition 1 and 0.32 [*sd* 0.29] for condition 2).

Discussion

Dialogue provides subtle clues to the robot's functionality and thus to adequate partner modeling. In the interactions with the nonverbal robot, users employed a huge lexicon and a large amount of out-of-domain vocabulary, which is very difficult to handle for state-of-the-art language technology. Furthermore, users did not gain a proper understanding of what the task was really about and what the robot needed to know. In contrast, very simple robot utterances provided users with the information they needed, even for the task in

Robot Output Version 1	Robot Output Version 2
INTRO:	
Yes, hello, how do you do?	Yes, hello, how do you do?
you can take us now to a place you want to name.	I'm ready.
FRIDGE:	
	I understood refrigerator.
	/ I understood fridge.
is this where you want to be to open it?	is this where you want to be to open it?
okay, thank you.	okay, thank you.
ARMCHAIR:	
I'm sorry?	I did not understand.
the armchair?	Did you say armchair?
Am I in front of it?	Am I in front of it?
Oh okay, I see which object you mean.	Oh okay, I see which object you mean.

Table 1. Example of Scripted Robot Output Sequences.

which the robot was not speaking itself. Moreover, participants in condition 2 regarded the robot more as a communication partner, as evident from the lower amount of off-talk, and regarded the situation more as joint action.

Robot Dialogue Design

In the previous section, we have seen that robot utterances may have a considerable impact on users' understanding of a task, the variability of linguistic structures chosen, the amount of out-of-domain vocabulary used and users' cooperativity in a given situation. In this section, we investigate the extent to which individual robot utterances contributed to this result.

Tasks

The tasks were identical to those in the previous experiment.

Procedure

The procedure was identical to the procedure in the previous experiment. The verbal output of the robot in this experiment in condition 1 is identical to the robot utterances used in the previous experiment. However, in condition 2 of this experiment, there are slight changes in the robot utterance design.

The robot utterances in both conditions were designed on the basis of conversational grounding theory (Clark and Schaefer 1989). In particular, utterances were so designed as to provide the participant with feedback on what had been understood about his or her instruction and thus to provide acceptance of the partner's contribution. The two versions now differ in two respects: First, the initial utterance by the robot in version 1, "you can take us now to a place you want to name," which was meant to guide the user into an appropriate understanding of the task, was replaced by a generic assertion of readiness. Second, the implicit means to provide feedback to the user in version 1 were replaced by explicit means. For example, the clarification question "the plant?" provides the user implicitly with a candidate understanding. In the second version, this was replaced by an utterance signaling feedback explicitly, namely "I understood plant." In addition, the robot uses additional feedback when reaching a location. Thus, the changes are indeed minimal (see table 1), the strategies all being regular conversational grounding mechanisms, just differing regarding their degree of explicitness.

Participants

Participants were 13 native speakers of German in condition 1 and 9 in condition 2. All were students at the University of Bremen who had previously been tested for their English language competence (all had attained C1 level).

Results

There are significant differences concerning the number of words, $F(1, 22) = 6.221453$, $p = 0.021494$, used by the participants; that is, users in the first condition talked significantly more to the robot (203.9 words on the average, $sd = 72.6$) than users in the second condition (132.8 words on average, $sd = 53.4$).

This corresponds to the smaller number of locations that participants in the second condition drove the wheelchair to ($F(1, 22) = 8.689840$, $p = 0.007955$); participants in the first condition explained ten locations on average to the robot ($sd = 2.8$) whereas users in condition 2 finished their instructions after only 6.67 locations on average ($sd = 2.3$). Thus, users in the first condition engaged in the interactions much more than users in the second condition. This corresponds to the finding that in condition 2, there are numerous instances in which a user ignores the robot utterances completely, for instance:

R031: sofa. — that's the sofa. — back,

Robot: I understood sofa.

R031: — backwards, —

R028: um next we turn,
— go backwards, — to um,
— the computer table. (2) (laughter)

Robot: Did you say computer?

R028: (3) um (3) 'kay, now, (1) we're going,

Robot: Did you say computer?

R028: (3) we're going to, (2)

Thus, the slight changes in the robot's utterances lead to less responsive behavior by the participants.

Discussion

The hypothesis regarding the two versions of robot output was that the implicit feedback of the first version contributes to the fluency and naturalness of the dialogues. This prediction is confirmed by the fact that participants in condition 2 spent significantly less time with the robot as evident from the fewer numbers of locations they took the robot to, as well as the significantly lower number of words. Thus, the exact wording of the robot's utterances plays a crucial role concerning users' engagement in human-robot interactions.

Conclusion: Human-Robot Dialogue as Uncertainty Reduction

I have introduced an approach to human-robot interaction that sees it primarily as an attempt of users to deal with what they understand as the affordances of the system and the requirements of the situation. Thus, in contrast to other approaches, the mindless transfer hypothesis or the register hypothesis, the current approach puts people in the position to actively react to various clues provided by the robot. People do bring different preconceptions into the interactions, which need to be accounted for in interaction design; yet in general, users are cooperative, but possibly lack information about what would be appropriate behavior. The goal of robot dialogue design thus has to be to reduce users' uncertainty and to guide them into appropriate partner models.

We investigated two areas in which a system designer can influence human-robot dialogue: appearance and dialogue design. Regarding appearance, I suggested that design should orient at providing clues to robot functionalities. Similarly, regarding dialogue design, we found that linguistic robot utterances are powerful means to guide the user into appropriate behaviors towards the robot. Second, we found that users make use of every single clue they can get to make sense of their communication partner and to find out how to solve the communication task posed to them efficiently. In particular, in the qualitative and quantitative analyses of the effects robot utterances may have on the interaction, we could see how useful robot utterances can be to provide users with a suitable partner model that matches the robot's real functionalities. In the analysis of the two versions of robot utterances, participants furthermore reacted sensitively to minimal linguistic cues, and thus utterance design has to be carried out very thoughtfully and perhaps even in iterations of usability testing.

The design implications are thus that human-robot interaction can profit considerably from carefully crafted linguistic robot output. Especially with naïve users of personal or welfare robots, linguistic output can facilitate uncertainty reduction and subtly guide users into appropriate behaviors, that is, behaviors that are interpretable for a robot. The dialogues used here were in fact scripted and completely insensitive to the users' utterances, yet the dialogues arising were fluent and natural, and participants did not only like the interactions but also changed their opinion about robots (Andonova 2006). Thus with the cheapest technical means a very good interactional result was achieved since the utterances used were carefully crafted. Consequently, attending more to the linguistic properties of individual robot utterances may even neutralize some of the technical shortcomings of current language technology.

Acknowledgements

This research has been partially funded by the European Union in the framework of the ITALK project under grant number 214668. The data reported on were elicited in the framework of the Special Research Area SFB/TR8 "Spatial Cognition," kindly funded by the German Research Foundation (DFG).

References

Andonova, E. 2006. On Changing Mental Models of a Wheelchair Robot. In *Proceedings of the Workshop on How People Talk to Computers, Robots, and Other Artificial Communication Partners*, SFB/TR8 Report 010-09_2006. Hansewissenschaftskolleg, Delmenhorst, April 21–23.

Baylor, A. L.; Rosenberg-Kima, R. B.; and Plant, E. A. 2006. Interface Agents as Social Models: The Impact of Appearance on Females' Attitude Toward Engineering. In *Extended Abstracts Proceedings of the 2006 Conference on Human Factors in Computing Systems*, 526–531. New York: Association for Computing Machinery.

Clark, H. H., and Schaefer, E. F. 1989. Contributing to Discourse. *Cognitive Science* 13(2): 259–294.

DiSalvo, C.; Gemperle, F.; Forlizzi, J.; and Kiesler, S. 2002. All Robots Are Not Created Equal: The Design and Perception of Humanoid Robot Heads. In *Proceedings of the Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, 321–326. New York: Association for Computing Machinery.

Fischer, K. (forthcoming): Recipient Design, Alignment, Feedback: The Role of the Addressee in So-Called Simplified Registers. Habilitation Thesis, University of Bremen, Bremen, Germany.

Fischer, K. 2006a. *What Computer Talk Is and Isn't: Human-Computer Conversation as Intercultural Communication*. Saarbrücken, Germany: AQ Verlag.

Fischer, K. 2006b. The Role of Users' Preconceptions in Talking to Computers and

Robots. In *Proceedings of the Workshop on How People Talk to Computers, Robots, and Other Artificial Communication Partners*, SFB/TR8 Report 010-09_2006, 112–130. Hansewissenschaftskolleg, Delmenhorst, April 21–23.

Fischer, K. 2008. The Role of Users' Concepts of the Robot in Human-Robot Spatial Instruction. In *Spatial Cognition V: Reasoning, Action, Interaction*. Lecture Notes in Computer Science 4387, 76–89. Berlin: Springer.

Fischer, K. 2010. Why It Is Interesting to Investigate How People Talk to Computers and Robots: Introduction to the Special Issue. *Journal of Pragmatics*, 42(9): 2349–2354.

Fischer, K. 2011. Interpersonal Variation in Understanding Robots as Social Actors. In *Proceedings of the 6th International Human-Robot Interaction Conference,* 53–60. New York: Association for Computing Machinery.

Fischer, K., and Bateman, J. A. 2006. Keeping the Initiative: An Empirically Motivated Approach to Predicting User-Initiated Dialogue Contributions in HCI. In *Proceedings of the 11st Conference of the European Chapter of the Association for Computational Linguistics*. Stroudsburg, PA: The Association for Computer Linguistics

Fischer. K.; Foth, K.; Rohlfing, K.; and Wrede, B. 2011. Mindful Tutors: Linguistic Choice and Action Demonstration in Speech to Infants and a Simulated Robot. *Interaction Studies* 12(1): 134–161.

Fischer, K., and Moratz, R. 2001. From Communicative Strategies to Cognitive Modelling. Paper presented at the First International Workshop on Epigenetic Robotics, Lund, Sweden, 17–18 September.

Fogg, B. J., and Nass, C. 1997. Silicon Sycophants: The Effects of Computers that Flatter. *International Journal of Human-Computer Studies* 46(5): 551–561.

Forlizzi, J.; Zimmerman, J.; Mancuso, V.; and Kwak, S. 2007. How Interface Agents Affect Interaction Between Humans and Computers. In *Proceedings of the 2007 International Conference on Designing Pleasurable Products and Interfaces*, 209–221. New York: Association for Computing Machinery.

Fussell, S. R.; Kiesler, S.; Setlock, L. D.; and Yew, V. 2008. How People Anthropomorphize Robots. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*, 145–152. New York: Association for Computing Machinery.

Gibson, J. J. 1977. The Theory of Affordances. In *Perceiving, Acting, and Knowing,* ed. R. Shaw and J. Bransford. Hillsdale, NJ: Lawrence Erlbaum, 67–82.

Grice, H. P. 1975. Logic and Conversation. In *Syntax and Semantics: Speech Acts*, Volume 3, 1–58. New York: Academic Press.

Hanna, J. E.; Tanenhaus, M. K.; and Trueswell, J. C. 2003. The Effects of Common Ground and Perspective on Domains of Referential Interpretation. *Journal of Memory and Language* 49(1): 43–61.

Ju, W., and Takayama, L. 2009. Approachability: How People Interpret Automatic Door Movement as Gesture. *International Journal of Design and Emotion* 3(2): 77–86.

Koda, T., and Maes, P. 1996. Agents with Faces: The Effects of Personification of Agents. In *Proceedings of the 5th IEEE Workshop on Robot and Human Communication,* 185–194. Los Alamitos, CA: IEEE Computer Society.

Krause, J., and Hitzenberger, L. 1992. *Computer Talk*. Hildesheim: Olms.

Lankenau, A., and Roefer, T. 2001. A Safe and Versatile Mobility Assistant. *IEEE Robotics and Automation Magazine* 7(1): 29–37.

Lee, K. M., and Nass, C. 2003. Designing Social Presence of Social Actors in Human Computer Interaction. In *Proceedings of the 2003 Conference on Human Factors in Computing Systems*, 1: 289–296. New York: Association for Computing Machinery.

Lee, K. M.; Jung, Y.; Kim, J.; and Kim, S. R. 2006. Are Physically Embodied Social Agents Better Than Disembodied Social Agents? The Effects of Physical Embodiment, Tactile Interaction, and People's Loneliness in Human-Robot Interaction. *International Journal of Human-Computer Studies* 64(10): 962–973.

Lee, M. K.; Kiesler, S.; and Forlizzi, J. 2010. Receptionist or Information Kiosk: How Do People Talk with a Robot? In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. New York: Association for Computing Machinery.

Nass, C. 2004. Etiquette Equality: Exhibitions and Expectations of Computer Politeness. *Communications of the ACM* 47(4): 35–37.

Nass, C., and Brave, S. 2005. *Wired for Speech: How Voice Activates and Advances the Human-Computer Relationship*. Cambridge, MA: The MIT Press.

Nass, C., and Moon, Y. 2000. Machines and Mindlessness: Social Responses to Computers. *Journal of Social Issues* 56(1): 81–103.

Paepke, S., and Takayama, L. 2010. Judging a Bot by Its Cover: An Experiment on Expectation Setting for Personal Robots. In *Proceedings of the 5th ACM/IEEE International Conference on Human Robot Interaction*. New York: Association for Computing Machinery.

Parise, S.; Kiesler, S.; Sproull, L.; and Waters, K. 1996. My Partner Is a Real Dog: Cooperation with Social Agents. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*, 399–408. New York: Association for Computing Machinery.

Powers, A., and Kiesler, S. 2006. The Advisor Robot: Tracing People's Mental Model from a Robot's Physical Attributes, 218–225. In *Proceedings of the First ACM SIGCHI/SIGART Conference on Human-Robot Interaction*. New York: Association for Computing Machinery.

Pratt, J.A.; Hauser, M.; Ugray, Z.; and Patterson, O. 2007. Looking at Human-Computer Interface Design: Effects of Ethnicity in Computer Agents. *Interacting with Computers* 19(4): 512–523.

Reeves, B., and Nass, C. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge, UK: Cambridge University Press.

Spenneberg, D., and Kirchner, F. 2002. SCORPION: A Biomimetic Walking Robot. *Robotik* 2002. Berlin: VDI Verlag.

Turkle, S. 2006. Robot as Rorschach: New Complicities for Companionship. In Human Implications of Human-Robot Interaction: Papers from the 2006 AAAI Workshop. Technical Report WS-06-09. Menlo Park, CA: AAAI Press.

Yee, N.; Bailenson, J.; Urbanek, M.; Chang, F.; and Merget, D. 2007. The Unbearable Likeness of Being Digital: The Persistence of Nonverbal Social Norms in Online Virtual Environments. *CyberPsychology and Behavior* 10(1): 115–121.

Zoltan-Ford, E. 1991. How to Get People to Say and Type What Computers Can Understand. *International Journal of Man-Machine Studies* 34(4): 527–647.

Kerstin Fischer is an associate professor of English language and linguistics at the University of Southern Denmark. She has been investigating how people interact verbally with computers and robots for many years; some of this work is documented in her book *What Computer Talk Is and Isn't* (2006). Within the ITALK project, she works on how language can be bootstrapped from human-robot interaction. Her current work comprises an extensive comparison between child-, foreigner- and robot-directed speech, targeting the underlying processes and mechanisms in adapting one's speech to a particular communication partner. Fischer is also coeditor of the book series *Studies in Pragmatics* (Emerald) and of the journal *Pragmatics and Society* (John Benjamins).