



Social and collaborative aspects of interaction with a service robot

Kerstin Severinson-Eklundh*, Anders Green, Helge Hüttenrauch

*Interaction and Presentation Laboratory (IPLab), Department of Numerical Analysis and Computer Science,
Royal Institute of Technology, SE-100 44 Stockholm, Sweden*

Abstract

To an increasing extent, robots are being designed to become a part of the lives of ordinary people. This calls for new models of the interaction between humans and robots, taking advantage of human social and communicative skills. Furthermore, human–robot relationships must be understood in the context of use of robots, and based on empirical studies of humans and robots in real settings. This paper discusses social aspects of interaction with a service robot, departing from our experiences of designing a fetch-and-carry robot for motion-impaired users in an office environment. We present the motivations behind the design of the Cero robot, especially its communication paradigm. Finally, we discuss experiences from a recent usage study, and research issues emerging from this work. A conclusion is that addressing only the primary user in service robotics is unsatisfactory, and that the focus should be on the setting, activities and social interactions of the group of people where the robot is to be used. © 2003 Elsevier Science B.V. All rights reserved.

Keywords: Service robots; Human–robot interaction; Social robots; Speech interfaces

1. Introduction

To an increasing extent, robots are being designed to become a part of the lives of ordinary people. Their tasks may range from entertainment or play, to assisting humans with difficult or tedious tasks. In these kinds of applications, the robot will interact closely with a group of humans in their everyday environment. This means that it is essential to create models for natural and intuitive communication between humans and robots.

Here we will focus on the use of a mobile service robot, whose main purpose is to help one or more individuals, some of them with special needs, in a home or workplace. The robot will move around in an environment, inhabited by a group of people, and perform

physical tasks, such as to fetch and deliver objects. It may also keep track of various kinds of information for the users. The robot is partly autonomous, and it carries out its missions in the immediate, shared environment, interleaved with rather frequent user interactions. The users, on the other hand, are busy with their ordinary activities while communicating with the robot.

The analysis of the interaction between human and robot, and the models to be used in design, should be based on an understanding of the context where the robot is to be used. This includes consideration of the group of people involved, their goals and activities, as well as the shared physical environment. The “division of labour” between human and robot is rarely given in beforehand, but may vary depending on the context. Users may prefer to do certain tasks themselves while they need assistance with others [19]. In other cases, users may be expected to assist the robot on its missions to compensate for limitations of autonomy (“collaborative control” [10]).

* Corresponding author. Fax: +46-8-10-24-77.

E-mail addresses: kse@nada.kth.se (K. Severinson-Eklundh), green@nada.kth.se (A. Green), hehu@nada.kth.se (H. Hüttenrauch).

More than traditional computer applications, the use of intelligent robots encourages a view of the machine as partner in communication rather than as a tool. This suggests that people can be expected to apply more social models of interaction than in ordinary human–computer interaction. According to Reeves and Nass [24], people always take advantage of their social skills when communicating with a computer, but this element can be expected to be particularly strong in human–robot contexts. The design of the robot can, in different ways, encourage a kind of social interaction.¹

The social character of human–robot interaction does not imply, however, that the robot and the human are assigned equal or comparable roles. In service robot contexts, it is natural to view the robot as assisting the human user. Neither does it imply that the robot has to have a humanoid appearance. A survey carried out in our research group [17] showed that the majority of the participants preferred to view a service robot as a smart appliance, although other perspectives, such as a personal assistant, were also rated as acceptable. The survey included 134 people from various demographic groups, and with both men and women participating to the same degree. When asked about the degree of independence of a robot, the majority of subjects preferred a robot that does only what it has been instructed to do, and does not act independently.

Many research problems are emerging with regard to how people may interact with a service robot, and how such robots may be designed to be “socially interactive” [9]. The research that already exists is sparse, and has usually not focused on robots in real-world settings. This paper discusses social and collaborative aspects of interaction with a service robot, based on our experiences of a recent project involving the design and evaluation of the service robot Cero. We will focus on the following issues:

- (1) the *personality* of a robot, and how it may invite to interaction;
- (2) the *paradigm of communication* between user and robot, and the extent to which it is based on human communicative principles;

¹ According to a taxonomy introduced by Breazeal [3], a robot that uses human-like social cues and communication modalities in order to facilitate interactions with people is “socially communicative”. See also [9] for an extension of the taxonomy.

- (3) how a robot can *mediate collaboration and social interaction* within a group of people, based on experiences from a longitudinal study of usage of the Cero robot.

2. The Cero robot project

We have been engaged in a 3-year project with the purpose of developing a fetch-and-carry-robot (Cero) for motion-impaired users in an office environment. The work has been user-centered, and has involved several steps to capture user needs during the design process. These have included task analysis and interviews with potential users, and simulations of the robot’s functionality to study users’ reactions. Prototypes of the robot’s exterior design and interfaces have been designed and evaluated in an iterative fashion. Finally, a long-term study of the use of the robot has been carried out, assisting a motion-impaired user as part of a workgroup setting.²

2.1. Basic features of Cero

The Cero robot has a simple, fetch-and-carry functionality. It is based on a Nomadic Super Scout platform with 16 sonar sensors to support navigation. The present version of the robot (see Fig. 1) has a physical design made to support transportation of light objects, with a special compartment for the objects to be transported. The objects can also be placed on top of the lid of the compartment, depending on the type of the object (e.g. size, weight, whether it is a personal belonging or contains private information). The robot can be controlled through a graphical interface, running on a PC/Windows, and a speech interface. As a complement, the top of the robot includes a small artificial character, CERO,³ capable of issuing simple movements. There is also a PDA-enabled graphical

² The development process of Cero, including the various stages of capturing user requirements, was described in [14]. The longitudinal usage study was reported in [15].

³ CERO stands for Cooperative Embodied Robot Operator. Originally, the name was chosen for the physical interface character, but it is now also used for the robot itself. In this paper, we try to separate them by using capital letters when referring to the interface character.



Fig. 1. The Cero robot.

interface that has been developed for situations when users are on the move.

Early in the project, a simulation study was performed in the Wizard-of-Oz framework [7] to explore how people would relate to a robot in the targeted type of scenario. This gave important input to the dialogue design, and comprised an initial iteration of the robot's physical design. A general conclusion from the study was that users need explicit feedback at several levels to understand the state of the robot; in particular, to what degree it has perceived the instructions given, and where it is headed when moving around. The latter problem was partly related to the fact that the provisional robot prototype had a cylindrical shape with few explicit indications of its ways of operation (see Fig. 2⁴). This led to the idea of having a physical character as part of the interface, which would both give the robot a visible front, and provide conversational feedback with simple gestures, complementing the other interface modalities.

2.2. Modes of operation of the robot

In the fetch-and-carry scenario, the Cero robot has two basic modes of operation. The user normally initiates interaction by instructing the robot

⁴ Children did not take part as users in the study; this picture was taken during a pause in the simulation trials, when the robot attracted the interest of a visitor to the lab.



Fig. 2. Preliminary robot prototype used in simulation study.

to get or deliver an object, or to go to a specified location. This is an autonomous navigation mode, in which the robot fulfills its task by creating a high-level plan, and activating a set of behaviors (such as avoiding obstacles) to reach its goal destination. Strictly speaking, the robot is only partly autonomous, since fetching or delivering an object is sometimes dependent on collaboration with or between users. This issue is further discussed in Section 5 below.

A second mode of operation is “near-navigation”. During the user trials, it became apparent that the users frequently needed to alter the robot's location. For example, the robot could not enter a user's office because of high doorsteps. Therefore, she often wanted to move it as close as possible to the door to be able to roll over on her chair and meet it at the doorstep. We thus decided to provide means for near-navigation of the robot in two ways. In the PDA-enabled graphical interface, the robot can be directly controlled in three directions via a joystick button (see Fig. 3). A related feature has been created in the speech interface, in the form of commands for a directive dialogue. For example, a user can move the robot by issuing the instructions “turn left”, or “go forward one meter”.

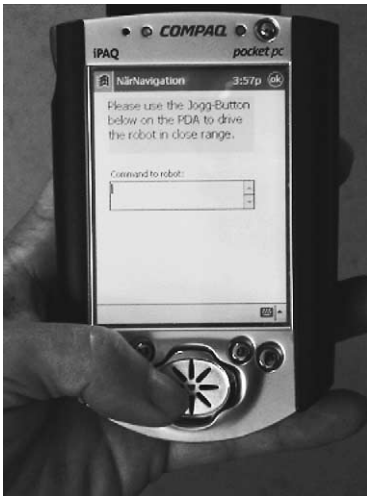


Fig. 3. Robot interface with direct control switch on a PDA.

3. Personality features of robots

An important aspect of a robot's way of communication with a user is the image of itself it projects. The survey mentioned earlier [17] showed that people's common images of robots are heavily influenced by science fiction robots in film and literature. However, when presented to the idea of a service robot in their own home, subjects gave priority to more practical characteristics, and preferred to view it as a smart appliance. Generally, there is a lack of useful models available for the design of service robots, which might give users a useful and adequate perception of them. In designing the interfaces of Cero, we wished to present the robot as intuitive, non-threatening and interesting to use. This led us to consider aspects of personality of robots in general.

People who encounter a mobile robot often attribute animate or human properties to it. This seems to be the case even when the robot has a modest construction and a simple way of operation. In particular, it appears that for people to perceive a robot has having animate qualities, the robot does not have to have a humanoid appearance. Our observations suggest that the mere fact that the robot moves around in a non-predictable way may cause people to interpret it as acting intentionally, and as having a kind of personality. This is a feature that designers can

take advantage of, to affect people's perception of the robot.

According to Norman [23], giving a robot a personality facilitates for users to understand its behavior. To be usable, a device must in general project a system image that can be used in developing a conceptual model of how it works [21]. Norman uses Sony's Aibo robot as an example: the fact that the robot has the personality of a puppy makes it natural that it sometimes does not understand or follow instructions.

It is important to design a robot's personality in ways that match its purpose. In an experiment conducted by Goetz and Kiesler [12], subjects were given instructions for a physical training pass, led by a robot with two different communication styles. The same robot was used in both conditions. In one condition, the robot had a cheerful behavior, made jokes etc. during the sessions, and treated the task as entertaining. In the other condition, the robot was serious, and mentioned health aspects of the task. The results showed that the participants liked the happy robot more, but that they followed the robot's instructions to a greater extent in the condition with the serious robot. It is conceivable that a robot's personality would contribute to a user's sense of trust in its instructions. In any case, a personality that is suited for one set of tasks may be inappropriate in another set of tasks.

What characteristics of a robot contribute to its personality? Evidently, there are numerous qualities that may affect the perception of a robot. Apart from the explicitly humanoid characteristics that a robot may have, it seems clear that its size, shape, colour, and the character of its movements may be important. In an experimental study performed by Butler and Agah [5], the speed of a robot as well as its way of approaching a human affected people's perceptions of it. Subjects were most comfortable with a slower robot, or a robot approaching indirectly rather than directly. Moreover, their reactions were affected by the robot's exterior shape: the level of discomfort perceived during a fast approach increased if the robot had a tall and humanoid body rather than a short, neutral cylindrical shape. Smooth movements were preferred over jerky ones, and the distance to the robot was also a factor contributing to the subjects' perception of comfort. Adding a tall, humanoid robot body attachment caused comfort levels to drop in all cases studied.




Fig. 4. The embodied interface character, CERO.

3.1. Designing Cero's personality and behavior

The CERO character can be seen as a representative for the robot, which is also a natural communication partner. Alternatively it can be seen as the “driver” of the robot. The robot itself is not humanoid, but the CERO figure has some simple human traits, with a head, arms and a body (see Fig. 4). We have developed a set of gestures that CERO is able to use in interaction. For example, it can nod to give conversational feedback, and move its arms as if it was walking, showing that the robot is on its way (see also Fig. 5).

The CERO figure is thought to contribute to the robot's affordances [11,22]. The idea is that it should encourage the use of human communication patterns. The same kind of character could be used in interfaces of other applications than robots. Green [13] has developed ideas of how a family of artificial characters (“Ceroids”) could be part of the speech interfaces to a set of home appliances, suggesting to the user that they have the same communicative paradigm.




Fig. 5. CERO's 4 degrees of freedom.

An animated figure can, through its gestures, express personal and emotional attitudes. It can be portrayed as shy, withdrawn, modest and so on, and thereby evoke reactions and associations that make people attribute a personality to it, and that affect how they engage in interaction. So far we have only worked with gestures that express the state of the robot and the human–robot dialogue, but emotional gestures are also of interest.

Our work with the CERO figure is related to research by Mizoguchi et al. [20]. They developed an expressive robot that uses gestures and movements to create a sense of familiarity to users. The purpose is that such a robot could be a part of interface to a range of devices. Their robot is technically more advanced than CERO, it is slightly larger (27 cm) and has 7 degrees of freedom (CERO has 4; see Fig. 5), which gives many parameters to control and therefore potential complexity in designing movements. An interesting feature of their work is that they designed a set of complex expressions, similar to “poses” in ballet, using parameter settings that were assessed with human subjects in experimental studies. They have also studied the effects of velocity parameters of gestures, to find settings that are experienced as comfortable to users.

4. Communicative principles guiding Cero's dialogue design

To what extent spoken dialogue interfaces for robots may be based on models of human conversation, and in what ways this might have a beneficial effect on

the design of service robots are still issues for ongoing research. In the work with the spoken dialogue interface for Cero, we have focused on providing a dialogue that takes human dialogue strategies into account. Rather than using a strictly command-oriented approach, where a natural language command is reactively followed by an action, we have developed a dialogue model based on the principles of grounding in human-to-human dialogue.

In human-to-human dialogue, the participants are engaged in a joint cooperative behavior to achieve a common goal [1,4,6]. When two participants are coordinating their actions using natural language, they try to reach a common ground relevant for the particular task at hand. This is referred to as *grounding* and is achieved by exchanging pieces of information needed to establish mutual knowledge, beliefs and suppositions [6].

At the level of task-specification, we use a *cautious* grounding strategy aimed at assuring that the user is certain about what instructions the robot has received and is about to carry out. This means that the robot acknowledges the user's request by reformulating it as a question, requesting confirmation by the user. This may in turn be confirmed as requested or rejected, the latter meaning that the robot cancels the mission in progress.

In the following example, the robot receives a command that is only partially understood by the dialogue system. The following turn by the user is aimed at providing more information for the system. If the user responds by specifying a location in the next turn, the system may infer that the task of getting and the object to be collected are part of common ground:

- U Robot, get a paper from John's office
 [get paper].
 R Get a paper from where? [requesting
 specification of location].

The directive instructions used for near-navigation; e.g. "turn left" are not grounded to the same degree as the fetch-and-carry tasks. Here, an *optimistic* strategy is used so that only the first directive command given by the user is grounded by asking for confirmation. If a new directive command is issued, it is carried out reactively, assuming that the user is in close vicinity of the robot and is monitoring its movements.

Conversational feedback of different kinds is one of the most important resources for grounding. Speech and body gestures can be viewed as the primary modes of production of feedback. Gestures may reinforce speech by introducing redundancy; they may also add information by expressing emotions and attitudes, or providing illustration of verbal content. Gestures also have an important role in communication management (e.g. nodding to show the current message has been understood).

At lower levels, the gestures of the CERO character provide simultaneous feedback of the robot's status, displaying information about the perceptual status of the system and information of low-level processing of natural language such as speech recognition and sentence parsing. The CERO character is interleaved with the speech system so that it is both capable of issuing conversational gestures reactively based on system states (e.g. raised amplitude, spoken vowels) and conventional gestures (e.g. emblems such as to nod or shake its head).

Brennan and Hulstén [2] proposed a taxonomy of eight categories or levels of feedback which correspond to different depths of grounding. Fig. 6 shows how these levels correspond to system states, with feedback on the various levels given through gestures and speech.

Level 0 concerns whether the system is active. The next level concerns if the user has the system's attention, or the system has detected a human voice. If the speech recognition system provides incremental recognition of commands, the next feedback level concerns partial results issued by the system (this may be expressed by small head nods issued by the CERO character). The categories three and four concern natural language processing and may produce different types of responses of the system (e.g. a low confidence score might trigger a request for the user to adjust the microphone, or asking the user to repeat the command). The fifth category of feedback is at the task level, e.g. asking the user to confirm an action. The sixth category concerns actions and is typically signalled by the movement of the robot itself. The last feedback category reports the current actions of the system (e.g. "Going to the kitchen").

When designing the feedback system of Cero's speech interface, we have not attempted to construct

Category	System event/state	Feedback gesture	Spoken Feedback
0. Not attending	– Robot off	<i>n/a</i>	<i>n/a</i>
1. Attending	– Robot on – Microphone sound detected – Speech detected	Positive: – raise head	<i>n/a</i>
2. Hearing	– Low confidence score for speech recognition – Partial commands	Negative: – Shake head Positive: – Nod	Negative: – “Repeat the last command”
3-4. Parsing/Interpreting	– Parsing errors and inconsistent commands	Negative: – Shake head	Negative: – “Please rephrase”
5. Intending	– Planner failures	Negative: – Shake head	Negative: – “Unable to perform X”
6. Acting	– Execution of tasks	Positive: – Nod	Positive: – “Go to X?”
7. Reporting	– Reporting on task execution	Positive: – Walking gesture using both arms	Positive: – “Going to X”

Fig. 6. Feedback levels according to Brennan and Hulteen, and examples of their realizations in the Cero robot system.

gestures corresponding to all eight levels. Rather, the animated CERO figure is intended to give feedback for which the speech synthesis is not adequate or sufficient, e.g. by displaying that speech has been detected or showing that the system is switched on.

There are several remaining issues regarding the speech interface that make longitudinal studies with an onboard system challenging. The availability of robust technology for far-field speech recognition should be acknowledged as one of the most important ones.

4.1. Consistency between interfaces

Consistency is one of the most important properties of a computer system with regard to learnability (see e.g. [8]). Designing for consistency means to provide users with cues that help them predict the system’s way of working, by using the same behavior in different parts of the system. In the design of Cero’s interfaces, we have strived for consistency between the graphical interface and the spoken interface. This is done in two ways: first, the graphical interface

has an animated CERO figure in the upper-left corner. Its movements provide feedback on the system state; e.g. when the robot is moving, the figure waves its arms. Secondly, the vocabulary of the graphical interface has been chosen in such a way that word choice and command style coincide as much as possible with the speech interface. This helps users learn the speech interface, by providing them with some words and phrases that can be used in commands.

5. Social and collaborative interaction in a workgroup setting

People who use a service robot in a workplace or home already have a social environment and an established way of cooperation in which the robot is to be included. Regardless of what are the robot’s main tasks, or the number of people using it, this overall social context is likely to influence its use. Designing interaction for a service robot therefore requires an understanding of the variety of activities going on in the shared environment.

5.1. Longitudinal usage study

The Cero robot has been evaluated in a 3-month field study,⁵ where one user from the target group of motion-impaired people used the robot as a part of her work. The aim of the study was to assess the long-term effects of a service robot, beyond the limited experiments and user evaluations already conducted in the project. In particular, we were interested in collecting data about the use (or non-use) of the robot during a regular workday, gather feedback from the user regarding the robot's performance and its interaction design, observe reactions from co-workers of the user as well as visitors to the work environment, and explore the setting for unanticipated events, to inform further design as well as human–robot interaction research in general (see [16] for a more comprehensive description of the study).

The study was conducted in an office workplace consisting of two adjoining corridors of about 70 m in total length, hosting a workgroup of about 30 people. Indirectly, another workgroup located at the end of one of the corridors was also affected by the robot usage.

Data was collected in several forms, including: (a) video recordings of four start-up sessions (about 1.5–2 h each); (b) direct observation in the environment throughout the study; (c) annotations about critical events, requests for help, and other comments made by the user directly to the researchers or via e-mail during the study; (d) logs from the system, monitoring the robot-system's internal state and all actions taken in its interfaces; and finally (e) a post-trial in-depth interview of 1.5 h with the subject 2 weeks after completion of the long-term study.

The main focus was upon one single user, although within the context of the entire workgroup setting. The

user was a middle-aged female academic, who had a neurological handicap resulting in a walking disability. She could walk short distances with two crutches, but was not able to carry any objects while walking. The robot was intended to assist with the transportation of objects such as books, papers, coffee cups etc. in the immediate work environment.

During the trial, the robot was started every morning during weekdays, and placed at a standby position close to the user's workplace. She could then use it whenever she needed it, at her own discretion. In contrast to a laboratory experiment, the complexities and situated nature of interacting with the robot were intensively studied at different levels. For example, the user was frequently observed attending to the robot's graphical interface while in parallel interacting with other people or doing other tasks.

When the user sent the robot to fetch or deliver something, she either coordinated it with someone at the destination, or accompanied the robot there herself. For example, to get coffee, she would send the robot to the kitchen, and then (perhaps somewhat later) go there to make sure there was coffee available, pour a cup, and put it on the robot's tray. If someone else was in the kitchen, they would sometimes offer to help with this. There was no possibility for the robot to do these tasks, because it lacked a grasping ability. Neither could it discover the presence of a person at the destination.

During the course of the study, the robot was driven for a total of 108 h, covering a distance of more than 14 km in the office environment. The system log registered 423 missions, 39% of which were transportations between the user's office and the kitchen. In the post-trial interview, the user confirmed that carrying coffee cups and glasses of water etc. was an important need of hers. This use of the robot was seen as a significant improvement over asking other people to do such tasks for her, providing the possibility to be independent of other people's assistance.

Through the data collected during the study, more than 100 different design issues were brought up, some cross-related to each other. These issues were annotated with regard to the context of their occurrence and categorized along different dimensions. Through this careful documentation procedure, the field notes had the important function of reporting back into the system design, and to bring forward and relate issues that

⁵ We are aware that the meaning of the term "field study" is somewhat unclear when applied to the use of mobile, embodied (social) robots. The term is chosen here to emphasize that the study was carried out in a natural office setting, as contrasted with experimental or laboratory conditions. Further, the study used a predominantly qualitative approach, where understanding is reached through interpretation of documented phenomena and observed actions of the participants involved. The nature of the data collected reflects a sensitivity to the context in which the investigation is conducted, and the analysis takes account of the richness, detail and complexity of the social setting.

were previously unknown in relation to human–robot interaction in an office setting.

5.2. Design aspects related to social interaction in a workgroup setting

During the longitudinal study, many unexpected events were observed that were a result of the various activities and social interactions going on in a workplace, and that made us reconsider details in the design of the robot. In particular, there were situations in which other people than the targeted user(s) attempted to interact with the robot, or needed to interact with it but did not know how to do it. Some of these situations were not prepared for in the interface, which was oriented towards the individual user and concentrated on the robot's missions.

While using the robot during daytime, the user in the field study was normally doing her usual work tasks. While the robot went on its missions, the user would continue writing, talking on the phone, or having meetings. As a consequence, she forgot several times that she had sent the robot on a mission. The robot was waiting at the destination, positioned in the middle of the corridor, but the user was unaware of this and continued with her work. Usually, this situation was not a problem, since people could pass the robot easily, and sometimes even found it amusing to find it there. However, there was one exception: the cleaning personnel could not pass the robot with their trolley; neither could they move the robot because they were not aware of how to operate it. In these and similar situations, someone had to alert the technical staff or the primary user who had to send the robot back again.

When the robot was navigating through the corridor, people who were passing by often tried to address it. They asked for its name, tried to put hats on it etc., and in some cases even obstructed its way (this would typically happen with student visitors). Generally, the robot attracted people's attention and was frequently the object of spontaneous conversations (as illustrated in Fig. 7). The robot, on the other hand, normally slowed down or stopped when encountering a person; this obstacle-avoiding behavior seemed to encourage people to interact with it. Similar things often happened also when the primary user accompanied the robot during a mission, resulting in frequent stops with conversations about the robot during a mission.



Fig. 7. Bystanders with the Cero robot in the office environment of the field study.

As a consequence of these observations related to the workgroup setting, we found that at least the following three types of situations need to be accounted for in the interface:

- (1) While the robot is performing a mission, situations may occur when the robot must be moved temporarily for reasons of safety or maintenance. There should be easy and transparent ways for authorised staff to operate the robot in such cases.
- (2) When the robot is about to fetch or deliver something, and the user cannot follow it, another person must help at the destination by placing objects on the robot's transportation tray. This means that several people need to cooperate and coordinate their actions within a mission. We may distinguish between the *primary user*, who initiates the mission and defines its overall goal, and *secondary user*, who is aware of the primary user's needs and agrees to help at the destination. There may be several people acting as secondary users. All of them must at least have the ability to assess the robot's state, to move it away from places not allowed, and to send it back to its standby location when the objects have been supplied.

The dialogue design could also be adapted so that the robot can ask if other users are present, and address them with clear instructions about how to help in completing the mission. We have prepared

for such collaborative missions in a modified interface of Cero, so that the robot can instruct a secondary user to help with a fetch or deliver task, using synthetic speech. In addition, a hardware switch has been placed on the robot's cover which can be used to send the robot to its final destination, after the fetch or delivery task has been carried out.

- (3) During interaction and navigation in the environment, there are other people crossing the robot's way, as mentioned above. Some of them are staff, for whom the robot may be familiar. Others are visitors, who may need information about the robot and its purpose. We call these users *bystanders*. For this category of users, a more social dialogue would be of interest, on top of the present task-oriented dialogue interface. This means that the robot would be able to e.g. respond to greetings and engage in a conversation about itself, the schedule of classes or seminars, or other current information that may interest visitors.

6. Issues for future research

The research about social qualities of human–robot interaction is still in its beginning. Several influential and impressive experimental systems have been built that provide examples for design, and help raise crucial questions. However, there is still little experience showing how people interact with and relate to robots in real-life settings. In this paper, we have focused on service robots, and how they can be designed to communicate and collaborate with users in a transparent and intuitive way. To advance knowledge in this area, there is a need for more extensive empirical study of robot use in real-life contexts, including efforts to create new concepts and models that work across different applications.

With regard to the personality of robots, there have been attempts to create a methodological arsenal to study people's perceptions of robots in experimental studies [18]. These will be helpful to study people's impressions of existing robots on a larger scale. However, it is also important to use more contextual approaches, by examining systematically how people perceive and relate to robots in their own natural environment, and how this relationship evolves over time,

to get richer information useful for design. These studies are likely to be more qualitative in nature, ranging from single case studies to more ambitious field studies of the use of service robots.

The dialogue between user and robot has already been studied by many researchers. Multimodal dialogue systems have been designed, but the state of knowledge in this field is still fragmented. The expectations on these systems have perhaps been too high, and we need rich empirical experience of the use of robust and simple systems in order to formulate new and relevant questions.

Finally, there are several issues pertaining to group collaboration and awareness that need to be studied with consideration of human–robot relationships. How do people view a robot as an embodied artifact, and how can it mediate their own activities in a group or organisation? How may the presence of one or more robots contribute to and shape their activity over time? These are broad questions that we cannot expect to find answers to within a near future. But as robots enter into the environments of more people, these questions will be gradually specified and clarified to provide a more solid basis for empirical study as well as design.

Acknowledgements

Erik Espmark created the physical design of the robot, including the CERO figure, and gave invaluable input during the initial stages of the project. We are also indebted to our system engineer Mikael Norman and to members of the Centre for Autonomous Systems, in particular Henrik Christensen, who gave us access to their sonar navigation system and provided technical assistance during the project.

The work described here has been funded by the Swedish Labour Market Board (AMS), the Swedish Foundation for Strategic Research (SSF), and the Swedish Agency for Innovation Systems (Vinnova).

References

- [1] J. Allwood, J. Nivre, E. Ahlsén, On the semantics and pragmatics of linguistic feedback, Technical Report 64, Gothenburg Papers on Theoretical Linguistics, Gothenburg University, 1991.

- [2] S.E. Brennan, E. Hulstén, Interaction and feedback in a spoken language system: a theoretical framework, *Knowledge-based Systems* 8 (1995) 143–151.
- [3] C. Breazeal, Toward sociable robots, *Robotics and Autonomous Systems* 42 (3–4) (2003) 167–175 (this issue).
- [4] H.C. Bunt, Dynamic interpretation and dialogue theory, in: M.M. Taylor, D.G. Bouwhuis, F. Neel (Eds.), *The Structure of Multimodal Dialogue*, vol. 2, John Benjamins, 1999.
- [5] J.T. Butler, A. Agah, Psychological effects of behavior patterns of a mobile personal robot, *Autonomous Robots* 10 (2001) 185–202.
- [6] H. Clark, *Using Language*. Cambridge University Press, Cambridge, 1996.
- [7] N. Dahlbäck, A. Jönsson, L. Ahrenberg, Wizard of Oz studies—why and how, *Knowledge-based Systems* 6 (4) (1993) 256–258.
- [8] A. Dix, J. Finlay, G. Abowd, R. Beale, *Human–Computer Interaction*, Prentice-Hall, Hemel Hempstead, Hertfordshire, 1993.
- [9] T. Fong, I. Nourbakhsh, K. Dautenhahn, A survey of socially interactive robots, *Robotics and Autonomous Systems* 42 (3–4) (2003) 143–166 (this issue).
- [10] T. Fong, C. Thorpe, C. Baur, Robot, asker of questions, *Robotics and Autonomous Systems* 42 (3–4) (2003) 235–243 (this issue).
- [11] W.W. Gaver, Technology affordances: use of familiar things in the design of interfaces, in: *Proceedings of the ACM CHI'91 Conference on Human Factors in Computing Systems*, 1991, pp. 79–84.
- [12] J. Goetz, S. Kiesler, Cooperation with a robotic assistant, in: *Proceedings of the ACM CHI'02 Conference on Human Factors in Computing Systems*, ACM Press, New York, 2002, pp. 578–579.
- [13] A. Green, C-Roids: life-like characters for situated natural language user interfaces, in: *Proceedings of the 10th IEEE International Workshop on Robot and Human Communication*, IEEE, Bordeaux/Paris, France, 2001, pp. 140–145.
- [14] A. Green, K. Severinson Eklundh, Task-oriented dialogue for CERO: a user-centered approach, in: *Proceedings of the 10th IEEE International Workshop on Robot and Human Communication*, IEEE, Bordeaux/Paris, France, 2001, pp. 146–151.
- [15] H. Hüttenrauch, A. Green, K. Severinson Eklundh, M. Norman, L. Oestreicher, Involving users in the design of a mobile office robot, *IEEE Transactions on Systems, Man and Cybernetics C*, to appear.
- [16] H. Hüttenrauch, K. Severinson Eklundh, Fetch-and-carry with Cero: observations from a long-term service robot trial, in: *Proceedings of the 11th IEEE International Workshop on Robot and Human Communication*, Berlin, Germany, September 2002.
- [17] Z. Khan, Attitudes towards Intelligent Service Robots, Technical Report IPLab-154, TRITA-NA-P9821, Royal Institute of Technology, Stockholm, 1998.
- [18] S. Kiesler, J. Goetz, Mental models of robotic assistants, in: *Proceedings of the ACM CHI'02 Conference on Human Factors in Computing Systems*, ACM Press, New York, 2002, pp. 576–577.
- [19] C. Laschi, G. Teti, G. Tamburrini, E. Datteri, P. Dario, Adaptable semi-autonomy in personal robots, in: *Proceedings of the 10th IEEE International Workshop on Robot and Human Communication (Roman'01)*, Bordeaux-Paris, 2001.
- [20] H. Mizoguchi, T. Sato, K. Takagi, M. Nakao, Y. Hatamura, Realization of expressive mobile robot, in: *Proceedings of the 1997 IEEE International Conference on Robotics and Automation (ICRA)*, Albuquerque, New Mexico, 1997.
- [21] D.A. Norman, Cognitive engineering, in: D.A. Norman, S.W. Draper (Eds.), *User Centered System Design*, Lawrence Erlbaum, Hillsdale, NJ, 1986, pp. 31–61.
- [22] D.A. Norman, *The Design of Everyday Things*, Currency/Doubleday, London, 1990.
- [23] D.A. Norman, How might humans interact with robots? Human Robot Interaction and the Laws Of Robotology, Notes for a Keynote Address to the DARPA/NSF Workshop on Human–Robot Interaction, San Luis Obispo, CA, September 2001. http://www.jnd.org/dn.mss/Humans_and_Robots.html (July 2002).
- [24] B. Reeves, C. Nass, *The Media Equation: How People Treat Computers, Television and New Media Like Real People and Places*, Cambridge University Press, Cambridge, 1996.



Kerstin Severinson-Eklundh is a professor of human–computer interaction at the Royal Institute of Technology (KTH) in Stockholm. She achieved her Ph.D. in Communication Studies at the University of Linköping in 1986. She is head of the Interaction and Presentation Laboratory (IPLab), an interdisciplinary environment for human–computer interaction research and education. One of her current

research projects is “Human interaction with intelligent service robots”, which is a collaboration with the Centre for Autonomous Systems. Other research areas at IPLab include computer-supported cooperative work, language technology, and computer-assisted writing.



Anders Green is a Ph.D. student of human–machine interaction at IPLab, the Royal Institute of Technology (KTH) in Stockholm. His research is focused on human–robot interaction using multi-sensory natural language user interfaces. He holds a Master of Arts in Computational Linguistics from Göteborg University, received in 1997. He participates in the Graduate School of Human–Machine Interaction and the Swedish National Graduate School of Natural Language Technology (GSLT).



Helge Hüttenrauch is a Ph.D. student in the human–robot interaction research group at IPLab, the Royal Institute of Technology (KTH) in Stockholm. His research interests concern usability and interface design for service robots. After receiving his diploma in computer science and media from the University of Applied Science in Furtwangen, Germany in 1995, he worked with Ericsson Business

Networks AB as a usability engineer for computer telephony applications until 1998. He is currently affiliated with the Graduate School for Human–Machine Interaction.