Robots in the Wild: Observing Human-Robot Social Interaction Outside the Lab

Selma Sabanovic
Science and Technology Studies
Rensselaer Polytechnic Institute
Troy, New York 12180
Email: sabans@rpi.edu

Marek P. Michalowski
Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
Email: michalowski@cmu.edu

Reid Simmons
Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
Email: reids@cs.cmu.edu

Abstract— **This paper discusses the use of observational studies of human-robot social interaction in open human-inhabited environments as a method for improving on the design and evaluating the interactive capabilities of social robots. First, we discuss issues that have surfaced in attempts to evaluate social interactions between humans and robots. Next, we review two observational studies involving robots interacting socially with humans and discuss how the results can be applied to improving robot design. The first is an analysis of a mobile conference-attending robot that performed a search task by augmenting its perception through social interaction with human attendees. The second is an analysis of a stationary robotic receptionist that provides information to visitors and enhances interaction through story-telling. Through these examples, we show how observational studies can be applied to human-robot social interactions in varying contexts and with differing tasks to quantitatively and qualitatively evaluate (and discover unanticipated aspects of) the social interaction. Finally, we discuss design recommendations suggested by insights gained through these analyses.**

I. Introduction

Assessment of the sociality and interactive functionality of different robots is a challenging issue in human-robot interaction and social robot design. Social robotics projects vary greatly in their stated scientific, technical, and social goals. Some researchers aim to acquire new knowledge about human sociality through the simulation of human behavior [1], [2]. Others seek to improve the quality of human-machine interaction by creating interfaces that will rely on social cues and therefore be more natural, intuitive and familiar for users [3]. Various methods of using the psychological, physiological and social effects of robots on humans for therapeutic purposes are also being investigated [4], [5]. Due to the variety of design goals and methodologies in social robotics, it is difficult to compare human-robot social interaction across projects and to present results that can be generalized from one social robot design to another [6].

We argue that observation and behavioral analysis of the social interactions between robots and people situated in real-world contexts is necessary for considering and reconciling the various social, scientific, and technical concerns relevant to social robot design. Unlike industrial and laboratory robots, which perform their tasks in closed and well-defined work environments populated by trained professionals, social robots are meant to function as part of the everyday life of a more generalized public. Roboticists are aware that systems that work well in the lab are often less successful in noisier field environments, which makes testing outside the laboratory a necessary step in evaluating robots with non-laboratory applications [7]. Interactions with robots in the laboratory, under the watchful eye and expert guidance of the robot's designers, do not provide insights into the aspects of human-robot interaction that emerge in the less structured real-world social settings in which they are meant to function. It is therefore necessary to evaluate human-robot interactions as socio-culturally constituted activities outside the laboratory, or "in the wild" [8].

Evaluations of the performance of social robots should also consider the direct (intended as well as unintended) impact on humans that interact with them [2]. As social robots become available to the general public, it is increasingly important to study the nature of the robots' interactions as well as their effect on people interacting with them outside the lab. In the case of the elder therapy robot PARO, which is being sold in Japan, long-term experiments have been performed with groups of nursing home residents [4]. It is not clear how individuals will be interacting with the robot in their homes without the constant attention and scaffolding[1] provided by roboticists and nursing staff. These effects can be investigated using observational ethnographic and behavioral research in situated interaction settings.

We discuss issues relevant to current methods of evaluating social robots and review a number of projects that have used observational analysis. We present our own observational analysis of two social robots interacting with people in open real-world settings. We then discuss the analysis in terms of its relevance to testing design assumptions and improving robot design.

II. Issues in Social Robot Evaluation

A common theme in discussions of social robot designs and their success as applications refers to the "humanlike" characteristics (appearance, emotions, personality) and skills (speech, gaze, gestures) of the robot [9], [10]. Each social robotics project is based, explicitly or implicitly, on a model of appropriate "humanlike" behavior constructed and shared by the roboticists. A focus on the individual robot's inherent "humanlike" abilities is an inconclusive approach to designing for human-robot social interaction, as the appropriate behavior is dependent both on the social environment and the participants in the interaction. These models often exclude aspects of human-ness that roboticists find uncomfortable or undesirable to the user [2]. We propose that a focus on analyzing the

[1]Scaffolding here refers to speech or actions that provide a supporting framework for interactions with the robot.

0-7803-9511-5/06/$20.00 ©2006 IEEE

robot's capabilities as being constituted and situated through interaction in context would result in more apposite measures of a robot's sociality that can inform robot design.

Social robots are most often evaluated in the same environment in which they are developed – the laboratory. When the designers and roboticists, rather than informed outsiders, are also the evaluators of a robot's performance, the result is an internalist view of the design that makes it difficult to refute or revise its original premises. The robot's social capabilities are judged using the analytical tools and assumptions about sociality that evolved in conjunction with its design [11]. These design assumptions are often based on the roboticists' own "conscious models" [12] regarding human social behavior that fail to take into account possible differences between commonsense narratives about social interaction and the social patterns that become visible through detailed study. Internalist laboratory evaluations of the robot's social interaction also obscure the elaborate techniques for scaffolding the robot developed by the scientists working on it daily. The ultimate test of the robot's capabilities lies in an open social environment in which it must work constantly and autonomously.

The hybrid social and technical agenda of social robotics projects raises the issue of the appropriateness of quantitative or qualitative methods for assessing a robot's performance [13]. The robotics community conventionally applies quantitative measures such as speed or turns to task completion to judge a robot's performance and compare among robot designs and architectures. The contextual nature of social interaction precludes reliance on quantitative measures of human-robot social interaction alone, as this would lead to the systematic exclusion of evaluations of phenomena that are not easily amenable to quantification. Reliance on quantitative metrics and controlled experiments alone has limited utility for understanding social interactions in contexts where the task boundaries and success criteria are not clearly defined. The example of the robot GRACE at AAAI05, to be discussed in more detail in Section IV, portrays a case in which quantitative metrics such as the time it takes the robot to complete its task are less relevant than its ability to interact with conference attendees.

It is necessary to evaluate and compare social robots by submitting their interactions with humans outside the laboratory to detailed observation and analytical scrutiny. This approach is commonly used in psychology, ethology, and sociology to study observable behavior such as activities, postures, gestures, facial expressions, movements, and social or human-system interactions. It should be performed by social scientists trained in observational analysis [12]. Both the relocation of the robot outside the lab and the distancing of the observer from the everyday tasks of programming and maintaining the social robot make for a more valid test of the robustness of the robot's interactive capabilities.

III. Systematic observation and analysis

Observational behavioral analysis has been used within the laboratory to study the efficacy of certain robot behaviors and people's reactions to them. Watanabe et al. [14] use behavioral analysis of human nonverbal interaction both to develop their InterRobot speech system and to evaluate the ease of resulting human-robot interactions. Breazeal et al. [15] employ behavioral analysis of task-oriented interactions between humans and the robot Leonardo to demonstrate the salience of nonverbal cues in cooperative task-oriented interactions between humans and robots. Kanda et al. [16] conducted two-week-long experiments with Robovie in a school setting and recorded video of the robot's performance, yet the analysis of the robot's performance relied on the robot logs. Further examples of observational research can be found in Urban Search And Rescue (USAR) projects (e.g. [17]). Within the field of social robotics, Dautenhahn and Werry [18] discuss the utility of behavioral observation performed in a controlled setting combined with statistical analysis for the assessment of a robot's effects on autistic children.

We suggest that an effective and appropriate approach to the evaluation of human-robot social interaction is to analyze it as a situated activity performed in the context of particular concrete circumstances [19]. This approach can be related to the concept of "situated robotics" [20], which is based on the notion that complex environments have a strong influence on the nature and complexity of the behavior of embodied robots that are embedded in them. In the case of socially situated robotics, the organization of situated action is emergent from the interaction among actors and between actors and their social and physical environments. The only way these emergent capabilities can be evaluated is to take the robots outside the scripted laboratory setting and to engage them in everyday action in human social contexts. Through understanding how social robots are situated in their social and physical environments, it will be possible to design "socially embedded" robots [21] that are structurally coupled with their surroundings.

Observations of human-robot interactions in naturally occurring social situations can be used to create new theoretical and practical models of appropriate social robot behavior and design. Along with the promise of novel technical challenges that can lead to developments in and beyond conventional approaches to robotics, social robotics has potential for advancing our knowledge about human sociality by observing how people explore and interact with social robots in non-laboratory social environments. As opposed to the scientist in the laboratory, the person interacting with the social robot in a museum, school, or hospital will often not be aware of its technical characteristics and limitations; they will be noticing the robot's behavior and creating or responding to social cues, affordances, and scaffolding specific to the interaction at hand. Focusing our attention on human-robot interactions in dynamic open environments will enable us to evaluate and design socially responsible and responsive robots that fit into an ongoing flow of coordination with humans in their everyday environment.

In the case of robots for which social human-robot interaction plays a focal role [6], fine-grained observational analysis of the robot interacting in a real-world environment can be used to analyze how humans react to and interact with the robot; how humans interact with each other while interacting with the robot; which aspects of the robot's, and human's, actions lead to breakdowns in the interaction; and how the robot succeeds and fails to engage humans in interaction. These analyses can provide detailed quantitative and qualitative data that can be used to improve socially situated/embedded robot interactions through iterative design

Fig. 1. The social robots GRACE and the Roboceptionist

processes.

IV. Observation of social robot open interaction

We observed interactions between the Carnegie Mellon University robot GRACE (Graduate Robot Attending a ConferencE) and conference attendees at AAAI05 in Pittsburgh, PA and between the Roboceptionist [22] and Robotics Institute students, staff, and visitors (both shown in Fig. 1). Our discussion focuses on showing how observational analysis can be used to understand situated interaction between people and robots, to reveal factors that surpassed or challenged the initial design assumptions about social interaction, to suggest changes in the robot's design, and to relate findings to more general applications in social robotics.

A. Robot design and assumptions

GRACE and the Roboceptionist are similar with respect to hardware but differ in their interactive abilities and tasks. The robots are RWI mobile bases with flat-panel LCD monitors that display animated faces. Both robots sense humans in the environment using SICK laser scanners and single PTZ cameras. The variations in their interactive capabilities influence their interactions by defining the ways the robots can act on their social and physical environment. The Roboceptionist is stationary, although the head (screen) can turn, while GRACE is fully mobile. GRACE has a touch-screen interface that people can use to give her directions and answer her questions, while the Roboceptionist relies on typed questions and comments from its interaction partners. While GRACE's verbal communication capabilities were limited to greeting and requesting directions, the Roboceptionist uses a more sophisticated natural-language system that can answer questions about the weather and office locations, present building maps, and relate stories.

GRACE and the Roboceptionist are situated in environments that are populated by humans and are therefore social as well as physical [19]. The robots' social interaction models were designed according to the team members' experientially based ideas about natural and entertaining ways to initiate and maintain the appropriate interactions. At the AAAI05 Robot Exhibition, GRACE's task was to locate a team member wearing a pink hat by interacting with conference participants and asking them for directions [23]. The Roboceptionist is located at the entrance to Carnegie Mellon's Newell-Simon Hall and can respond to visitors' questions while using storytelling to foster long-term repeated interactions with visitors. GRACE's target interactions were shorter – she had only to interact briefly until she received directions.

B. Data collection and analysis

We used similar data collection and analysis methods in both cases. The interactions between robots and voluntary participants were recorded by a handheld camera (in the case of GRACE at AAAI05) and by two security cameras mounted on the booth (in the case of the Roboceptionist) (Fig. 2). For GRACE, we analyzed 3.6 hours of video taken over the course of three days at the conference. We obtained data about the Roboceptionist's social interactions from 12.51 hours of video recorded over a period of two days.

The videotaped data was coded using behavioral analysis software [24]. Manual codes were temporally applied to videotaped micro-behaviors such as utterance, spatial movement, gesture, and gaze as performed by the robots, the people who interacted with them directly, and those who were in close proximity but did not interact directly (i.e. passed by the robots). In the Roboceptionist's case, coded logs obtained through analysis of the videotaped data were aligned and merged with logs of the robot's internal state and behaviors. The data was subjected to statistical analysis to describe the frequency of various types of events, while lag-sequential analyses were used to determine the incidence of certain events being preceded or followed by others.

Fig. 2. Sample frames from videos of GRACE and the Roboceptionist.

C. Observational Results

Observational analysis focusing on the socially situated nature of the robots' tasks and designs revealed notable influences of the spatial and social environment on social human-robot interaction. We discuss interaction spaces, group interaction, interpersonal interaction, and rhythmicity as salient factors that ought to be considered when designing for human-robot interaction.

1) Spaces of interaction: In contrast to the original design assumptions, our analyses show that the different physical and social qualities of the spaces in which the robots interacted had significant effects on their performance. With GRACE, we identified three categories that varied in spatial configuration and social use: the reception, a social event held in a large hall in which people were contained and crowded (31.25 minutes recorded); the hallway, a place through which people walked on their way to the various conference presentations (104 min.); and the banquet, a social gathering held in the hallway (82.4 min.).

Lag sequential analyses were used to study how participants responded to GRACE within 5 seconds of her attempts to engage (move toward or speak to person) and disengage from an interaction (move or turn away), and when it was wandering through the conference space (Fig. 3). The reception

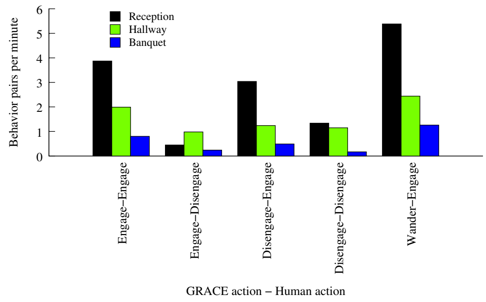

Fig. 3. A 5-second lag sequential analysis of people's responses to GRACE's actions.

had the highest overall rate of participants' engagement with GRACE. Furthermore, when GRACE made an attempt to start an interaction, participants were more likely to respond positively (by turning to the robot, looking at it, or touching its screen) in the reception and banquet. In the hallway, on the other hand, people were equally likely to engage as they were to disengage from the interaction in response to GRACE's interactive behavior. This can be attributed to the transitory nature of the hallway; the other two situations involved a longer time spent socializing in one space. Moreover, people at the reception were more likely to attempt to continue interacting with the robot (through movement, gaze, or via touch-screen) even after GRACE made a disengaging action. While GRACE was reasonably successful at getting help from conference participants (86% of 391 interactions), it is apparent that she was more effective in instigating interaction in certain environments than in others.

For the stationary Roboceptionist, space is a heuristic for identifying interaction partners. The Roboceptionist's design is based on a model of interaction that primarily considers people's location within the hallway to be an indicator of their interactive state. Our observations show that, using this model, the robot overestimates the number of people who engage and interact with it – while it saw 1500 people as engaged and 772 as interacting, our observations evaluated 195 people as interacting and 178 as engaged with the robot. While such overestimation could be advantageous for GRACE's short-term interaction goals, the Roboceptionist's aim of encouraging long-term interaction may be better served by a stricter correspondence of the robot's responses to interaction events. The categorization of people standing in front of the robot as interacting also proved inaccurate. Of 184 instances of participants standing in front of the Roboceptionist, 49% percent are followed by the person directly interacting with the robot (type, turn to) and 19% by gaze in the robot's direction. 32% of the time, people standing in front of the Roboceptionist do not interact. This may be partially due to the similarities between the entryway and GRACE's hallway, since both are generally used by people going to other destinations rather than socializing.

2) Just looking?: Along with using spatial dynamics to identify potential interactions, our observations emphasize gaze as the most common interactive behavior displayed by people toward both robots. Out of 2000 separate instances of GRACE-oriented behaviors coded, gaze interaction accounted for 30% (Fig. 4). Gaze was followed by direct interaction in 82% of cases involving GRACE. With the Roboceptionist, 1923 instances of gaze behavior were identified. The Roboceptionist experienced a significantly lower proportion (27.7%) of people who looked at and then interacted directly with the robot. The transitory nature of the Roboceptionist's location may have been influential, as GRACE also experienced the greatest incidence of non-interactive movement in the hallway (78% of all non-interactive movement).

The higher correlation between gaze and direct interaction with GRACE may be attributable to novelty value; since the Roboceptionist has been occupying its spot since 2003, people may be less intrigued by it. Even though the incidence of direct interaction following gaze behavior seems quite variable, the high frequency of this display of attention to the robot merits consideration. Given more effective vision and the ability to detect gaze, the robots (especially the Roboceptionist) might engage more people in interaction by identifying onlookers and turning to them.

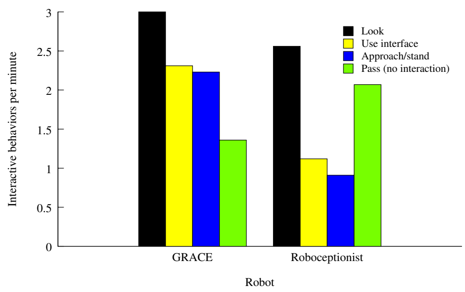

Fig. 4. The types of robot-oriented behaviors exhibited by human interaction partners.

3) Dyadic and group interactions: Dyadic (one-on-one) interaction is implicit in the designs of both robots. GRACE was programmed to wander around until she identified a person or a participant initiated a touch-screen interaction. The robot could only identify, query, and receive directions from one person at a time. Likewise, the Roboceptionist senses many people in the environment but recognizes as interacting only those who are standing directly in front of it and/or typing on the keyboard. For both robots, when one person is identified as interacting, others are ignored until the ongoing interaction is terminated.

Contrary to design assumptions, our observations show that it is equally likely for the robots to be approached by groups of people as by solitary individuals. With GRACE, contrary to expectation, 53% of 171 touch-screen interactions involved more than one person gathering around GRACE and taking turns giving her directions or helping each other understand the task and locate the person in the pink hat. While the banquet and hallway both had very similar distributions of interactions, in the reception hall there were more interactions with GRACE involving two people (37%) than one person (20%). This may be partially accounted for by the social atmosphere and the spatially contained and crowded conditions at the reception, where people came into close contact with each other and with GRACE quite frequently. As for the Roboceptionist, out of 120 interactions, 54% were with one individual, 29% with two people, and 17% with three or more people(Fig. 5).

Although our videos were unaccompanied by audio, during coding we were able to recognize that the groups interacting with the Roboceptionist had different modes of interaction –

Fig. 5. The size of groups with which GRACE and the Roboceptionist interact.

some people were showing the Roboceptionist off to their friends, others were explaining how it functions to a tour group, some were taking turns interacting with the robot, etc. These number- and task-related categories are not perceived by the robot, which is configured to identify only one person to interact with. Input to the robots (via touch-screen for GRACE and via keyboard for the Roboceptionist) can be given by only one person at a time. This creates difficulties in interaction because the dyadic design cannot cope with the qualitatively different modes of interaction that occur in larger groups [25]. The design of the booth and immobility of the keyboard and robot force people to shuffle around when they try to interact in groups – one person steps aside as the other slides in front of the keyboard to type, then they switch again. This is confusing for the robot, and sometimes leads to disruptions in ongoing group interactions.

4) Interpersonal interaction and scaffolding: During interaction with the Roboceptionist and GRACE people often helped each other interact with the robot by discussing its reactions, pointing out the salient aspects of the robot and its environment, and taking turns in interacting and getting responses from the robots. Parents who stopped by the Roboceptionist with their children needed to type in their stead, but placed the children in the space directly in front of the robot. With GRACE, children were picked up and held so they could reach the screen. Adult visitors also interacted with each other while interacting with the robots, pointing out different parts of the Roboceptionist to each other 96 times; with GRACE this happened 43 times. In group interactions, especially when there were two people in front of the Roboceptionist, people took turns typing on the keyboard and standing in front of the robot. With GRACE, participants took turns giving her directions.

Interactions between conference participants differed qualitatively depending on the social and spatial location in which they occurred. We recorded 331 interactions between conference participants while they were interacting with GRACE. Conversation between participants was the most common form (44%), followed by interaction through gaze (33%), spatial movement (walking, standing, turning) (17%), and gesture (touch, waving, pointing) (6%). The distribution over these different types of interactions is similar across the three spaces, but the reception saw a higher frequency of interpersonal interaction (3.3 per minute) than the hallway (1.7 per minute) and the banquet (0.6 per minute). The rate of interpersonal interaction among people gathered around the Roboceptionist was comparatively low (0.53 per minute). The most common type of interaction between people who were interacting with the Roboceptionist was gaze (50%), followed by conversation (42.5%) and movement (standing next to, turning to: 7%).

5) Rhythmicity in interaction: The concept of rhythmicity (rhythmic synchronization and entrainment) is central to human, and therefore human-robot, social interaction [26]. A social robot needs to be able to microcoordinate movements with its interaction partners. This will give it the capability to anticipate interactive behavior rather than merely react to it [27]. This may be especially relevant to the Roboceptionist, since it aims to sustain longer interactions and encourage repeated visits. Also, the Roboceptionist is located in a space where people are mostly passing by and are less likely to interact.

Regarding rhythmicity of interaction, our observations show that the Roboceptionist responds belatedly to people who pass next and close to it. The Roboceptionist turned to a person in response to their proximity on 816 occasions. Only 20% of the Roboceptionist's turns to nearby people happened within one second of their passing, when they could still see the robot if they are walking away at an average pace. Another 29% of the turns were made within 2 seconds following the person's passing, which means that half the turns took 3 seconds or more to initiate. 41% of the Roboceptionist's turns to people passing in its immediate vicinity were followed by an interactive response by the person. The most common reaction was to look in the robot's direction (78%), while movement towards the robot accounts for 14% of interactions. 83% of the responses were made to turns directed at the person in question, while only 17% of people's interactive actions followed turns made in another direction.

The Roboceptionist's model of interaction identifies the person as someone with whom to interact just as they are passing by, after which it turns its screen towards the person. Because the behavior is a reaction to, rather than in anticipation of, the person passing, by the time the robot starts turning the person is often too far away to be engaged by the robot's behavior. Our results show that, rather than reacting to any movement by the Roboceptionist, people respond to motions directed towards them in a timely manner. The Roboceptionist needs to be more rhythmically in sync with the people it aims to interact with, and it must be able to anticipate their actions as well as react to them in order to have smoother interactions.

V. Design recommendations

Roboticists approached the design of GRACE and the Roboceptionist with a "conscious model" [12] of social interaction in mind, which was tailored to accommodate the technical capabilities of the robots. The models can be made more contextual by considering the impact of varying social or spatial environments on the robots. The observation results that showed the influence of the social and spatial environment on GRACE's performance suggest certain ways the robot's design could be improved to create more interaction possibilities for the robot and its interaction partners. They also portray some of the problems with robots that are not "socially embedded" [21] or "structurally coupled" with the environment: their behavior is not very robust and they cannot adapt well to changes in the spatial or social environment. With the results of the observational analysis in mind, we can suggest design modifications to make the robots more adaptive and robust.

Keeping in mind the particular contexts in which the two robots will be interacting, our observations suggest different design modifications for them. GRACE could obtain basic information about the crowd density and motion patterns of people in the room (moving quickly through the hallway, standing around in groups at the reception and banquet) and approach them accordingly. If the room is crowded, the robot should be more assertive in moving through the crowd and asking for space (as in [28]), while fewer people in the room can signal to the robot that it can ask for directions more often and pursue individual people more assertively by standing in front of them.

The Roboceptionist would benefit from a more precise interaction model with which it could identify individuals open to interaction rather than bothering every passerby. The robot will currently turn to a person that types but will not recognize another individual standing next to it, looking and commenting on its performance. In order to cope with group interaction, the robots need to be able to switch attention from one person to the other and follow the rhythm of the interaction, beyond simply responding to typing behavior. Although this is challenging, the robots should be designed such that they can take advantage of affordances in the environment (such as people looking at them) as well as understand patterns of surrounding motion (such as the difference between interactors taking small steps around the booth and walking away from the robot entirely).

These types of suggestions are made possible by using observational analysis to discover properties of the interaction that might not have been imagined or anticipated earlier in the design process. After changes are implemented, comparative studies can be done to evaluate their performance and the results thereof can be included in the next step of the iterative design process.

VI. Conclusion

We suggest that social robots should be observed, objectively and analytically, by trained social scientists in real-world environments with untrained interactors. We have presented a set of results from behavioral analysis in which the spatial environment had a significant effect on the social interactions. We discussed other salient issues in social robot design that emerged from our analysis, such as the use of gaze as a sign of interest in interaction, the differences between dyadic and group interactions, interpersonal interaction and scaffolding interactions with the robot, and the importance of rhythmicity for social interaction. These results challenge original design assumptions and suggest appropriate modifications that might be used to improve the robot's interactive effectiveness.

Acknowledgements

This project was partially supported by National Science Foundation ITR projects #SES-0522630, #IIS-0329014, #IIS-0121426, and a GRFP.

References

[1] C. G. Atkeson, J. Hale, S. Kotosaka, F. Pollick, M. Riley, S. Schaal, T. Shibata, G. Tevatia, A. Ude, S. Vijayakumar, and M. Kawato, "Using humanoid robots to study human behavior," *IEEE Intelligent Systems*, vol. 15, pp. 46–56, 2000.

[2] S. Restivo, "Romancing the robots: Social robots and society," in *Robot as Partner: An Exploration of Social Robots Workshop, International Conference on Intelligent Robots and Systems*, 2002.

[3] C. Breazeal, *Designing Sociable Robots*. MIT Press, 2002.

[4] K. Wada, T. Shibata, T. Saito, , and K. Tanie, "Analysis of factors that bring mental effects to elderly people in robot assisted activity," in *IEEE International Conference on Intelligent Robots and Systems*, 2002.

[5] K. Dautenhahn and I. Werry, "Issues of robot-human interaction dynamics in the rehabilitation of children with autism," in *From Animals to Animats, The Sixth International Conference on the Simulation of Adaptive Behavior*, 2000.

[6] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots: Concepts, design and applications," Carnegie Mellon University Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-02-29, November 2002.

[7] T. Salter, K. Dautenhahn, and R. te Boekhorst, "Robots moving out of the laboratory - detecting interaction levels and human contact in noisy school environments," in *IEEE International Workshop on Robot and Human Interactive Communication*, Kurashiki, Okayama, Japan, 2004.

[8] E. Hutchins, *Cognition in the Wild*. MIT Press, 1995.

[9] B. R. Duffy and G. Joue, "I, robot being," in *Intelligent Autonomous Systems Conference*, Amsterdam, The Netherlands, 2004.

[10] D. Hanson, A. Olney, I. A. Pereira, and M. Zielke, "Upending the uncanny valley," in *National Conference on Artificial Intelligence (AAAI '05)*, Pittsburgh, PA, 2005.

[11] I. Hacking, "The self-vindication of the laboratory sciences," in *Science as Practice and Culture*, A. Pickering, Ed. Chicago: University of Chicago Press, 1992, pp. 29–64.

[12] D. Forsythe, *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford University Press, 2001.

[13] B. R. Duffy, G. Joue, and J. Bourke, "Issues in assessing performance of social robots," in *WSEAS International Conference on Robotics*, Skiathos Island, Greece, 2002.

[14] T. Watanabe, M. Okubo, and H. Ogawa, "A speech driven embodied interaction robots system for human communication support," in *Int. Conf. on Systems, Man, and Cybernetics (SMC2000)*, 2000.

[15] C. Breazeal, C. D. Kidd, A. L. Thomaz, G. Hoffman, and M. Berlin, "Effects of nonverbal communication on efficiency and robustness in human-robot teamwork," in *Proceedings of IROS*, Barcelona, 2005.

[16] T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro, "Interactive robots as social partners and peer tutors for children: A field trial," *Human Computer Interaction (Special issue on human-robot interaction)*, vol. 19, no. 1-2, pp. 61–84, 2004.

[17] J. L. Burke and R. R. Murphy, "Human-robot interaction in usar technical search: Two heads are better than one," in *IEEE International Workshop on Robot and Human Interactive Communication*, Kurashiki, Okayama, Japan, 2004.

[18] K. Dautenhahn and I. Werry, "A quantitative technique for analyzing robot-human interactions," in *Proceedings of the International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, 2002.

[19] L. Suchman, *Plans and Situated Actions: The Problem of Human/Machine Communication*. Cambridge: Cambridge University Press, 1988.

[20] M. J. Mataric, "Situated robotics," in *Encyclopedia of Cognitive Science*. Nature Publishing Group, Macmillan Reference Ltd., November 2002.

[21] K. Dautenhahn, B. Ogden, and T. Quick, "From embodied to socially embedded agents - implications for interaction-aware robots," *Cognitive Systems Research*, vol. 3, pp. 397–428, 2002.

[22] R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. C. Schultz, and J. Wang, "Designing robots for long-term social interaction," in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS '05)*, August 2005.

[23] M. Michalowski, D. Busquets, C. DiSalvo, L. Hiatt, N. Melchior, R. Simmons, and S. Sabanovic, "Social tag: Finding the person with the pink hat (technical report)," in *Proceedings of the National Conference on Artificial Intelligence (AAAI '05) Mobile Robot Competition Workshop*. AAAI Press, July 2005.

[24] "Noldus Information Technology: The Observer," URL: http://www.noldus.com/products/observer.

[25] L. Caporael, "The evolution of truly social cognition: The core configuration model," *Personality and Social Psychology Review*, vol. 1, pp. 276–298, 1997.

[26] R. Collins, *Theoretical Sociology*. New York: Harcourt, Brace, Jovanovich, 1988.

[27] S. Restivo, "Robots 'r' us: The sociology of social robots," Keynote Lecture, Hixon/Riggs workshop on Social Studies of Social Robots, Harvey Mudd College, March 27-30 2003.

[28] S. Thrun, J. Schulte, and C. Rosenberg, "Interaction with mobile robots in public places," *IEEE Intelligent Systems*, pp. 7–11, July/August 2000.