Introduction to this Special Issue* on

Human-Robot Interaction

Sara Kiesler

Carnegie Mellon University

Pamela Hinds

Stanford University

Sara Kiesler is an experimental social psychologist with interests in technology (including robotics), communication, groups, and organizations. She is a professor in the Human Computer Interaction Institute, Carnegie Mellon University, Pennsylvania. **Pamela J. Hinds** studies the impact of technology on individuals and groups. She is an Assistant Professor in the Department of Management Science and Engineering at Stanford University.

*Special Issue of *Human-Computer Interaction*, Volume 19 (2004), Numbers 1 & 2.

Human-Robot Interaction

Human-computer interaction (HCI), as a field, has made great strides toward understanding and improving our interactions with computer-based technologies. From the early explorations of direct interaction with computers, we have reached the point where usability, usefulness, and an appreciation of technology's social impact, including its risks, are widely accepted goals in computing. HCI researchers, designers, and usability engineers work in a variety of settings on many kinds of technologies. Recent proceedings of the CHI conference give evidence of this diversity. Topics include not only the office systems where HCI work began, but also tiny mobile devices, web and Internet services, games, and large networked systems. This special issue introduces a rapidly emerging technology and new focus for HCI – autonomous robots and the human-robot interactions required by these robots.

Until recently, HCI researchers have done little work with robots. Keywords related to robots or to human-robot interaction have not been included in the standard list of CHI topics, and this year was the first in which robots were a theme. This state of affairs was reasonable. As Sebastian Thrun's opening essay in this special issue explains, today's workhorse robots are mainly programmable industrial machines that offer modest challenges in human-computer interaction. Now, advances in computer technology, artificial intelligence, speech simulation and understanding, and remote controls have led to breakthroughs in robotic technology that offer significant implications for the human-computer interaction community.

Autonomous mobile robots can identify and track a user's position, respond to spoken questions, display text or spatial information, and travel on command while avoiding obstacles. These robots will soon assist in a range of tasks that are unpleasant, unsafe, taxing, confusing, low paid, or boring to people. For example, nurses making rounds in assisted living facilities spend much of their time sorting and administering medications. A robotic assistant could do some of this work, as well as chores that are difficult for elderly people such as fetching newspapers and mail, getting up and down stairs, getting things out of high or low cabinets, and carrying laundry, enabling elderly people to be independent longer. Robotic assistants in the future might act as guards, help fight fires, deliver materials on construction sites and in mines, and distribute goods or help consumers in retail stores. Robots might even provide high-interaction services such as taking blood and coloring hair.

Autonomous robots like these will need to carry out social and intellectual as well as physical tasks. Ideally, these robots will create a comfortable experience for people, gain their cooperation, encourage healthy rather than overly dependent behavior in clients, customers, and co-workers, and provide appropriate feedback to remote operators and others involved in the robotic system. Although roboticists are gaining practical experience with mobile, autonomous robots in settings such as museums (Thrun et al., 2000), we lack a principled understanding of how to design robots that will accomplish these more ambitious goals.

Human-Robot Interaction in HCI

HCI, and its sister discipline, human factors, offer a rich resource for research and design in human-robot interaction. Much has been learned in the last three decades about how people perceive and think about computer-based technologies, about human constraints on interaction with machines, about the factors that improve usability, and about the primary and secondary effects of technology on people and organizations. Much of this work will be applicable to robots. Nonetheless, autonomous robots are a distinctive case for several reasons:

First, people seem to perceive autonomous robots differently than they do most other computer technologies. People's mental models of autonomous robots are often more anthropomorphic than are their models of other systems (Friedman, Kahn, & Hagman, 2003). The tendency for people to anthropomorphize may be fed, in part, by science fiction and, in part, by the powerful impact of autonomous movement on perception (Scholl & Tremoulet, 2000). When we build autonomous robots to look human, we may encourage anthropomorphic mental models of these systems. As Hinds, Roberts, and Jones (this issue) explain, some roboticists argue that humanoid robots provide for a more natural interface than more mechanistic robots. Thus, humanoid robotics are the focus of much research and development.

A second major reason autonomous robots are a distinctive case in HCI is that robots are ever more likely to be fully mobile, bringing them into physical proximity with other robots, people, and objects. As two articles in this issue make clear, mobile robots will have to negotiate their interactions in a dynamic, sometimes physically challenging, environment (e.g. Burke, Murphy, Coovert, & Riddle, this issue; Yanco, Drury, & Scholtz, this issue). If one or more remote operators partly control the robot, they must help the robot negotiate its interactions in the remote space, creating a complex feedback system. In figure 1, we show one such futuristic scenario in a medical setting. Here, we see a robot whose task is to sort and dispense medications interacting with an elderly client. At the same time, the robot is designed to sense its clients' state, using indicators such as unusual posture, gestures, or eye movement indicating illness. A remote medical worker monitors this information, adjusting the robot's route or tasks as needed and watching for signs of problems in client states. In this example, the interfaces of interest involve the robot-client, robot-operator, and even multiple person or robot interactions.

A third reason that autonomous robots are a distinctive case for HCI is because these robots make decisions, that is, they learn about themselves and their world, and they exert at least some control over the information they process and actions they emit. Of course, many computer agents in desktop, automotive, and other computer systems make decisions, and the use of agents is increasing rapidly. Computer agents present interesting HCI issues, for example, to what extent the agent should display confidence intervals for the decisions it makes. An autonomous robotic system will add even more complexity because it must adjust its decisions sensibly and safely to the robot's abilities and to the options available to the robot in a given environment. The system also must detect and respond to changes in the environment and its users. Imagine a robotic walker that guides a frail person and detects when its user is ill or falling or when the environment is dangerous. How much control should such a walker take? How sure of itself should it be? How should it respond if the user wants to turn back, stop, or opposes the robot's suggestions? As these questions suggest, designing an appropriate interaction scheme and interface for such a system requires an understanding of the people who will use such a system, and of their world. As the ethnography of elders in this volume shows (Forlizzi, DiSalvo, & Gemperle, this issue), designing these robots appropriately will require a deep understanding of the context of use and of the ethical and social considerations surrounding this context.

We do not claim that these problems are entirely new. Design explorations and research in human-robot interaction existed in the field of robotics since at least the mid 1990s. At Interval Research Corporation, for example, Mark Scheef and his colleagues built Sparky, a "social robot," and studied children's and adults' reactions to it (Scheef, Pinto, Rahardja, Snibbe, & Tow, 2000). Today, many such developments are taking place in Europe and in Japan, for instance, the humanoid Robovie robot described in this issue (Kanda, Hirano & Eaton, this issue). MIT's Robotic Life project is an example of design explorations at the edge of robotics and HCI, in which Cynthia Breazeal and her colleagues are trying to create capable robotic creatures with a lifelike presence. Another example in quite a different domain is the work of Brian Scassellati (2000), who builds human-like robots to investigate models of human development. Other domains include space exploration and the military. Over the last few years, research on human-robot interaction has gained increasing attention and funding. The National Science Foundation and the Defense Department's DARPA recently co-sponsored an interdisciplinary workshop in which participants discussed problems of human-robot interaction for Robonaut, a robot that will help astronauts outside a space capsule, and for search and rescue robots (Murphy & Rogers, 2001). Two yearly conferences now provide a forum for papers on human robot interfaces – the *Humanoid Robots Conference* and the *IEEE RO-MAN Workshop*.

In planning this special issue, we noted that despite the many prior and ongoing activities in robotics related to human-robot interaction, most of the development and the published literature in this area is concerned with technical advances in robotics and computer science that make human-robot interaction possible. Our goal for this issue was to stretch the field of inquiry by focusing especially on behavioral, cognitive, and social aspects of human-robot interaction and the social contexts surrounding human-robot interaction. For example, we hope this special issue will encourage researchers in the field to think about what useful tasks robots can and should do in real social environments, and how to improve how robots communicate and respond to ongoing human communications and behaviors. We invited work that examined human-robot interaction in its social context. We imposed another bias too: Given the comparative absence of systematic empirical investigation in the field, we gave preference to systematic empirical studies and to interdisciplinary collaborations. We also encouraged authors to reflect on the social and ethical issues raised by the deployment of robots in work or everyday life. The HCI community is an especially appropriate place to carry out this kind of analysis because of its legacy of applying behavioral and social science to technical problems and of doing interdisciplinary research and design.

Contents of this Special Issue

In this issue of *Human-Computer Interaction*, we present six articles in the emerging area of human-robot interaction.

The first article, an invited essay by Sebastian **Thrun**, provides a technical context for the articles that follow. The author reviews the state of the art in robotics, suggests advances that are likely in the future, and points out some challenges faced in robotics that impinge on human-robot interaction. The author has suggested a useful framework for HCI researchers' work in human-robot interaction, that is, a framework that differentiates among three kinds of robots – industrial robots, professional service robots that will operate in work organizations and public settings, and personal service robots. These three kinds of robots have different capabilities, different user groups, and different contexts of use. This framework will help the HCI community identify some of the greatest opportunities for research in human-robot interaction

The first empirical article, by **Forlizzi, DiSalvo, and Gemperle**, offers a theoretical ecological framework for the design of personal service robots in homes of elderly people. The authors use this framework to show how aging occurs within a local ecology that includes the elder person, the home, products within the home, and important people in the elderly person's life. The authors describe a fascinating ethnography of elders in which they explore how products maintain or lose their usefulness and value for well and ill elders. More generally, the study and the framework should help designers and researchers to consider, and design for, the social context of personal service robots.

The next article, by **Kanda, Hirano, and Eaton**, presents a field study of two robots that visited a children's elementary school in Japan for two weeks, with the purpose of teaching children English. This article is a good example of a field trial with robots. The trial exemplifies the risks and advantages of studying peoples' responses to robots over time in a real social setting. The authors had to understand and cope with problems of a noisy environment and rambunctious children, but they were able to track interactions and the effects of these interactions on learning over time. The children's enthusiasm for the robots waned over the two-week period, but those children who continued to interact with the robot (mainly those who could understand a bit of the robot's English to begin with), learned from it. Although the effects are modest and the time was short, the results of this study are impressive because this study is the first practical demonstration that students can learn from a humanoid robot.

The third empirical article, by **Burke, Murphy, Coovert, and Riddle**, reports on an opportunistic field study of search and rescue robots used as part of a night rescue training exercise. The authors made careful observations of how remote operators interacted with the robots and one another, and then developed a systematic coding scheme to analyze these interactions. To their surprise, the main human-robot interaction problem was not remote navigation but rather understanding the situation the robot had encountered. The authors describe the interactions among team members who helped the operator understand the state of the robot and the environment. This article is not only an interesting account of the people and robots used in disaster search and rescue, but also points to some of the main human-robot interaction problems in this domain.

The fourth empirical article by **Yanco, Drury, and Scholtz**, offers a different perspective on human-computer interaction for search and rescue robots. The authors took advantage of a yearly robotics IEEE competition for search and rescue robots held to encourage advancements in this field. They developed metrics to compare the usability of the human-robot interface across competitors, and compared their observations using these metrics with performance scores in the competition. The authors argue that usability standards for other kinds of computer interfaces are only partly applicable to human-robot interfaces. For example, as did the authors of the previous article, these authors concluded that one of the biggest problems in the human robot interface for search and rescue robots is that the remote operator often lacks accurate situation awareness of the robot's state and the state of the environment in which the robot is located. This problem seems to us to be unique to human-robot interaction, and especially difficult because of simultaneous changes taking place in the operator, the robot, and the robot's environment.

The fifth empirical article, by **Hinds, Roberts, and Jones**, is an experimental laboratory study. The authors explore how people who have to work closely with professional service robots will perceive and work with these robots. This study is one of the first controlled experiments to examine the effect of a humanoid robot appearance on peoples' responses, with a machine-like robot used as a comparison. The study suggests that people may be more willing to share responsibility with humanoid as compared with more machine-like robots, a possibility that has important implications for collaborations in which the robot makes key decisions about the task.

Taken as a whole, these articles represent some of the first systematic empirical research in human-robot interaction. We hope these articles show that human-robot interaction offers many interesting and important problems for the HCI community. More interdisciplinary collaboration between roboticists, behavioral and social scientists, and designers is important, we believe, to advancing the field of human-robot interaction. Roboticists understand the technology and its applications; behavioral scientists and others can provide theory and methods. But there are plenty of opportunities even for those far from a robotics laboratory. For instance, research on computer agents, avatars, and other ways of representing autonomous, computer-based assistance, will contribute to our understanding and design of robots. Useful studies also can proceed using commercial robots such as AIBO and the HelpMate robot (King and Weiman, 1990), conducting laboratory studies using robot shells and Wizard of Oz methods, or performing ethnographic studies of the contexts to which robots may be applied. In general, we see many opportunities for researchers of all stripes and believe that leadership from the HCI community could advance research in human-robot interaction in important ways, influencing the development of the field and the design of robots.

NOTES

Support. We acknowledge the support of the National Science Foundation (IIS-0121426) in preparing this special issue.

Authors' Addresses. Sara Kiesler, Human Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA, 15213, Email: kiesler@cs.cmu.edu. Pamela J. Hinds, Management Science & Engineering, Terman 424, Stanford University, Stanford, CA 94305-4026, Email: phinds@stanford.edu.

ARTICLES IN THIS SPECIAL ISSUE

Burke, J. L., Murphy, R. R., Coovert, M. D. and Riddle, D. L. (2004). Moonlight in Miami: A Field Study of Human-Robot Interaction in the Context of an Urban Search and Rescue Disaster Response Training Exercise. *Human-Computer Interaction*, *19*, xxx-xxx.

Forlizzi, J., DiSalvo, C. and Gemperle, F. (2004). Assistive Robotics and an Ecology of Elders Living Independently in Their Homes. *Human-Computer Interaction*, *19*, xxx-xxx.

Hinds, P. J., Roberts, T. L. and Jones, H. (2004). Whose Job is it Anyway? A Study of Human-Robot Interaction in a Collaborative Task. *Human-Computer Interaction*, 19, xxx-xxx.

Kanda, T., Hirano, T. and Eaton, D. (2004). Interactive Robots as Social Partners and Peer Tutors for Children: A Field Trial. *Human-Computer Interaction*, *19*, xxx-xxx.

Thrun, S. (2004). Toward a Framework for Human-Robot Interaction. *Human-Computer Interaction*, *19*, xxx-xxx.

Yanco, H. A., Drury, J. L. and Scholtz, J. (2004). Beyond Usability Evaluation: Analysis of Human-Robot Interaction at a Major Robotics Competition. *Human-Computer Interaction*, *19*, xxx-xxx.

REFERENCES

Friedman, B., Kahn, P.H., & Hagman, J. (2003). Hardware companions? – What online AIBO discussion forums reveal about the human-robotic relationship. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems*, 273-280. New York: ACM.

King, S. and Weiman, C. (1990). Helpmate autonomous mobile robot navigation system. *Proceedings of the SPIE Conference on Mobile Robots 235,,*190–198.

Murphy, R., & Rogers, E. (2001). *Human-robot interaction: Final Report for DARPA/NSF study human-robot interaction.* Retrieved June 5, 2002, from http://www.csc.calpoly.edu/~erogers/HRI/HRI-report-final.html

Scassellati, B. (2000). How robotics and developmental psychology complement each other. *Proceedings of the NSF/DARPA Workshop on Development and Learning*. Lansing, MI: Michigan State University.

Scheef, M., Pinto, J., Rahardja, K., Snibbe, S., Tow, R. (2000). Experiences with Sparky, a social robot. *Proceedings of the 2000 Workshop on Interactive Robot Entertainment*.

Scholl, B. J. & Tremoulet, P. (2000). Perceptual causality and animacy. *Trends in Cognitive Science, 4,* 299-309.

Thrun, S., Beetz, M., Bennewitz, M., Burgard,W., Cremers, A., Dellaert, F., Fox, D., Hahnel, D., Rosenberg, C., Roy, N., Schulte, J., and Schulz, D. (2000). Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *International Journal of Robotics Research*, 19, 972–999.