Whose Job Is It Anyway?
A Study of Human–Robot Interaction in a Collaborative Task

Pamela J. Hinds
Stanford University

Teresa L. Roberts
PeopleSoft, Inc.

Hank Jones
Stanford University

Pamela 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. **Teresa Roberts** is a professional in human–computer interaction, with interests in user-centered design and computer-mediated communication; she has most recently been a senior interaction designer at PeopleSoft and at Sun Microsystems. **Hank Jones** is a robotics researcher with an interest in user-centered design of human–robot interactions; he recently completed his doctorate in Aeronautics and Astronautics from Stanford University in the Aerospace Robotics Laboratory.

CONTENTS

ABSTRACT

The use of autonomous, mobile *professional service robots* in diverse workplaces is expected to grow substantially over the next decade. These robots often will work side by side with people, collaborating with employees on tasks. Some roboticists have argued that, in these cases, people will collaborate more naturally and easily with humanoid robots as compared with machine-like robots. It is also speculated that people will rely on and share responsibility more readily with robots that are in a position of authority. This study sought to clarify the effects of robot appearance and relative status on human–robot collaboration by investigating the extent to which people relied on and ceded responsibility to a robot coworker.

In this study, a 3×3 experiment was conducted with human likeness (human, human-like robot, and machine-like robot) and status (subordinate, peer, and supervisor) as dimensions. As far as we know, this study is one of the first experiments examining how people respond to robotic coworkers. As such, this study attempts to design a robust and transferable sorting and assembly task that capitalizes on the types of tasks robots are expected to do and is embedded in a realistic scenario in which the participant and confederate are interdependent. The results show that participants retained more responsibility for the successful com-

pletion of the task when working with a machine-like as compared with a human-oid robot, especially when the machine-like robot was subordinate. These findings suggest that humanoid robots may be appropriate for settings in which people have to delegate responsibility to these robots or when the task is too demanding for people to do, and when complacency is not a major concern. Machine-like robots, however, may be more appropriate when robots are expected to be unreliable, are less well-equipped for the task than people are, or in other situations in which personal responsibility should be emphasized.

1. INTRODUCTION

Advances in artificial intelligence and speech recognition, less expensive yet more sophisticated mobile computing hardware, and even such mundane changes as increasing ubiquity of ramps in public buildings have combined to make *professional service robots*—robots that assist workers—more practical than ever before. Autonomous mobile robots made with current technology can identify and track people and objects, understand and respond to spoken questions, and travel to a destination while avoiding obstacles (see Fong, Nourbakhsh, & Dautenhahn, 2002). Robots can be built to have abilities that complement human abilities. They can go to places that are toxic or unsafe and can tolerate repetitive, mundane tasks. They can have large databases of knowledge and can connect through networks to vast sources of additional data.

With these ongoing advances, the use of robots in the workplace is likely to grow substantially. The workplace in the near future will increasingly contain robots and people working together, each using their own stronger skills, and each relying on the other for parts of the tasks where the other has the better skills. In a recent report (United Nations, 2002; see also Thrun, 2004), the United Nations indicated that the use of these professional service robots will grow substantially in the next few years in fields as diverse as the military, medical services, and agriculture. Autonomous robots, for example, are expected to work in tandem with military personnel so that soldiers can better understand the dangers of the battlefield; robots also will supply troops with ammunition and provide surveillance (Squeo, 2001) and assist astronauts in investigating distant planets (Ambrose, Askew, Bluethmann, & Diftler, 2001). Already, robots perform the mundane chore of delivering medications from pharmacies to nursing stations in hospitals, using their intelligence to avoid obstacles as they travel (Okie, 2002; Siino & Hinds, 2004); people, however, are required for loading and unloading the medications, and for programming the robot's destination. Sheridan (1992) described an "optimistic scenario" in which robots will

grow in number and variety, becoming available to us to do our beck and call in our homes, schools, and government facilities, in our vehicles, our hospitals, and across the entire spectrum of our workplaces—factories, farms, offices, construction sites, mines, and so on. (p. 336)

In many instances, these robots will share the same physical space with people and work closely with people to accomplish joint tasks as part of their day-to-day work.

Professional service robots, this newer class of robots, are specifically designed to assist workers in accomplishing their goals (see Thrun, 2004). These robots differ from industrial robots and many other technologies found in the work environment (e.g., appliances, computers, navigation systems, etc.) because they are mobile, they do things without being commanded, and they are interactive. These differences suggest that professional service robots may affect the work environment in socially important ways. Because of their ability to move with apparent intentionality in physical space, they are likely to be perceived as animate, triggering social responses (for a review, see Scholl & Tremoulet, 2000). Their ability to travel between different departments also may allow the unplanned movement of information between distant coworkers.

If professional service robots are to share the workplace with people, we need to understand what the interaction between them is likely to be like. Will people trust robots to perform operations that the robots are capable of, without oversight? If things go wrong, will people take appropriate responsibility to correct the problem, or will they abdicate responsibility to the robot? In the face of uncertainty, will people ask for and accept the guidance of expert robots? What aspects of the design of the robot will affect the way people and robots work together? The better we understand these questions, the better we can design robots to be effective work partners.

For the study we report here, we conducted a laboratory experiment designed to look at the effects of the robot's appearance and the relative status of the robot on how people work with robots. We also compare human–robot interaction with human–human interaction to better understand how interacting with robotic partners may alter the current work environment. We studied the effects of robot appearance because roboticists are currently making at least *de facto* decisions about appearance without the benefit of information on the ramifications and perhaps with misconceptions of their effects. We chose relative status as a second dimension because it also can be relatively easily manipulated when introducing a robot into a team, and because status can have a powerful effect on relationships between coworkers. Although these are only two of many possible considerations for the design and implementa-

tion of robots (e.g., speech mechanisms, autonomy levels, sensor types, etc., are other important design considerations), we believe appearance and status are particularly important to the design of professional service robots because they are, for most robots, "free variables" for the robot designer that are minimally dependent on technological advances.

2. THEORY AND HYPOTHESES

2.1. Collaboration With Human-Like Versus Machine-Like Robots

Current work in the field of robotics is flooded by efforts to make robots more human-like. Roboticists are designing robots with heads, faces, necks, eyes, ears, and human-like voices based on the premise that a humanoid robot is the most appropriate form for human–robot interaction (Ambrose et al., 2001; Brooks & O'Reilly, 2002; Hashimoto, Narita, Sugano, Takanishi, Shirai, & Kobayashi, 2002; Ishiguro, 2003; Simmons & Nourbakhsh, 2002). Researchers argue that a humanoid form will ease interaction because rules for human social interaction will be invoked, and thus, humanoid robots will provide a more intuitive interface (Breazeal & Scassellati, 1999; Brooks, 2002). Brooks, for example, suggested that, "it will be useful for a large mobile appliance and a person to be able to negotiate who goes first in a tight corridor with the same natural head, eye, and hand gestures all people understand already" (p. 38). The premise that the humanoid form is the appropriate form for human–robot interaction, however, remains largely untested. Opponents of a humanoid form suggest that robots are machines and that humanoid features may generate unrealistic expectations and even fear (see Dautenhahn, 1999). Turkle (1984) observed that it is important to people that we be able to see ourselves as different from machines, asserting that the blurring of the line between people and machines can be disturbing and frightening. Brooks also suggested that the current infatuation with humanoid robots may be a phase through which we need to pass as we learn more about human–robot interaction. These considerations about humanoid robots, both positive and negative, may affect people's response to professional service robots in the workplace, particularly with regard to their willingness to rely on robots to help them achieve their goals. Therefore, it is important to conduct empirical studies of human-like versus machine-like robots and to explore the trade-offs.

In this study, we examine how the appearance (humanoid vs. machine like) of a robot might affect people's willingness to rely on and share responsibility with their robotic partner. We choose to focus on these dependent variables because they are central to the collaboration process.

The first response we examine is reliance on work partners. People rely on others (both other people and machines) when those others have capabilities that they, themselves, do not have (e.g., trusting sums computed on a calculator). What is more variable, however, is the extent to which people rely on others when their relative abilities are less obvious a priori. The extent to which people rely on a new technology instead of their own or other people's input is crucial to the success of technology and to the benefits of the technology being realized (e.g., Wall, Jackson, & Davids, 1992). Although overreliance on technology can have disastrous effects (e.g., a 1994 midair collision resulted when one pilot neglected to take manual control from the automated system; Sparaco, 1994), we focus on underreliance. Our reasoning is that getting people to rely on robots is the more pressing concern. There is substantial evidence that people resist technologies that are programmed to augment human decision making even when the technology is more accurate. Gawande (2002), for example, reported that patients preferred the judgment of a cardiologist for interpreting electrocardiogram reports even when an automated system provided correct responses 20% more often than the cardiologist. At this point, however, little is known about when and why people will rely on robots as compared with people, particularly advanced robots that have the ability to engage in collaborative tasks and the discretion to make decisions.

The second response we examine in this research is the extent to which people assume responsibility for the task. Roberts, Stout, and Halpern (1994; also Grabowski & Roberts, 1997) extensively discussed the importance of accountability and responsibility for organizational tasks. They argued that accountability may improve the quality of decisions because decision makers who feel responsible consider more alternatives (see also Tetlock, 1985). They also, however, pointed out that too much responsibility can be unpleasant and can lead to rigidity (Roberts et al., 1994).

Responsibility and reliance could be inversely related. That is, as people rely more on a robot, they may assume less responsibility for the task, and might care less about the resulting success and failure of the work. However, just as a person might rely on a spell checker to provide correct spelling (e.g., "there" and "their," but not "thier"), and might even allow a grammar checker to suggest the right homonym, the person retains ultimate responsibility for picking the correct spelling in context. Therefore, we argue that reliance on a robot does not necessarily breed complacency or abdication of responsibility, and that these two constructs and their antecedents must be examined separately. The most appropriate mix of reliance and responsibility in human–robot collaboration, for example, may be one in which the human relies on the robot for maximum input but does not abdicate responsibility.

We anticipate that people will rely more on and cede more responsibility to human-like as compared with machine-like robots. Over the last decade, re-

search has suggested that people may respond to computers and other technology using the same social rules that they use with people (see Nass, Steuer, Tauber, & Reeder, 1993; Reeves & Nass, 1996). People, for instance, are polite to computers, use norms of reciprocity, and apply gender stereotyping (Reeves & Nass, 1996). People respond to technology using social rules in part because the primary model people have for dealing with an intelligent, autonomous "other" is human-to-human social interaction. Although they do not necessarily believe that computers and other technologies are human, they are drawn to interact using social rules because cues such as natural language usage and interactivity trigger these responses. Further, these researchers argued that the extent to which social rules are applied depends, in part, on the number and strength of cues conveyed by the technology. Steuer (1995), for example, claimed that there are five characteristics that cue people to interact as though their partner is a social actor: natural language use, interactivity, human social roles, human-sounding speech, and human-like physical characteristics. This line of thinking suggests that people also may use human social rules when interacting with autonomous robots. As more of the factors Steuer listed are exhibited, people may respond to robots in ways that more closely mirror human–human interaction. Therefore, more human-like robots as compared with machine-like robots should elicit higher levels of reliance.

Another reason we expect more reliance on human-like as compared with machine-like robots is because human-like robots may be perceived as more predictable or responsive than machines, and thus, people may be more comfortable interacting with them. When assigned collaborative tasks with collocated colleagues, it is generally considered appropriate for people to share ideas, interact with one another, and engage in collaborative decision making (see Kraut, Fussell, Lerch, & Espinosa, 2002; Olson & Olson, 2000). Such collaboration, however, requires an understanding of how the person interacts and makes decisions, and of the person's knowledge and capabilities. Extensive research on how people reach common ground with others has established that people estimate the knowledge of others based on cues that they receive from that person and from the environment, and that they subsequently tailor their own communications according to the common knowledge they believe is shared (see Fussell & Krauss, 1992; Issacs & Clark, 1987).

People also will make estimations of the capabilities of robots as they develop a mental model of what the robot knows. Human-like characteristics are likely to engender a more human mental model of the robot (see Kiesler & Goetz, 2002). That is, the conceptual framework that people use to predict and interpret the robot's behavior may be more similar to that used to predict and interpret the behavior of people. With a more human mental model, people are more likely to assume human-like traits and capabilities. Therefore, people may assume that more common ground is

shared with the human-like as compared with the machine-like robot, thus reducing uncertainty and facilitating collaboration. With human-like as compared with machine-like robots, people may also feel a stronger sense of shared identity. Parise, Kiesler, Sproull, and Waters (1996), for example, found that participants cooperated more with human-like agents and less with dog-like agents, although they found the dog-like agents more likeable. Parise and his colleagues argued that this difference occurred because people felt more similar to agents that were more human like, thus increasing their sense of shared social identity.

The aforementioned lines of reasoning suggest that people will be more at ease collaborating with human-like robots. Perceived common ground and shared identity with a human-like robotic partner will facilitate collaboration because the person is likely to be more confident in his or her estimates of the robot's knowledge and in his or her ability to interact effectively with it. Therefore, we predict that when people are collaborating with robots on ambiguous tasks, they will rely more on human-like as compared with machine-like robots. Using the same logic, we anticipate that people will relinquish more of their sense of personal responsibility for the task to human-like as compared with machine-like robots.

Hypothesis 1a: People will rely on a human-like robot partner
                    more than on a machine-like robot partner.
Hypothesis 1b: People will feel less responsible for the task
                    when collaborating with a human-like robot
                    partner than with a machine-like robot partner.

2.2. Relative Status of Robot Coworkers

The status hierarchy has historically been one of the more pronounced features of social and organizational life. Our perceptions of others' status can determine our perceptions of the target's capabilities (see Swim & Sanna, 1996) and performance (Pfeffer, Cialdini, Hanna, & Knopoff, 1998), our willingness to defer to the target's opinion (e.g., Strodtbeck, James, & Hawkins, 1957), and our willingness to assume responsibility versus allowing another to assume it (Roberts et al., 1994).

Outside of science fiction, technology typically plays a subservient, lower status role relative to those who use it. Technology products, including robots, typically are perceived as servants or tools designed to help us to achieve our goals. As robots gain more autonomy, however, there may be cases in which the robots need increased authority to encourage people to defer to the robots' expertise (see Nass, Fogg, & Moon, 1996). For example, in complex environments, people may not have complete information or the capacity to process

information as rapidly as robots. In such cases, deferring to the robot may improve the likelihood of task success. Consistent with this idea, Goetz, Kiesler, and Powers (2003) recently reported that people complied more with a serious, more authoritative robot than with a playful robot when the task itself was serious. It seems that how the robot is presented to those collaborating with it may affect the extent to which people are willing to rely on it.

Research on status effects clearly demonstrates that even arbitrarily assigned status labels (e.g., leader, supervisor, expert, etc.) cause people to attribute more competence to those of higher status. Surprisingly, this effect holds even outside of the target's domain of expertise. For example, research many years ago on jury decisions indicated that people rely more on the opinions of those who hold more prestigious, although unrelated, professional positions (e.g., Strodtbeck et al., 1957). More recent research shows that when people are labeled as leaders, even when the label is clearly arbitrary, observers are more likely to see the targets as evincing leader-like behaviors (Sande, Ellard, & Ross, 1986).

Research also examines how workers' sense of responsibility shifts when they are in leadership positions. Supervisors and leaders typically see themselves as more competent and more responsible for the assigned task (see Sande et al., 1986). Also, when supervising or when a supervisor is involved in a task, people view the work product as better (Pfeffer et al. 1998). Often times, organizations hold supervisors responsible for the actions of their subordinates.

Assuming that people respond to robots' roles using social rules similar to those used with people, we hypothesize that people will rely more on the robot partners and assume less responsibility for the task when working with robots that are supervisors as compared with robots that are peers and subordinates.

Hypothesis 2a: People will rely on the robot partner more when it is characterized as a supervisor than when it is characterized as a subordinate or peer.

Hypothesis 2b: People will feel less responsible for the task when collaborating with a robot partner who is a supervisor than with a robot partner who is a subordinate or peer.

2.3. Interaction Between Human Likeness and Status

Although we expect a main effect for status, we also anticipate an interaction between status and human likeness. Given our earlier hypotheses, we expect that people's sense of responsibility for a task will be highest when the

partner is more machine-like *and* in a subordinate role. We reason that people will view the machine-like robotic partner as a tool intended to help them with their task. Therefore, they should treat the robot as they would a pen, hammer, or shovel—tools that have clearly defined, mechanical abilities but no will of their own and can assume no responsibility. Therefore, we posit that people will feel most responsible for the task when they are working with machine-like subordinate robots. We present no interaction hypotheses predicting reliance, however, because we reason that people are accustomed to relying on tools to help them in accomplishing their work.

> Hypothesis 3: People will feel the greatest amount of responsibility when collaborating with machine-like robot subordinates as compared with machine-like robot peers and supervisors; and as compared with human-like robot subordinates, peers, and supervisors.

3. METHOD

To test our hypotheses, we conducted a 3×3 laboratory experiment. The experiment was a between-subject design, manipulating human likeness (human, human-like robot, machine-like robot) and status (subordinate, peer, supervisor) with the human condition as the baseline. Each participant was asked to collaborate on a task with a confederate who reflected one of the nine cells in the design. The confederate used the same script for all conditions and was unaware of the status manipulation. In the robot conditions, we used a *Wizard of Oz* approach in which the robot was teleoperated, appearing to be operating autonomously. The same man teleoperated and spoke for the robot in the two robot conditions, and he acted as the human confederate. The experiment was videotaped with cameras suspended from the ceiling of the experimental lab.

3.1. Participants

Participants were 292 students recruited on a university campus, randomly assigned to condition, and paid for their participation. The mean age of participants was 20.51 years old. Fifty-nine percent of the participants were women. Because we thought it was important that participants believed that our robots operated autonomously, the last question we asked those in the two robot conditions was how they thought the robot worked, from a technical standpoint. Forty-two (21.5%) of the participants who worked with one of the robots expressed suspicion about whether or not the robot was autonomous. Suspicious

participants were approximately equally spread between the human-like and machine-like robot conditions. When these cases were removed from the analysis, there was no effect on the pattern of results, so the analyses we present here include data from all participants.

3.2. Tasks and Procedures

In the experiment, we asked participants to work with a partner in a parts depot for a company that develops innovative remote-control devices. They were told that their job was to collect the parts required to assemble various objects that would be assembled by another team of workers. The task entailed working with the confederate to jointly collect objects from a list, put them into bins, and take the bins to a table near the door. Participants were told that the confederate was familiar with the location of the parts and could carry the bins on its tray, but did not know what was needed and would not be able to collect parts. The division of labor helped establish interdependence between the participant and the confederate, and created a plausible story for why the confederate was not able to open drawers and pick up items. The task was also designed to capitalize on the unique capabilities of a robot (e.g., carrying materials, moving around a room, and remembering detailed information about the location of objects), although still making sense for a human confederate. Finally, the task was one that could be credibly conducted in ways consistent with each of the possible status conditions without any modification of the script.

On arriving at the lab, participants were given a packet of instructions. They completed a brief demographic survey and then were provided detailed instructions on the task. After reading the instructions, participants were given four pages, each containing a list of items that were to be collected during the task; the items on each page were to be collected into a single bin, one bin per page. Then, they were introduced to the confederate (the human, the human-like robot, or the machine-like robot) with the experimenter saying, "I'd like to introduce Chip, who will work with you on this study." In all cases, the confederate entered the room ready to begin the task and, after acknowledging the participant's greeting, started the task by asking, "What's first on the list?" After that, the pace of the task was determined by the participant reading the parts lists. The participant read out the items from the list and the confederate identified which cabinet and which drawer the parts were in. The participant collected the parts from the drawers and put them into a bin on the tray that the confederate was holding. The confederate was prepared to answer questions about ambiguous items, if asked. The task took about 20 min to complete.

Some ambiguity was built into the task to increase uncertainty and cause the participant to make explicit decisions about whether or not to rely on the confederate for more than just the rote aspects of the task. The opportunity for errors also provided a basis on which the participant could assign responsibility and blame. For example, in one case, there were not enough parts of the specified color. In another case, the confederate allegedly misunderstood the participant's instruction and directed the participant to the wrong drawer. Figure 1 provides a sample transcript for one session in which the participant was interacting with a machine-like robot. In this scenario, there were not enough "four-slot connectors" of the required color, so the participant had to figure out how to handle this anomaly.

When all four sets of parts had been collected, the experimenter returned to the room and removed the four bins. The confederate also left at this time. After waiting a short time, the experimenter returned with a handwritten score-sheet showing how well the dyad had performed. The scores were always the same regardless of the participant's actual performance. The scores for the four bins ranged from 72% to 100%, so that the participant would perceive both failure and success.

After receiving their scores, the participants filled out a survey with questions about their experience on the task.

3.3. Manipulations

Human Likeness

Human likeness had three levels in this study. The human baseline condition was created by having a human confederate play the role of the partner. For the human-like and machine-like robot conditions, we used a single robot that could be teleoperated and whose appearance could be easily altered. The robot, which is available commercially and is frequently used for trade shows, stood on a circular base that was 23.50 cm high.[1] The base contained wheels that allowed the robot to move. Radio controls allowed the operator to make it roll forward and backward, and to turn. The link between the operator and the robot was completely wireless. A camera mounted in the robot's head allowed the operator to see the experimental lab from the robot's view. The operator interacted with the participant (in the next room) through microphones and speakers on the robot and on the

---

1. The robot was purchased from The Robot Factory (**http://www.robotfactory.com/**).

Figure 1. **Excerpt from a session including a participant (P) and a machine-like robot (R).**

P: [Reading from packet.] Twelve bright green four-slot connectors. Bright green—
R: That's in this cabinet here.
P: [Walks to specified cabinet.] This cabinet here?
R: In the left column, second drawer from the top.
P: [Pulls a piece out of the drawer to show to Chip.] Is this—whoa man—is this a bright green four-slot connector?
R: Yes.
P: Okay. I'm going to put 12 of these into, uh, B. [Counting quietly to self, only gets to 11.] We might have a problem. [Counts pieces again, still only has 11 pieces.] Yeah, is there any other drawer that has the bright green, uh, four-slot connectors?
R: No.
P: We're, uh, can we use a dark green four-slot connector?
R: Yes.
P: [Puts all pieces in the bin.] Okay, I'm gonna use one dark green four slot connector then. Okay. Now we need …

---

operator's headset. Using this *Wizard of Oz* technique, the experimenter and operator acted as if the robot were autonomous.

To manipulate the human likeness versus machine likeness of the robot, we altered the robot's appearance by replacing the outer covering. Pictures of the human-like and machine-like robots are provided in Figure 2. In the human-like robot condition, the robot had a face with eyes, nose, and mouth. It also had ears and a full head of hair. The main part of the robot had a torso, arms, and legs. Its texture was soft fabric. It had a White male appearance and wore a denim shirt, khaki pants, and a baseball cap (the human confederate was a White male and dressed in similar clothes). In the machine-like robot condition, the robot covering was metallic and angular. The main part of the machine-like robot was encased in a silver box. From research conducted by DiSalvo, Gemperle, Forlizzi, and Kiesler (2002), human-like facial features such as a nose, eyelids, and a mouth account for most of the variance in the perception of human likeness in robot heads. Our human-like robot had a nose, eyelids, and a mouth, whereas our machine-like robot had none of these features.

To confirm that our manipulation of robot appearance was effective, we conducted a pilot study. We took each robot (human-like and machine-like) to a public plaza at Stanford University, where it interacted briefly with people. The interaction consisted of the robot approaching a person and saying that it was developing its language skills. It asked the respondent to describe three objects that it was carrying on its tray. After the respondent described the objects (it did not matter what the person said, although if the person did not say very much, the robot prompted the person with, "Can you tell me more about

Figure 2. **Photographs of the human-like and machine-like robots.**

it?"), the robot directed the person to a table, where she or he filled out a survey that contained a series of questions asking his or her opinions about the robot. Each person was rewarded with a premium-quality chocolate bar. The same operator was used for both conditions, and he always followed the same script and guidelines for what to say. Our dependent variable for the pilot study consisted of seven phrases describing the robot as either human like or machine like (see Figure 3). For each question, there was a 7-point scale ranging from 1 (*strongly disagree*) to 7 (*strongly agree*). A reliable scale ($\alpha = .80$) was created for human likeness by averaging across the seven items. In the pilot study, there were 94 respondents: 46 interacting with the human-like robot and 48 interacting with the machine-like robot. Forty-six percent of the respondents were men, and the mean age was 26.95 years. The results of the pilot study confirmed that our human-like robot was perceived as significantly more human like compared with our machine-like robot. Our human-like robot was rated on average 3.69 ($SD = 0.91$), and our machine-like robot was rated on average 2.90 ($SD = 0.83$) on our 7-point scale of human likeness. An analysis of variance (ANOVA) shows a strong statistical difference between the ratings, $F(1, 92) = 18.95$, $p < .001$.

Status

Status was manipulated in the written instructions by telling the participant that their partner was their supervisor, their peer, or their subordinate. Such minimal labels have been used successfully in previous research to create status effects (see Sande et al., 1986).

To check our status manipulation, we asked participants two questions about the extent to which they were assigned a leadership versus a subservient role on the task ($\alpha = .78$). Consistent with our planned manipulation, the results indicate that participants who were told that they were working with a subordinate confederate rated themselves as 4.54 ($SD = 1.45$), whereas those who were told that they were working with a supervisory confederate rated themselves as 3.58 ($SD = 1.67$) on our 7-point leadership scale. When told that they were working with a peer, participants rated their own leadership in the middle, 4.36 on average ($SD = 1.39$). A regression analysis suggests a strong linear trend, $\beta = (1, 291) = -.24$, $p < .001$, in the desired direction.

3.4. Measures

Our primary dependent variables were reliance on the partner and sense of responsibility for the task. Reliance was measured based on behaviors coded from the videotapes recorded during the task. We looked particularly at reliance in the more ambiguous situations, those in which the participant could

Figure 3. **Table of survey questions used to create scales.**

Scales	Cronbach's α
Human likeness	.80
To what extent does the robot	
have human-like attributes?	
look like a machine or mechanical device?[a]	
have characteristics that you would expect of a human?	
look like a person?	
have machine-like attributes?[a]	
act like a person?	
act like a machine?[a]	
Responsibility	.77
To what extent did you feel it was your job to perform well on the task?	
To what extent did you feel ownership for the task?	
To what extent did you feel that your performance on this task was out of your hands?[a]	
To what extent did you feel that good performance relied largely on you?	
To what extent did you feel obligated to perform well on this task?	
Attribution of credit	.66
Our success on the task was largely due to the things I said or did.[a]	
I am responsible for most of the things that we did well on this task.[a]	
Our success on this task was largely due to the things my *partner* said or did.	
My *partner* should get credit for most of what we accomplished on this task.	
Attribution of blame	.85
I hold my *partner* responsible for any errors that we made on this task.	
My *partner* is to blame for most of the problems we encountered in accomplishing this task.	

Note. Where *partner* is indicated, this word was replaced with either *subordinate*, *peer*, or *supervisor* depending on the status condition

[a] The item was reverse scored.

choose to solicit input or not. We designed the task such that there were five anomalies, providing the basis on which we could assess the extent to which the participant relied on the confederate in ambiguous situations. Therefore, we coded for behaviors indicating that the participant neglected to consult the confederate when these anomalies occurred and the confederate had better information. We then reverse scored this variable. The videotapes were coded

by a single rater, but 10% were coded by another rater to assess reliability (Cohen's $\kappa = .81$).

Sense of responsibility was measured directly and indirectly from questions on the posttask survey. All survey items were measured on 7-point scales ranging from 1 (*less*) to 7 (*more*) of the item. Our first measure asked participants directly about how responsible they felt for the task and for performance on the task (see Figure 3). These five items were then averaged to create a scale measuring their sense of *responsibility*. As a less direct indicator of responsibility, we measured the extent to which people attributed credit to their partner and to themselves (reverse scored). We then averaged across these four items (see Figure 3) for a measure of *attribution of credit*. We also reasoned that, although blame is not equivalent to the abdication of responsibility, people who feel more responsible for a task are less likely to attribute all of the blame for errors to their partner (see Goodnow, 1996), so we measured the extent to which participants assigned blame to their partners. Two items (see Figure 3) were averaged together to create a reliable *attribution of blame* scale.

To better evaluate the theory underlying our predictions, specifically that human-like robots as compared with machine-like robots would be relied on more because people would feel they were more similar to themselves, we coded the videotapes for *shared social identity*. We did this by counting the number of times participants used individualistic language such as *I, my, you,* and *your* and the number of times participants used more collective language such as *us, we,* and *our*. We reasoned that participants who felt a stronger sense of shared identity with the confederate would use more collective language, and that those who felt more distant from the confederate would use more individualistic language. In coding for individualistic and collective language, the coders covered the video screen, coding only from the audio, so that they were blind to condition and not biased by nonverbal cues from the tapes. For both measures, the videotapes were coded by a single rater, but 10% were coded by another rater to assess reliability (Cohen's $\kappa = .99$ and .95, respectively). In our *individualistic language* variable, there were two extreme outliers (in excess of 100 uses of individualistic language). Although these outliers worked in favor of our hypotheses, we removed them to allow a more conservative analysis of the data.

Similarly, to support theories about common ground, we coded the tapes for the number of factual questions asked (Cohen's $\kappa = .85$). We reasoned that more factual questions would be asked to establish common ground when it was thought to be missing or felt less strongly.[2]

---

2. We also coded nonverbal indicators from the videotapes, including the proximity of the participant to the robot and the extent to which the participant appeared engaged in the task. These measures resulted in no significant effects, so they were excluded to simplify presentation of the results.

We also included two control variables in our analyses—gender and mood. Mood can have a significant effect on the attributions that people make (e.g., Forgas, Bower, & Moylan, 1990), so we included a six-item indicator of the participant's mood ranging from 1 (*a very negative mood*) (depressed, sad, etc.) to 7 (*a very positive mood*) (happy, excited, etc.). Mood correlated with, and was therefore included in, regressions predicting responsibility ($r = .15$, $p < .01$). Gender was negatively correlated with attribution of credit ($r = -.13$, $p < .10$), with men attributing somewhat less credit than women; therefore, gender was included in regressions predicting attribution of credit.

4. RESULTS

Figure 4 displays the descriptive statistics for and correlations between variables. As expected, participants relied more on human as compared with robotic confederates. Reliance and responsibility, however, were not strongly correlated ($r = -.06$) and could, therefore, be treated as separate constructs. Consistent with our arguments, responsibility was associated with less attribution of credit and blame, suggesting that when people feel more personally responsible for a task, they attribute less credit and blame to others.

4.1. Effects of Human Likeness

Not surprisingly, we found that participants relied more on the human partner ($M = 4.73$, $SD = 0.56$; see Figure 5) than on robot partners ($M = 4.50$, $SD = 0.86$), and the difference was significant in a two-way ANOVA with human versus robot and status as factors, $F(2, 272) = 5.09$, $p = .03$. We found, however, little difference in peoples' feelings of responsibility ($M = 4.81$ vs. 4.83), and only small differences in the extent to which they attributed credit and blame to the human versus the robot partner ($M = 4.54$ vs. 4.31 and $M = 3.27$ vs. 3.34, respectively). An ANOVA with human confederate versus robot predicting responsibility confirms that this effect is not significant even when controlling for mood, $F(2, 285) = .03$, $p = .88$. A similar analysis confirms a small but nonsignificant effect when predicting the attribution of credit to human as compared with robot partners, $F(2, 287) = 3.32$, $p = .07$, and no significant effect for blame, $F(2, 287) = .18$, $p = .68$.

In developing our hypotheses, we argued that the extent to which the robot appears human like as compared with machine like will affect participants' willingness to rely on it. There was, however, little difference in the extent to which people relied on the human-like as compared with the machine-like robot ($M = 4.60$ vs. 4.42), $F(2, 180) = 1.84$, $p = .18$. We also predicted that people would cede more responsibility to a human-like as compared with a machine-like robot. Our analyses of participants' sense of responsibility support this hypothesis. When collaborating with the human-like as compared with the machine-like robot, participants re-

Figure 4. **Descriptive statistics and correlations.**

Variable	M	SD	1	2	3	4	5	6	7	8	9	10	11
1. Human likeness[a]	—	—											
2. Human versus robot[b]	—	—											
3. Human-like vs. machine-like robot[b]	—	—											
4. Status[a]	—	—											
5. Reliance[c]	4.58	.78	.16*	.14**	.10	-.04							
6. Responsibility[d]	4.82	1.08	-.08	-.01	-.18**	-.01	-.06						
7. Attribution of credit[d]	4.39	1.00	.17*	.11***	.20*	-.13**	.09	-.30*					
8. Attribution of blame[d]	3.32	1.39	-.01	-.03	.03	.17*	-.06	-.15**	.18*				
9. Collective language[e]	3.06	6.06	-.13**	-.11***	-.08	-.03	.08	.08	-.03	.04			
10. Individualistic language[f]	8.10	9.70	-.24*	-.25*	-.04	-.07	.06	.14**	-.12***	-.04	.28*		
11. Mood[g]	6.57	1.34	-.03	-.01	-.06	-.05	.04	.15**	.07	-.05	.12**	.08	
12. Gender[h]	.41	—	<.01	<.01	-.02	.05	-.12***	.11***	-.13**	.08	.09	-.04	.03

[a]Human likeness and status are ordered scales in which human and supervisor = 3, human-like robot and peer = 2, and machine-like robot and subordinate = 1. [b]These are dichotomous indicators in which human = 1 and robot = 0, and human-like = 1 and machine-like = 0, respectively. [c]The scale for reliance was 1 to 5 with 5 equal to high levels of reliance. [d]These variables all were measured on 7-point scales ranging from 1 (low levels of the item) to 7 (high levels of the item). [e]The scale for collective language measured the number of utterances of terms such as *us* and *we*. It ranged from 0 to 47. [f]The scale for individualistic language ranged from 0 to 48 utterances. [g]Mood was measured on a 10-point scale in which 10 equates to a very happy, positive mood. [h]We used a 0, 1 scale for gender where 0 = female and 1 = male.

*p < .01. **p < .05. ***p < .10.

169

Figure 5. **Means and standard deviations by human likeness and status.**

Condition	Reliance		Responsibility		Attribution of Credit		Attribution of Blame		Collective Language		Individualistic Language	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Human partner												
Subordinate[a]	4.69	.64	4.65	.87	4.51	1.26	3.20	1.36	2.31	4.09	4.72	3.85
Peer[a]	4.81	.48	4.85	1.37	4.76	1.09	3.12	1.41	2.35	4.13	6.39	8.49
Supervisor[b]	4.69	.54	4.95	1.15	4.35	1.06	3.50	1.67	1.69	2.44	5.03	4.88
Total human partner	4.73	.56	4.81	1.14	4.54	1.14	3.27	1.48	2.13	3.64	5.38	6.06
Human-like robot												
Subordinate[b]	4.60	.81	4.62	1.11	4.70	.98	3.08	1.20	3.59	5.17	11.90	12.27
Peer[a]	4.70	.70	4.68	1.04	4.45	.83	3.13	1.35	2.30	3.06	10.00	9.32
Supervisor[a]	4.48	.78	4.60	1.34	4.35	.81	3.94	1.54	2.93	5.92	8.70	8.92
Total	4.60	.76	4.64	1.16	4.50	.88	3.39	1.42	2.93	4.82	10.23	10.26
Machine-like robot												
Subordinate[a]	4.42	.93	5.24	.72	4.24	1.06	3.18	1.31	5.00	10.99	13.15	14.54
Peer[b]	4.50	.92	4.92	.79	4.34	.62	3.00	1.21	2.31	5.20	9.41	10.66
Supervisor[b]	4.34	.97	4.88	1.13	3.80	.99	3.72	1.20	4.81	7.76	10.78	8.57
Total	4.42	.93	5.02	.90	4.13	.93	3.30	1.27	4.05	8.36	11.13	11.55
Total robot partners	4.50	.86	4.83	1.05	4.31	.92	3.34	1.34	3.52	6.91	10.71	10.94

Dependent Variables

[a]$n = 33$. [b]$n = 32$.

ported lower levels of personal responsibility ($M=4.64$ vs. 5.02). An ANOVA contrasting these two conditions confirms that the effect is significant, $F(2, 189)=6.37$, $p=.01$, even when mood is included in the analysis, $F(2, 187)=5.29$, $p=.02$. Our measure of attribution of credit shows a similar pattern with less credit being attributed to the robot in the machine-like as compared with the human-like conditions ($M=4.13$ vs. 4.50). An ANOVA contrasting the human-like and machine-like robot conditions suggests that the effect on attribution of credit, $F(2, 189)=8.41$, $p=.004$, is significantly different in the two conditions even when gender is included as a covariate, $F(2, 188)=8.35$, $p=.004$. There was not, however, a significant difference when predicting attribution of blame, $F(2, 188)=.24$, $p=.63$. Therefore, although little support is provided for Hypothesis 1a, substantial support is provided for Hypothesis 1b in which we argue that robots with a human-like as compared with a machine-like appearance will reduce the extent to which people feel responsible for the task.

In an additional analysis of the effect of appearance, we estimated linear regression equations including the human confederate (baseline condition) as part of the human-likeness scale (machine-like robot, human-like robot, human confederate). Doing so allowed us to test the reasoning that people would rely more on partners and share more responsibility for the task when the partners were more human-like (including being human). Although no strong linear effect was found for responsibility, there was a strong positive relation with reliance ($\beta=.16$, $p=.007$) and attribution of credit ($\beta=.17$, $p=.004$), indicating that people exhibited more reliance and attributed more credit to their partners as their partners displayed more human-like (or human) characteristics.

In the logical arguments leading to our hypotheses, we reasoned that people will feel a stronger sense of social identity with human-like robots than they will with machine-like robots, and that shared identity might contribute to more reliance and shared responsibility. From the videotapes, we coded individualistic language (e.g., *I*, *me*, *you*, *your*) and collective language (e.g., *us*, *we*, *our*) as a means of measuring the extent to which the participants were expressing a sense of shared social identity with the robot. Although not statistically significant, participants used less individualistic language with human-like robots than with machine-like robots ($M=10.23$ vs. 11.13), $F(2, 177)=.31$, $p=.58$.[3] However, they also used less collective language ($M=2.93$ vs. 4.05), $F(1,$

---

3. We also conducted the analyses including the two outliers (in excess of 100 uses of individualistic language). Both of the outliers were in the human-like robot supervisor conditions. When included, the mean for the human-like robot supervisor condition is 17.76 ($SD=35.04$), and the mean for the human-like robot conditions is 13.18 ($SD=21.99$). When we include the outliers in the analysis of variance, the difference between the human-like and machine-like robot conditions is not statistically significant, $F(2, 179)=.68$, $p=.41$.

179) = 1.42, p = .25. These data provide conflicting results and, thus, little support for our arguments about shared identity being at the root of differences in reliance and responsibility when working with human-like versus machine-like robots. Normalizing the data to get the fraction of all pronouns used that were collective or individualistic produced no statistically significant effects. The pattern of these data suggests that people were more talkative overall with machine-like robots as compared with human-like robots. The data could be interpreted as support for a common ground explanation. That is, when people perceive less common ground between themselves and a partner, they tend to talk more (Fussell & Krauss, 1992). They talk more, in part, to provide the partner with the background required to interpret future interactions and, in part, to gather more information about what the partner knows. To evaluate this possibility further, we coded the videotapes for the number of factual questions that the participant asked of the confederate. Consistent with a common-ground explanation, participants who worked with the human-like robot asked fewer factual questions on average (M = 2.58) than those who worked with the machine-like robots (M = 3.45), although this difference was not significant, $F(1, 182) = 1.29$, p = .26.

4.2. Effects of Status

In Hypotheses 2a and 2b, we argued for a main effect of status. We posited that people will rely more on a robotic partner and feel less responsible for the task when the partner is assigned a supervisory as compared with a subordinate or peer position relative to the participant. We found little support for Hypothesis 2a. People relied more on the robot peers than they did robot subordinates or supervisors (M = 4.41 supervisors vs. 4.60 peers vs. 4.51 subordinates), and the difference between the supervisor and the other status conditions was not significant, $F(1, 182) = 1.13$, p = .29. Analyzing Hypothesis 2b, participants reported feeling less responsible when collaborating with a robot supervisor as compared with a robot peer or subordinate (M = 4.74, 4.80, 4.94, respectively), but participants also reported that less credit was due to the partner when it was a supervisor, which is the opposite of what we had hypothesized (M = 4.08, 4.39, 4.47, respectively). Although the effect for responsibility was not significant, $F(1, 191) = .66$, p = .42, when conducting two-way ANOVAs contrasting the supervisor condition with the other status conditions, the effect of status on attribution of credit was significant, $F(1, 191) = 6.73$, p = .01. That is, participants attributed significantly less credit to the robot supervisor as compared with the robot peer and subordinate. Paradoxically, we also found that participants were more likely to blame robot supervisors as compared with robot peers and subordinates for errors and mistakes that were made (M = 3.83, 3.07, 3.13, respectively) and that this differ-

ence was highly significant, $F(1, 191) = 13.53$, $p < .001$. Overall, it seems that participants were much more critical of the robot in a supervisory position.

As with our Hypotheses 1a and 1b, we had an underlying linear assumption in our status variable, suggesting that status would increase the extent to which people relied on the partner and would decrease their sense of responsibility for the task. Therefore, we conducted linear regressions with status as a scale (supervisor, peer, subordinate) and human likeness as a variable. The only significant linear effects were found when predicting attribution of credit ($\beta = -.137$, $p = .001$) and attribution of blame ($\beta = .21$, $p = .003$), suggesting that as the status of the robot target increases, people attribute less credit for successes and more blame for failures in performance.

4.3. Interaction Between Human Likeness and Status

Our final hypothesis predicted an interaction effect between human likeness of the robots and the robot's status. In Hypothesis 3, we argued that people will feel most responsibility for the task when they work with a machine-like robot subordinate. To test this hypothesis, we conducted a two-way ANOVA analysis with only the robot conditions included in the analysis. We found a significant effect in the expected direction for responsibility when contrasting the machine-like subordinate conditions with all other robot conditions, $F(1, 193) = 6.37$, $p = .01$. We found little effect for attribution of credit, $F(1, 193) = .26$, $p = .61$; or attribution of blame, $F(1, 193) = .58$, $p = .45$. These analyses provide some evidence that people will feel most responsible for the task when they are collaborating with a machine-like robot that is presented as a subordinate. A similar test predicting reliance showed no significant effect, as we expected, $F(1, 184) = .36$, $p = .55$.

A summary of the support found for each of our hypotheses is detailed in Figure 6.

5. DISCUSSION

As far as we know, ours is one of the first systematic, controlled experiments comparing people's responses to human coworkers versus robot coworkers; and to more humanoid versus less humanoid robots. Our findings suggest that there are significant differences in the extent to which people will rely on robots as compared with human work partners. When working with a person instead of a robot, participants relied more on the partner's advice and were less likely to ignore their counsel. We found, however, only marginal support for the idea that people would feel less burden of responsibility when interacting with another person as compared with a robot. It appears from these data that

Figure 6. **Summary of hypotheses and results.**

Hypotheses	Results
Human-like versus machine-like robots	
Hypothesis 1a: People will rely on a human-like robot partner more than on a machine-like robot partner.	Not supported
Hypothesis 1b: People will feel less responsible for the task when collaborating with a human-like robot partner than with a machine-like robot partner.	Supported
Relative status of robot coworkers	
Hypothesis 2a: People will rely on the robot partner more when it is characterized as a supervisor than when it is characterized as a subordinate or peer.	Not supported
Hypothesis 2b: People will feel less responsible for the task when collaborating with a robot partner who is a supervisor than with a robot partner who is a subordinate or peer.	Mixed support
Interactions between human-likeness and status	
Hypothesis 3: People will feel the greatest amount of responsibility when collaborating with machine-like robot subordinates as compared with machine-like robot peers and supervisors; and as compared with human-like robot subordinates, peers, and supervisors.	Mixed support

participants collaborated more with the human partners than with the robot partners but still did not necessarily cede responsibility to them.

Comparing robots with different appearances, our data show that interacting with a more machine-like robot may increase the personal responsibility that people feel for the task. This effect was increased when status was added to the manipulation. Our data indicate that participants felt most responsible when interacting with the machine-like subordinate, suggesting that a machine-like appearance coupled with the framing of a subordinate position may result in the highest levels of human responsibility. Knowing this may be useful when it is important for workers to feel the full weight of responsibility for the task in which they are engaged. When people feel more responsibility for the task, mishaps may be avoided because people explore more options and are more diligent about finding an appropriate solution (see Roberts et al., 1994). When people feel responsibility for the task, they may also be less likely to trust the robotic partner to perform tasks for which the robot is ill equipped or when the robot becomes ill equipped due to unanticipated changes to the task requirements or the environment. On the other hand, our results suggest that humanoid robots may be appropriate for situations in which the burden of responsibility can or should be attenuated for the people involved in the task.

For example, humanoid robots may be preferred when human complacency is not a concern or when the consequences of a risky task would be too difficult for a human to bear.

We hypothesized, as have others (e.g., Parise et al., 1996), that one of the reasons that people may respond differently to human-like technologies than to machine-like technologies is because they feel more similar to the former and thus experience more shared identity with them. Our behavioral measures of shared identity, however, provided little evidence to this effect. We did not see a difference in the individualistic or collective language used by participants in the different conditions. On the other hand, participants in the machine-like robot conditions appeared to talk more with the robot than did participants in the human-like robot conditions, suggesting that they might have perceived less common ground with these robots, and felt they had to explain themselves more, or provide more instruction.

Our status manipulation generated mixed effects. When collaborating with supervisors, participants attributed less credit and more blame to their partner. This effect suggests that, as in the popular Dilbert cartoon (see www.dilbert.com), the supervisor was viewed as undeserving and was blamed for most of the problems encountered on the task. Additional research is needed to assess the robustness of this effect. It is possible that the effect in our study is specific to the task being performed and to the role we assigned to the supervisor. The task was relatively straightforward, albeit with some ambiguity, and did not allow the supervisor to display a particularly impressive level of skill or competence. Our results may suggest that when people or robots are put in supervisory positions without commensurate skills or authority, their subordinates will respond negatively. A situation in which the partner has substantially more skill or knowledge relative to the participant might reveal somewhat different effects. It is also possible, however, that the comparatively young student population in our sample have beliefs that are consistent with the "Dilbert effect," maintaining fairly negative impressions of those in supervisory roles.

This study is an early foray into the examination of people's responses to professional service robots. As such, many questions remain, and additional studies are needed to fully understand how people will respond to and work with professional service robots on collaborative tasks. One area that merits exploration, for example, is the nature of the task being performed. The task we utilized here was a relatively mundane parts-sorting task that required movement and knowledge on the part of the robot. Even in the ambiguous parts of the task, participants demonstrated high levels of reliance on their partner ($M = 4.58$ out of 5). We expect that reliance was high because the task was routine and posed little risk, and the robot was clearly equipped to perform the role assigned to it. Over time, robots are likely to assume tasks that

are substantially more complex, risky, and uncertain than this experimental task (e.g., Burke, Murphy, Coovert, & Riddle, 2004). Situations in which people working with a robot are already cognitively overloaded and do not have the capacity to monitor the robot's actions are also likely. For example, search and rescue robots may work in tandem with rescue workers in extreme weather conditions, such as in the aftermath of a hurricane. In such situations, people are overworked and experience stress, uncertainty, and physical danger. Based on the research reported here, we anticipate that humanoid robots may be appropriate for tasks that are complex or risky because people will more readily delegate responsibility to them. Although our task did not present high risk to the participants, our participants reported being reasonably well motivated to participate ($M = 6.97$ on a 10-point scale); we anticipate that stronger motivation and higher risk might strengthen these results. Substantial research, however, is needed to fully understand the interplay between the design of the robot, the task being performed, the interaction between the person and the robot's skill and knowledge, the amount of perceived risk, and people's willingness to rely on the robot.

Although the focus of our research was on robots in the work environment, this study was conducted in a controlled laboratory setting. Doing so enabled us to maintain control over the conditions being tested. At the same time, realistic aspects of the work environment were intentionally designed out. For example, participants in our study worked in dyads and did not interact with other co-workers or friends. We believe that people's responses to robots in the work environment will be significantly influenced by the social and organizational context in which they are embedded (see Siino & Hinds, 2004). Robots also could have a significant and unanticipated effect on the dynamics of work teams. Existing research suggests that the nature of effective team processes may be different when automated systems are introduced (see Bowers, Oser, Salas, & Cannon-Bowers, 1996). It will be fruitful to investigate the effect of human-like and machine-like robots on the dynamics of teams and organizations.

In examining human likeness versus machine likeness of the robot, we chose to create robots that were a composite of a variety of human-like and machine-like features. The human-like robot had facial features, a torso, arms, and legs. It had the appearance of a man clothed in a denim shirt, khaki pants, and a baseball cap. The machine-like robot was metallic and angular with none of the physical human-like features previously described. We created these two conditions to make an initial determination of the impact of human likeness in a robot partner. We believe, however, that it will be important to decompose and examine the independent effects of features such as eyes, mouth, and legs. One also could look individually at features that we held constant: a human-like voice, the absence of human-like gestures, and the robot's movements (e.g., rolling vs. walking). Along with others (e.g., Jensen, Farn-

ham, Drucker, & Kollock, 2000; Nass & Lee, 2000), we think it quite likely that each of these features may have some effect on perceived human likeness and on perceptions of the robot coworker. Given the power of voice, it is possible that the effects that we found could be duplicated by manipulating only the voice (natural human voice vs. machine-generated voice). Similarly, behavioral manipulations may provide a powerful way to convey human likeness in robots (see Breazeal, 1999) and may generate similar effects. Research into each of these dimensions would be a contribution to this field of work.

It is also important to note that although we manipulated human likeness of the robot, the robot we used was nowhere near as human like as advanced robots are and can be. For example, Honda's ASIMO has a more humanoid form than we were able to produce in this study. In our pretest, our machine-like robot was rated on average 2.90, and our human-like robot on average 3.69, on a 7-point scale of human likeness. It is possible that our effects for reliance are weak because our human-like robot was not extremely human like and the task did not require advanced human-like behaviors. We anticipate that future research examining human reliance using increasingly humanoid robotic forms and behaviors will find stronger effects. This work also will help us understand the so-called uncanny valley—a space in which robots evoke expectations of human likeness but are not quite human and, therefore, create discomfort (see DiSalvo et al., 2002; Reichard, 1978).

Other factors that we anticipate will have a significant effect on human–robot interaction are the expectations that people have of robots and the experience they have with them. As people gain experience, the novelty of the technology wears off and people develop ways to adapt the technology to better fit their needs (e.g., Barley, 1986; Orlikowski, 2000). In this study, having some experience with robots (e.g., a class or two) did not affect our results. Having taken a class or two about robotics, however, may not be enough to establish clear expectations about the capabilities of robots in general or about a specific robot in particular. As people gain more experience with robots and with the particular robot with which they are working, we expect they will develop new mental models of the robots' capabilities, revise their perceptions of how the robot fits into the work environment, and make alterations to the robot or their use of it to better accommodate their needs. This gain in experience will no doubt affect people's collaboration with the robot. Ideally, longitudinal studies will help to inform these questions.

Although there is substantial research required to fully understand human–robot interaction on collaborative tasks, the research reported here provides some initial findings that suggest how people may respond to humanoid professional service robots and the conditions under which humanoid robots may be preferable to machine-like robots. These findings provide input to

guide the decisions of those designing robots and determining the roles that robots will play in the work setting.

More broadly, we view this research as an early effort using the theoretical and empirical foundations of social psychology and organizational behavior to inform the design of robots for the work environment. We believe that continued work in this area and ongoing collaboration between social scientists and roboticists to identify and explore questions of importance can make a significant contribution to the field of robotics and improve the likelihood for successful implementation of professional service robots.

NOTES

Authors' Present Addresses. Pamela J. Hinds, Management Science & Engineering, Terman 424, Stanford University, Stanford, CA 94305–4026. E-mail: **phinds@stanford.edu**. Teresa L. Roberts, PeopleSoft, Inc., 4460 Hacienda Drive, Pleasanton, CA 94588–8618. E-mail: **terry_roberts@peoplesoft.com**. Hank Jones, Aerospace Robotics Laboratory, Durand 250, Stanford University, Stanford, CA 94305. E-mail: **hlj@arl.stanford.edu**.

REFERENCES

Ambrose, R. O., Askew S. R., Bluethmann, W., & Diftler, M. A. (2001). Humanoids designed to do work. *Proceedings of the IEEE 2001 International Conference on Humanoid Robots, Robotics and Automation Society.* Tokyo, Japan.

Barley, S. (1986). Technology as an occasion for structuring: Evidence from observation of CT scanners and the social order of radiology departments. *Administrative Science Quarterly, 31,* 78–108.

Bowers, C. A., Oser, R. L, Salas, E., & Cannon-Bowers, J. A. (1996). Team performance in automated systems. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 243–265). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Breazeal, C. (1999). Robot in society: Friend or appliance? *Proceedings of the Agents 99 Workshop on Emotion-Based Agent Architectures.* Seattle, Washington.

Breazeal, C., & Scassellati, B. (1999). *How to build robots that make friends and influence people.* Paper presented at the 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS–99), Kyonjiu, Korea.

Brooks, R. (2002). Humanoid robots. *Communications of the ACM, 45*(3), 33–38.

Brooks, R., & O'Reilly, U.-M. (2002). *Humanoid robotics group, MIT artificial intelligence laboratory.* Retrieved November 11, 2002, from http://www.ai.mit.edu/projects/humanoid-robotics-group/

Burke, J. L., Murphy, R. R., Coovert, M. D., & 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,* 85–116. [this special issue]

Dautenhahn, K. (1999). Robots as social actors: AURORA and the case of autism. *Proceedings of the Third International Cognitive Technology Conference.* San Francisco, CA.

DiSalvo, C. F., Gemperle, F., Forlizzi, J., & Kiesler, S. (2002). All robots are not created equal: The design and perception of humanoid robot heads. *Proceedings of the DIS 2002 Conference on Designing Interactive Systems.* New York: ACM.

Fong, T., Nourbakhsh, I., & Dautenhahn, K. (2002). *A survey of social robots* (CMU Robotics Institute Technical Report CMU–RI–TR–02–29). Pittsburgh, PA.

Forgas, J. P., Bower, G. H., & Moylan, S. J. (1990). Praise or blame? Affective influences on attributions for achievement. *Journal of Personality and Social Psychology, 59,* 809–819.

Fussell, S., & Krauss, R. M. (1992). Coordination of knowledge in communication: Effects of speakers' assumptions about what others know. *Journal of Personality and Social Psychology, 62,* 378–391.

Gawande, A. (2002). *Complications: A surgeon's notes on an imperfect science.* New York: Holt.

Goetz, J., Kiesler, S., & Powers, A. (2003). Matching robot appearance and behavior to tasks to improve human–robot cooperation. *Proceedings of the 12th IEEE Workshop on Robot and Human Interactive Communication. RO-MAN 2003.* San Francisco, CA.

Goodnow, J. J. (1996). Collaborative rules: From shares of the work to rights to the story. In P. Baltes & U. Staudinger (Eds.), *Interactive minds* (pp. 163–193). Cambridge, UK: Cambridge University Press.

Grabowski, M., & Roberts, K. (1997). Risk mitigation in large-scale systems: Lessons from high reliability organizations. *California Management Review, 39,* 152–162.

Hashimoto, S., Narita, S., Sugano, S., Takanishi, A., Shirai, K., & Kobayashi, T. (2002). *History of humanoid robot at Waseda University* (Humanoid Robotics Institute, Waseda University, Japan). Retrieved November 11, 2002, from http://www.humanoid.waseda.ac.jp/history.html

Isaacs, E., & Clark, H. (1987). References in conversation between experts and novices. *Journal of Experimental Psychology: General, 116,* 26–37.

Ishiguro, K. (2003). *Intelligent robotics laboratory, Osaka University.* Retrieved July 28, 2003, from http://www.ed.ams.eng.osaka-u.ac.jp/

Jensen, C., Farnham, S. D., Drucker, S. M., & Kollock, P. (2000). The effect of communication modality on cooperation in online environments. *Proceedings of the CHI 2002 Conference on Human Factors in Computer Systems.* New York: ACM.

Kiesler, S., & Goetz, J. (2002). Mental models of robotic assistants. *Proceedings of the CHI 2002 Conference on Human Factors in Computer Systems*. New York: ACM.

Kraut, R., Fussell, S., Lerch, F., & Espinosa, A. (2002). *Coordination in teams: Evidence from a simulated management game* (Working Paper). Retrieved May 26, 2003, from http://www.andrew.cmu.edu/user/sfussell/Manuscripts_Folder/Team_Coordination.pdf

Nass, C., Fogg, B. J., & Moon, Y. (1996). Can computers be teammates? *International Journal of Human–Computer Studies, 45,* 669–678.

Nass, C., & Lee, K. M. (2000). Does computer-generated speech manifest personality? An experimental test of similarity-attraction. *Proceedings of the CHI 2000 Conference on Human Factors in Computer Systems*. New York: ACM.

Nass, C., Steuer, J., Tauber, E., & Reeder, H. (1993). Anthropomorphism, agency, and ethopoeia: Computers as social actors. *Proceedings of the INTERCHI 93 Conference on Human Factors in Computer Systems*. New York: ACM.

Okie, S. (2002, April 3). Robots make the rounds to ease hospitals' costs. *The Washington Post,* p. A03.

Olson, G. M., & Olson, J. S. (2000). Distance matters. *Human–Computer Interaction, 15,* 139–179.

Orlikowski, W. (2000). Using technology and constituting structures: A practice lens for study technology in organizations. *Organization Science, 11,* 404–428.

Parise, S., Kiesler, S., Sproull, L., & Waters, K. (1996). My partner is a real dog: Cooperation with social agents. *Proceedings of the CSCW 96 Conference on Computer-Supported Cooperative Work*. New York: ACM.

Pfeffer, J., Cialdini, R., Hanna, B., & Knopoff, K. (1998). Faith in supervision and the self-enhancement bias: Two psychological reasons why managers don't empower workers. *Basic and Applied Social Psychology, 20,* 313–321.

Reeves, B., & Nass, C. (1996) *The media equation*. Cambridge, UK: Cambridge University Press.

Reichard, J. (1978). *Robots: Fact, fiction, and prediction*. New York: Penguin.

Roberts, K. H., Stout, S. K., & Halpern, J. J. (1994). Decision dynamics in two high reliability military organizations. *Management Science, 40,* 614–624.

Sande, G. N., Ellard, J. H., & Ross, M. (1986). Effect of arbitrarily assigned labels on self-perceptions and social perceptions: The mere position effect. *Journal of Personality and Social Psychology, 50,* 684–689.

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

Sheridan, T. B. (1992). *Telerobotics, automation, and human supervisory control*. Cambridge, MA: MIT Press.

Siino, R., & Hinds, P. (2004). *Making sense of new technology: Organizational sensemaking of autonomous mobile robots as a lead-in to structuring*. Unpublished manuscript, Stanford University.

Simmons, R., & Nourbakhsh, I. (2002). *Social robots project, Carnegie Mellon University*. Retrieved November 11, 2002, from http://www-2.cs.cmu.edu/~social/

Sparaco, P. (1994). A330 crash to spur changes at Airbus. *Aviation Week and Space Technology, 141*(6), 20–22.

Squeo, A. (2001, December 13). Meet the newest recruits: Robots. *The Wall Street Journal,* p. B01.

Steuer, J. (1995). *Vividness and source of evaluation as determinants of social responses toward mediated representations of agency.* Unpublished doctoral dissertation, Stanford University, California.

Strodtbeck, F. L., James, R. M., & Hawkins, D. (1957). Social status in jury deliberations. *American Journal of Sociology, 22,* 713–719.

Swim, J. K., & Sanna, L. J. (1996). He's skilled, she's lucky: A meta-analysis of observers' attributions for women's and men's successes and failures. *Personality and Social Psychology Bulletin, 22,* 507–519.

Tetlock, P. E. (1985). Accountability: The neglected social context of judgment and choice. In L. L. Cummings & B. M. Staw (Eds.), *Research in organizational behavior* (Vol. 7, pp. 297–332). Greenwich, CT: JAI.

Thrun, S. (2004). Toward a framework for human–robot interaction. *Human–Computer Interaction, 19,* 9–24. [this special issue]

Turkle, S. (1984). *The second self: Computers and the human spirit.* New York: Simon & Schuster.

United Nations and The International Federation of Robotics. (2002). *World Robotics 2002.* New York: United Nations.

Wall, T. D., Jackson, P. R., & Davids, K. (1992). Operator work design and robotics system performance: A serendipitous field study. *Journal of Applied Psychology, 77,* 353–362.