

Final Report for the DARPA/NSF Interdisciplinary Study on Human–Robot Interaction

Jennifer L. Burke
Vladimir J. Lumelsky

Robin Roberson Murphy
Jean Scholtz

Erika Rogers

Abstract—As part of a Defense Advanced Research Projects Agency/National Science Foundation study on human–robot interaction (HRI), over sixty representatives from academia, government, and industry participated in an interdisciplinary workshop, which allowed roboticists to interact with psychologists, sociologists, cognitive scientists, communication experts and human–computer interaction specialists to discuss common interests in the field of HRI, and to establish a dialogue across the disciplines for future collaborations. We include initial work that was done in preparation for the workshop, links to keynote and other presentations, and a summary of the findings, outcomes, and recommendations that were generated by the participants. Findings of the study include—the need for more extensive interdisciplinary interaction, identification of basic taxonomies and research issues, social informatics, establishment of a small number of common application domains, and field experience for members of the HRI community.

An overall conclusion of the workshop was expressed as the following—HRI is a cross-disciplinary area, which poses barriers to meaningful research, synthesis, and technology transfer. The vocabularies, experiences, methodologies, and metrics of the communities are sufficiently different that cross-disciplinary research is unlikely to happen without sustained funding and an infrastructure to establish a new HRI community.

Index Terms—Man-machine systems, mobile robots, research and development, technology social factors.

I. INTRODUCTION

WHILE human–robot systems have always been an active area in robotics, in the last ten years we have seen even more emphasis in this area. The reasons for this are multiple, and sometimes seemingly contradictory, and have to do with weaknesses and strengths of the two components of the human–robot

pair, i.e., the relatively slow advance in recent robotic intelligent systems, the realization that more reliance on automation at the expense of human skills is a likely way to succeed in some critical areas of technology (such as military or space exploration systems), the realization that instilling intelligence in robots is harder than we first thought, while on the other hand, human intelligence is not infallible and in some areas (such as spatial reasoning) trails robot/computer intelligence. Today more than ever, human–robot systems seem to present a viable response to many pressing needs of automation. Realizing these needs and creating effective human–robot teams will require joining forces of the communities that have had relatively little contact so far—researchers in robotics and computer science, human–computer interaction, cognitive science, psychology. By creating teams in which the expertise of humans and robots can be leveraged, we will have systems that are more capable than either humans or robots alone.

Slowly (as some may feel) but surely, robots continue being introduced into many aspects of our society, from military uses to medicine; from entertainment to home and office laborers; for use on land, sea, air, and space. The hardware, the software, and the sensor technologies are developing rapidly.

Robot teleoperation, still the primary mode of operation in today’s human–robot systems, can be highly successful and irreplaceable (think of the Shuttle robot arm), but these systems are also very limited and expensive. Some of them require two operators per robot and the cognitive demands of operating the robots constrain other tasks that the users might need to perform. Communication channels with robots are often limited which makes remote operations difficult. Operating in a physical, often harsh environment is difficult for robots and operator intervention is needed to assess the situation, and help the robot maneuver correctly to recover. Often it requires stopping the robot, thus, introducing interruptions and delays.

One of the big deterrents to integrating robots into human–robot systems is the current design focus. Robots have been designed from the robot point of view. While this focus was appropriate in developing the existing hardware and software robot platforms, it is not team-centric. We need to now look at how robots and humans can form synergistic teams. This does not imply that current research in robotics should stop. From the discussion of problems, it should be evident that more work is needed. This discussion will likely point to new approaches in the design of human–robot systems: more flexible kinematics from manipulators; intelligent software for more (or less) autonomy in robots; better understanding

of the division of responsibilities between human and robot intelligence; learning algorithms; and better sensors and fusion algorithms.

The new research directions require interdisciplinary work. Groups of researchers from robotics, human–computer interaction, cognitive science, psychology, and social and behavioral science are needed to begin to look at designing synergistic teams of humans and robots where team members perform tasks according to their abilities. As with teams of humans, responsibilities and roles in human–robot teams will likely be dynamic. Team members (humans and robots) need to recognize changing situations and adapt to ensure that the team mission is successful. Research is needed to bring about such capabilities. The challenge in this interdisciplinary research is to understand how to effectively combine current theories and where new theories and models need to be developed to support human–robot teams.

A list of proposed research directions would include the following:

- studies of human intervention with different levels of autonomy;
- developing and delivering cues to facilitate remote perception based on cognitive studies;
- cognitive studies on limitations of human intelligence in typical human–robot tasks (such as limitations in spatial reasoning, reaction speed, consistency, effects of fatigue, etc.);
- interaction modalities, both input and output, that depart from today’s typical means—keyboards, mice, displays—and can be used in various physical environments;
- appropriate levels of abstraction for effective but intuitive command and control of robots;
- development of roles for robots and humans within teams, based on studies of human roles, role switching, and handoff behaviors;
- adaptability of humans, robots, and human–robot teams according to the dynamic nature of situations;
- scalable user interfaces to allow one human to work efficiently with a team of multiple robots;
- designing tools for developing human–robot interfaces;
- Robot architectures and world models that support robot evolution;
- Evaluation methodologies and metrics to assess the progress of research in human–robot teams.

We make no claim that this list is comprehensive, or even correct. It is rather a starting point that researchers can use to approach the work before them. While the capabilities of robots will continue to evolve, existing systems are a good point for establishing and advancing research on human–robot projects. Although a number of research directions could be interpreted as being specific to a particular discipline, we want to emphasize that interdisciplinary coordination is a necessary condition for successful research in interactive human–robot systems.

II. STUDY OBJECTIVES

Over the last five years, numerous workshops and meetings have started to define a future vision for robotic and intelligent

machines, which must work together with humans to achieve common goals. During the same period of time, there has been an increasing interest in human-centered systems, which address issues of how to achieve synergism between man and machine, and more specifically, which take a philosophical stance on building technology that serves human needs. The term “human-centered robotics,” as used in the media, does not necessarily refer to a human-centric approach to robotic research and development. One goal of the study was to bring these sometimes disparate communities into closer contact, and to examine similarities and differences in philosophical foundations, key issues and technologies, methodologies, applications and outcomes assessment. An overview of some recent related meetings and a source of literature on human–robot interaction, including additional papers and materials provided by the participants, can be found at <http://www.crasar.org>.

The purpose of this study was to provide a forum where roboticists could interact with psychologists, sociologists, cognitive scientists, communication experts, and human-computer interaction specialists to explore the social interaction “space” between humans and robots. A sharing of ideas and philosophies, as well as more concrete methodologies and metrics, were expected to lead to the following outcomes:

- development of a taxonomy of issues, including identification of various dimensions of the field;
- identification of “grand challenges” in the area of human–robot interaction;
- development of a preliminary interdisciplinary roadmap which can guide and encourage research and development for this domain;
- establishment of a repository of resources and information tied to the roadmap.

A set of preliminary questions was provided for the participants’ consideration. These were designed to initiate ideas about issues and to promote further discussion during the workshop itself.

- In systems where humans and robots work as a team, how are tasks/responsibilities divided between the partners? How can we identify the skills needed by the robot? When can the team members (robotic or human) be interrupted? Can the robot be “over-tasked”?
- Can we establish a taxonomy of human–robot relationships, and identify what levels of “interpersonal skills” the robots will need to perform effectively in these roles?
- Does the physical form of the robot and/or its “personality” affect how people respond to it? Does the context of the relationship also play a role (e.g., workplace versus home, safety-critical versus low-impact, remote versus local, hazardous versus benign, dependent versus independent)?
- How are issues of safety and reliability impacted by the human–robot relationship? For example, what if people attribute more intelligence to the robot than it actually has?
- What about robots which are modeled on animal behaviors or which are new generation animal/machine hybrids—what kind(s) of interface(s) will allow humans to direct/control/interact with these types of robots or robotic communities?

- How can human–robot relationships be effectively studied (e.g., can the principles of user-centered design be applied to this domain)? Are there ways to build simulations in which people can *physically* interact with robots in realistic settings?
- Can we map current interface techniques (e.g., speech, vision, gesture, augmented/virtual reality, direct manipulation graphical user interfaces (GUIs), etc.) to the various types of relationships, or do we need to develop entirely new kinds of interfaces?
- What kinds of methods/metrics can be developed to gauge the utility of different types of human–robot relationships?
- Is there a role/need for a “human-centered” approach in this domain, where human-centered implies that the purpose of technology is to serve human needs?
- Who is responsible for failure, and how does this impact the kinds of interfaces needed/desired?
- Can we identify the different kinds of people who may encounter robots during their daily lives, and map the range of their attitudes toward intelligent robots?
- Can we determine where these attitudes come from, and how they are influenced by demographic factors such as culture, gender, education, age, socio-economic status?

III. ORGANIZATION OF THE STUDY

The purpose of the study was to initiate contact and discussion between traditional roboticists, cognitive scientists, and human factors experts. This interaction was organized into three phases:

- 1) preliminary online discussion phase;
- 2) workshop phase;
- 3) analysis and synthesis phase.

The bulk of the data collection and findings were generated during the workshop phase. This consisted of an invitation-only workshop, where fifty-six participants met for one and a half days (September 29–30, 2001) at California Polytechnic State University in San Luis Obispo. An additional nine participants were included remotely through video-conferencing with DARPA headquarters in Washington, DC. Invitees included representatives from government, academia, and industry, and a special effort was made to include a number of graduate students specifically working in the area of HRI. Results of the discussions were posted to the website, and the subsequent analysis and synthesis phase was conducted by the co-chairs following the workshop.

During the preliminary phase, a steering committee was formed, a website created, and a list of invited participants created. The steering committee consisted of V. Lumelsky (NSF), J. Scholtz (DARPA), R. Arkin (Georgia Institute of Technology), C. Breazeal (Massachusetts Institute of Technology), C. Nass (Stanford University), M. Peshkin (Northwestern University), and D. Woods (the Ohio State University). Preliminary informational materials were generated by the steering committee and posted on the website for invitees to examine. As part of their invitation, participants were directed to the website and asked to submit additional issues and discussion topics, links to related papers, and a summary of their expertise

and interests. The preliminary online discussion phase resulted in an expanded list of issues and several partial taxonomies, as well as links to relevant papers.

The workshop opened with introductory remarks, keynote talks, and video/special presentations. The remainder of the workshop was then spent with break-out groups (BOGs), which were used to generate discussion and foster connections among participants. The break-out groups were organized as follows: participants were divided ahead of time (by the organizers) into six groups, each of which included representatives from academia, government, industry, and graduate students, with a balance between roboticists and nonroboticists. A seventh group was made up of the invitees who were teleconferencing from Washington, and included representatives from DARPA, NSF, Science Applications International Corporation (SAIC), National Institute of Standards and Technology (NIST), and the University of Pennsylvania. Each group was given a copy of the preliminary issues and taxonomy materials previously generated, and presented with a common charge. The groups were supported by volunteer Cal Poly students who recorded the discussions for use during the analysis and synthesis phase. Each group was asked to pick a discussion leader and a recorder. At the end of the discussion session, each group presented a summary of its findings to the entire workshop followed by an open, moderated discussion.

The first BOG discussion was devoted to reviewing and expanding upon the issues surrounding HRI. It was an opportunity for the participants to acquaint themselves with each other’s areas of expertise and to try to develop a common basis for discussion. Although preliminary “raw” issues materials had been provided ahead of time, it was clear that the different disciplines viewed these issues in different ways, with different priorities, and with widely varying vocabularies.

The second breakout group session was intended to identify grand challenges (if possible) and to develop a preliminary roadmap for directing research in this area. The follow-up discussion focused on eliciting the technical goals of HRI in the next year, three years, five years, etc., identifying some of the major obstacles to achieving the goals, and gathering suggestions and ideas about how to move forward in the promotion of this nascent interdisciplinary dialogue.

IV. FINDINGS

In addition to the four expected outcomes of the study (a preliminary taxonomy, identification of grand challenges, a roadmap, and a repository) described in the next section, the study produced additional findings. As was expected, the findings raised more questions which reinforce the need for HRI to be treated as a unique area of inquiry.

A. *Interdisciplinary Interaction*

More extensive interdisciplinary interaction must be motivated. HRI is an intrinsically cross-disciplinary endeavor. There is a perceived need for cross-disciplinary education and joint work. Each community brings a different set of skills and expertise needed for the HRI enterprise. For example, the cognitive science community brings in modeling, representations, results

from human studies, and metrics of human performance. The AI/robotics community has a different set of representations, plus algorithms, embodied systems experience, and metrics of robot performance. The HCI community has concrete methodologies and usability studies, as well as metrics for measuring the performance of the interaction.

It was determined that there are many areas that the communities need to work on jointly. In particular, work is needed in system architectures that explicitly include the human in the loop. How HRI teams can cooperatively deal with uncertainty and incomplete knowledge is another major issue. Knowledge representation is a further topic, especially with respect to how we express ability in terms of humans and robots. Likewise, mental models of humans, robots, and teams are needed in order for each agent to understand the other. Situationally correct interfaces and modes of interaction are needed, including research into the modalities, mechanisms of presentation, timing, the amount of information to be conveyed, and the level of autonomy. While each community has metrics for performance, there needs to be new metrics for the HRI system.

The cross-disciplinary nature of HRI studies led the participants to formulate the following pragmatic questions which can steer the development of an HRI infrastructure.

- What would it take to get cognitive scientists to work with roboticists? What is the payback to the social science community? Is there sufficient satisfaction to be gained from the theoretical work being implemented practically in robotics?
- How do we provide roboticists with the cognitive science background they need? What kinds of cognitive experiments are appropriate for robotics platforms? Are certain applications (e.g., elder care) of more interest to cognitive science?
- What can cognitive scientists learn from AI? Are there specific cognitive theories that are amenable to robotics? Or is it better to identify the goal first?

It was determined that the programmatic context must be viewed from both the robot side and the human side in order to adequately capture the issues. The robotics puzzle can be considered as a set of problems in the areas of mobility, communications, power, robustness, reliability, sensing, perception, and understanding. An important question is: How can the human-robotic team compensate for less than perfect robotic performance in these areas? The flip-side of the coin is consideration of how HRI issues relate to the human puzzle. Unlike robots, humans exhibit fatigue and stress. They routinely handle interruptions, can perform multiple channel information fusion, multitask, deal with complexity and uncertainty, generate mental models, conduct spatial and temporal reasoning, and have situational awareness. These abilities are influenced by the user interface and training. An important question that stems from the human puzzle is: What do we know from here that can be used as a baseline for human-robotic teams?

B. Research Issues and Basic Taxonomies

Basic taxonomies and research issues must be identified. Research-related issues which should be addressed within the next

TABLE I
NUMERIC RELATIONSHIPS: HUMAN-ROBOT RATIOS

Humans	Robots
One person	One robot
One person	Many robots
Many people	One robot
Teams of people	Teams of robots

TABLE II
SPATIAL RELATIONSHIPS: INTIMACY AND VIEWPOINT

Role	Human's POV	Spatial Relationship
Commander	God's-eye	Remote
Peer	Bystander	Beside
Teleoperator	Robot's eye	"robo-immersion"
Developer	Homunculus	Inside

three years are: metrics, toolboxes for interfaces, establishment of principles of user-centered design, and how to incorporate the contributions from broader communities (AI, engineering, psychology, etc.)

One common complaint from more engineering-oriented participants was the lack of toolboxes and explicit examples of how to apply principles of user-centered design. Although they do not currently exist, it was clearly indicated that such toolkits would be of enormous benefit to the HRI community.

Basic research issues which should dominate the field for the next five years or longer fall into three categories: representation, cognition, and control. The need for research into representations, particularly mental models, was stressed. Representation issues include the traditional application of representation: how to define the task for both humans and robots, how to represent each agent's internal state. In the case of HRI, it is important to be able to represent the *situational context* as well as *the inter-agent trust, expectations, and/or social models*. Cognition studies should be extended to consider how agents adapt to physical constraints and produce resilience in the face of failure; these topics have not been well explored in the past. Control is also an important issue, spanning levels of control, coordination (and communication to effect that coordination), and social roles (and how they impact information exchange).

While HRI has many research facets, it is important not to overlook the relationship between humans and robots from different viewpoints. Three basic relationship taxonomies were identified: numeric, spatial, and authority relationships. The "numeric" relationship involves the ratio of humans-to-robots in the task (see Table I).

Spatial relationships can be quite subtle, with the expectations or desired working representation of one agent quite different from another as seen in Table II.

Just considering the ratio of humans to robots does not adequately capture the coordination mechanism or social informatics between the agents. Instead, it may be helpful to look at the authority relationships as shown in Table III.

C. Social Informatics

Social informatics is a critical, unexplored arena. While emotional intelligence is needed from some applications, it may be

TABLE III
AUTHORITY RELATIONSHIPS: LEVELS OF CONTROL

Authority Relationship	Function	Context required
Supervisor	Commands “what”	Tactical situation
Operator	Commands “how”	Detailed perception
Peer	Cross-cueing	Shared environment, functions
Bystander	Interacts	Shares environment

inappropriate for others; therefore, both the issues of how to embody emotional intelligence and when it is useful were suggested as technical goals for the next three years.

Social issues include the following: Who is accountable for actions? What is acceptable for a robot to do and for what type of person? What type of duties do we want to turn over to robots? An example is putting a robot in an elderly person’s apartment: What responsibilities/skills are expected from the person, and which must be given to the robot? Accountability is a big problem and limited by the different classifications of consequence of actions. In HCI, it is standard to talk of whether it is easy to undo a command, and the question was raised whether there was an analogy with HRI teams. Acceptable actions for a robot or a person can depend on many factors. Age is a common factor in human tasks, and robots and humans have different ability distinctions. Another parameter might be whether the interaction is friendly or hostile. An interesting practical question is: what are the economic implications of different social informatics? It is possible that a user might prefer a more social robot but consequently get less done?

While the study has not produced any answers to these questions, it was noted that in some cases emotions and more naturalistic human interaction modes would enhance teamwork, while in other cases these might be inappropriate. It is clear that we need to better understand when naturalistic social interaction facilitates performance and acceptance and when it does not.

D. Application Domains

It is essential to define a small number of common application domains. Research in HRI has reached the point where appropriate domains are needed for rigorous evaluation and comparison of results.

The following five cross-cutting applications were suggested which represent the space of human–robot interaction:

- search and rescue robots;
- personal assistants;
- museum docents;
- fleets of robots;
- physical therapy robots.

Specific, well-understood domains for HRI study are needed for several reasons. First, knowledge acquisition is the foundation of modeling, yet it is a bottleneck. Participants expressed concern with the need to become subject matter experts in complex application domains in addition to conducting the HRI research, and suggested the inclusion of domain practitioners in the constitution of interdisciplinary teams. Domains for HRI can be characterized in terms of the ability to capture and model relationships, the impact of interactions on performance, the frequency of interaction between agents, the richness of interaction

relationships (not simply “master–slave”), the amount of communication, clear mechanisms for evaluation of usability, and the types of end-users included.

In our first application domain, robots for urban search and rescue have the humanitarian nature of personal assistant robots and the challenge of working with “average” end-users. Search and rescue robots are interesting because of the time pressure and the requirement that they must fit into the existing organizational and information rescue hierarchy. Research is already underway in the HRI aspects of rescue robots at the University of South Florida which will aid in modeling the relationships. The frequency of interaction between humans and robots differs from personal assistants and docents: search opportunities are sporadic and short, often only three or four episodes of activity, less than ten minutes in duration each over a twelve hour shift. The interaction is brief and intense. The role of social informatics is an intriguing research question since it is uncertain as to whether users should consider robots a tool to be sacrificed or should bond with them to get enhanced performance, like dog handlers. The end-user is someone who can undergo only limited training and may have some resistance to robots in the workplace.

Personal assistant robots also offer many opportunities for exploring HRI as well as making a contribution to society. Personal assistant robots are already being developed by NASA to aid astronauts by carrying gear and holding parts for assembly and by other institutions for aiding the handicapped. Other applications include military operations, where man-packable aerial robots can give an infantryman a personal “eye in the sky,” and a carrier for search and rescue gear. Personal assistant robots are an attractive domain because humans must work side-by-side with the robots for large amounts of time. The robots’ relationships with humans are servile, but personal. The tasks are limited enough that they can be modeled and evaluated. The end users are often ordinary people who cannot be expected to become robot experts. This domain has also been explored by the Swedish research team at the Interaction and Presentation Laboratory of the Royal Institute of Technology in Stockholm, and provides an opportunity for more international collaboration.

Museum docents are quite different from personal assistant robots. Docent robots offer a one-to-many relationship with humans, rather than one-to-one, and must get humans to do things that they may not do if left to themselves (such as interact with parts of a museum). At least two museum docents are already in existence: Minerva and Sage, developed by Carnegie Mellon University and Rhino in the Deutsches Museum in Bonn, Germany. The frequency of interaction is high and the robot must contend with a wide variety of people in differing emotional states.

The issue of the ratio of robots to humans is not well explored by the first three of these domains, prompting a call for considering swarms of robots. Applications of swarms include humanitarian demining, where multiple robots work under the direction of a single (or few) humans. Fleets of robots offer a set of different interactions, mostly that of interruption to the human and cooperation with other robots. Once the robots are tasked, they should perform their job autonomously until some

anomaly occurs, then the human must be alerted. This may take the human unawares and generate an incorrect or delayed response because of not comprehending the context of the problem. While the robots have a near-peer relationship with the human, they have the possibilities of a range of relationships among themselves, depending on how the swarm is organized. They may be cooperative, have a hierarchy, etc. While fleets of robots score highly on interaction, the role of social informatics appears to be limited. The end users are expected to be highly trained.

Another proposed domain is physical therapy robots, combining many of the attributes of personal assistants with strong humanitarian contributions. Physical therapy robots, however, are expected to work in constant direct physical contact with the human and must respond to subtle social informatics signals. The authority relationship is challenging because such a robot must make sure the patient receives the care even if he/she does not want it, yet be sensitive to the patient's needs and fears.

E. Field Experience

Members of the human-robot interaction community need field experience. One of the major drawbacks in HRI is the cost of working with robots. Robots are expensive and require specialized maintenance. As a result, there are few robots capable of HRI within the robotics community, and cognitive scientists and HCI researchers often have no access. One solution is to focus on a grand challenge task hosted by multiple institutions which can maintain the robots.

V. OUTCOMES AND RECOMMENDATIONS

In addition to the five findings, the study met its overall objectives: development of a taxonomy of issues, identification of grand challenges, development of a preliminary roadmap, and establishment of a repository of resources (preliminary). Each is detailed below. However, the major outcome of the study is more basic than these objectives. The discussions and concerns repeatedly raised by participants during the BOGs and in discussions led to the following conclusion.

HRI is a cross-disciplinary area, which poses barriers to meaningful research, synthesis, and technology transfer. The vocabularies, experiences, methodologies and metrics of the communities are sufficiently different that cross-disciplinary research is unlikely to happen without sustained funding and an infrastructure to establish a new HRI community. The workshop showed that there is research in almost every area of the taxonomy, however, these advances cannot be capitalized upon because of the disparities between the communities, i.e., the left hand does not know what the right hand is doing. It was a clear sentiment among the participants that HRI simply will not happen without an infrastructure.

A. Taxonomy of Issues

The first objective of this study was the development of a taxonomy of issues, including identification of various dimensions

of the HRI field. These are, communication, modeling, teamwork, usability, reliability, and safety evaluation standards and metrics, application domains, and representative end-users.

1) *Communication Issues:* Communication has many facets. *Direct human-robot communication* is possibly the most obvious issue. Modalities include, speech, vision, gesture, and teleoperation, though there may be other forms. *Mediated human-robot communication* is another topic. This arises from virtual environments, graphical user interfaces, and can be enacted by collaborative software agents. The physical *interaction and interfaces* impact communication. These methods include physical interaction between robots and humans, mixed-initiative interactions between humans and robots, and dialog-based interaction.

There are many aspects of interaction and interfaces which need to be explored. Inferring intent of an agent was noted as being critical. Technology transfer is needed to improve the state of robot interfaces, especially adding speech recognition. Studies are needed to determine what types of interfaces make interaction most efficient and most tolerant to high workloads. In terms of visuo-motor control, there is a need for basic research on how humans interact with machines, particularly with visual stimuli. Effects of delays, poor synthesis of information, and dynamic interactions are also important to HRI. A related question is: To what extent can people adapt to increased visuo-motor delays? The type of interaction will obviously be influenced by how many different kinds of robots we expect people to interact with. If the number of types is typically "one," then the user can invest a lot of effort into learning how to cope with that one kind. If the answer, as in the graphical UI world, is "many," then some common language or form of communication will be necessary to simplify the learning of each robot's control.

Other communication questions considered the role of other research trends to HRI interfaces: What about robots which are modeled on animal behaviors or which are new generation animal/machine hybrids—what kind(s) of interface(s) will allow humans to direct/control/interact with these types of robots or robotic communities? Can we map current interface techniques (e.g., speech, vision, gesture, augmented/virtual reality, direct manipulation GUIs, etc.) to the various types of relationships, or do we need to develop entirely new kinds of interfaces? How is believability (of the type used in animation) applied to robotics?

2) *Modeling:* Modeling issues spanned traditional concerns (cognitive, task, and environment modeling) to more HRI-specific concerns. Cognitive modeling of human reasoning, behavior, intention and action is needed for imitation (i.e., the robot learns how to behave from the human) and for collaboration (i.e., the robot understands what the human is doing within the context of the task). Task and environment modeling are needed as a basis for performance. Other modeling issues include social relations, learning, and methods.

An interesting modeling issue is that of *social relations*. One aspect is whether (and when) social relationships are necessary. For example, can "no personality" in an intelligent agent (software or robot) be perceived by humans as a cold, insensitive, indifferent agent? If so, do these perceptions differ by specific groups of people, differentiated by age, gender,

culture, etc.? On the flip side, does the personality of the human affect how the human interacts with the robot? If so, how? Does it arouse specific emotions, behaviors? Which ones, in what contexts? Are these effects (above) positive, or negative? Can we improve on these toward the positive? How so? In order to understand when these relationships are needed or when perceptions of such relationships need to be changed, social relations must be modeled.

Another aspect is the social relationships themselves, what they are and how they relate to situations and context. It is important to establish a taxonomy of human–robot “relationships,” and identify what levels of “interpersonal skills” the robots will need to perform effectively in these roles. Likewise, the impact of the context of the relationship (e.g., workplace versus home, safety-critical versus low-impact, remote versus local, hazardous versus benign, dependent versus independent) must be studied. Another cognitively-oriented question is whether robot personalities and affective states *have* to correspond to human personalities and affective states in order to be useful. And if not, what are the pros and cons of having robot personalities mimic human qualities? Regardless, the effective study of these relationships will require models.

While *learning* is not usually thought of as a modeling issue, per se, it does have a role in how to model human–robot relationships and components. Learning is needed to improve performance, exchange skills, and to adapt models of interaction. A basic question is: What is the role of learning in human robot interactions? A direct offshoot, given the evidence of rarity of true imitation in nature (excluding humans), is what tools will best facilitate learning between humans and robots? Answers to questions such as the following also require learning models: Can we design simple human–robot communication/interaction mechanisms that would help robots generalize from multiple learning experiences (e.g. teacher-provided feedback through speech, gestures or other)? Can we increase the expressiveness of teacher-robot demonstration experiences (and therefore the complexity of the tasks to be taught to a robot) through tighter human–robot interaction methods (e.g. speech, meaningful symbols, etc.)? This, in turn, raises the questions of: How will humans monitor the learning state of the robots? How will humans learn with learning robots?

Methods for modeling vary. User-centered software engineering modeling, methodologies and techniques have already been established. Techniques for collection of user-centered data (e.g., interviews, surveys, talk-aloud reports, video protocols, etc.) and analysis of user-centered data (e.g., protocol analysis, task analysis, etc.) exist, although they are not commonplace in the robotics community. A wide range of rapid prototyping techniques may be useful: storyboard mockups, wizard-of-oz techniques, and simulations. An unanswered question is whether new techniques are needed especially for robotics. It is believed that it is essential to have an iterative involvement of users throughout the development process, not just at the end (user-centered software engineering techniques), and that well-known usability guidelines should be followed throughout the design and development process. The evaluation of usability should measure the adherence to established

guidelines, effectiveness of the human-robot communication, and effectiveness of human–robot performance.

3) *Teamwork*: Teamwork issues can be subdivided into two areas: architectures and task allocation. Teamwork is particularly relevant for control of unmanned platforms for the military. Previous work in mixed-initiative systems may provide some insight into the dimensions of task allocation. Architectures focus on the optimal organization of teams of multiple robots and a single human, multiple humans and a single robot, and multiple robots-multiple humans. Research into architectures is expected to determine situations which require an authoritarian, hierarchical structure, or a more “democratic” structure. While architectures are being investigated by the multiagent community, they often neglect questions of how a single robot can work with more than one human, balancing multiple demands, and how tasks can be traded between humans and robots as needed.

Task allocation in human-robot teams is nontrivial because each partner in the team has skills that the other lacks, including intelligence skills. One example was given of determining the correct partitioning of skills. In surgery robotics, it may be easy to determine that a human’s hand shakes and that visual acuity is limited, but it is much more difficult to detect deficiencies in spatial reasoning. It is not clear what we need to know about environments, tasks, humans, and robots to be able to optimize mission performance even if we knew the capabilities of the human and robot.

Task allocation is unlikely to be static. For example, human workload may affect the human–robot interaction. Some robot designs may even work better when humans have a high workload. On the other hand, it may be possible to “over-task” the robot. Likewise, people may be smarter than machines in some situations. Given a dynamic situation, who has the ultimate control: the robot or the human?

4) *Usability, Reliability, and Safety Evaluation Standards and Metrics*: A major issue in HRI is the overall utility of such systems. HCI has typically considered this in terms of usability, while roboticists consider the reliability and robustness of the system to be critical. As a result, open issues include the determination of appropriate metrics of evaluating the success, effectiveness, and quality of human–robot teams and establishing whether such metrics can be task-independent. The need for metrics also emphasizes the need for benchmark problems where work in the HRI arena can be directly compared and where the effectiveness of different human–robot interfaces can be measured. Usability studies are also warranted. Task analysis of users is needed as well as measures of the utility of different types of human–robot relationships.

Fault tolerance and failure management must be considered from several angles. One viewpoint is which agent is responsible for the failure and how does this impact the kinds of interfaces needed to diagnose and recover from the problem. Another consideration is the impact on each agent of an interruption in order to deal with a failure.

5) *Application Domains*: The participants in the study found that there is an expectation that human–robot interaction applications will soon become ubiquitous. However, the applications themselves create an important research issue: the choice of good applications will be very helpful in identifying

further issues in HRI and confirming models. The applications cited were: urban search and rescue (USAR), military applications—digital battlefield and/or robotic forces, personal care and service robots, home appliances (lawnmowers, vacuum cleaners, etc.), medical applications (robotic surgery, hospital delivery systems, etc.), entertainment robots (toys, pets, parks guides, etc.), driving robots, humanoid robots, space exploration, and hazardous environments collaborations.

6) *Representative End-Users*: End-users span the spectrum of human ability. At the highest level of specialization and robot familiarity are researchers, graduate students, and specialized robotics technicians (e.g., JPL and NASA personnel). The next level is skilled workers in other areas (little or no robotics or even computer experience required), e.g., search and rescue workers, miners, manufacturing workers, etc. These workers are expected to have a collaborative relationship with the robot. The next lower level of robot familiarity is unskilled workers, who may encounter the robots as part of the workplace even though they are not working directly with the robots. When they encounter the robot, they have to establish their relationship with it. Another level of end-user is disabled or elderly people, which have little or no robot or computer experience required. In this case robots are personal service robots, and have a prosthetic reliance relationship, i.e., people rely on robots to improve their quality of life and help them do basic tasks they wouldn't normally be able to do. The next level is ordinary people with little or no robot or computer experience required who use personal service robots. Here the robots act as "staff" like maids, butlers or robotic appliances like vacuum cleaners, lawn mowers, etc. The lowest level of ability is children who have little or no robot or computer experience but operate robot pets, toys, entertainment 'bots, e.g., in an entertainment park.

B. Grand Challenges

The study generated two types of grand challenges: applications and isolated technologies. Two well-defined grand challenge applications were the following. 1) Develop a robot search and rescue system that is the system of choice for the majority of search and rescue units by 2010, and 2) meet the existing AAI grand challenge of a robot attending conference and delivering presentations.

Additional applications challenges were less concrete. 1) Create a robot for service positions. Service positions were expected to have interaction with ordinary people as well as have natural metrics of performance. 2) Develop a robot for military operations in urban terrain (MOUT) and reconnaissance, surveillance, targeting, and acquisition (RSTA) applications or for service positions.

The technology challenges were varied.

- 1) Picking up human social cues (attentional state, body language) and interpreting human behavior (intent, goals, desires) would be impressive demonstrations.
- 2) Along those lines, showing the elements of team cohesion by understanding the operator state and the environment would be worthy.

- 3) Likewise a study of group dynamics versus team dynamics might be useful, for example whether a mixed human/robot soccer team could beat pure robot or human teams and the results analyzed.
- 4) A demonstration of a robot receiving instruction from a human would also be notable.
- 5) Advances in physical interfaces such as wearable computing (gestures, voice) and instrumented garments were also mentioned, as was development of a standard tasking language.

Other participants felt that any demonstration of communication, representation, and cognition would be a worthy challenge.

It was suggested that possibly an essay contest could be held to encourage researchers to articulate and flesh out these Grand Challenges.

C. Preliminary Roadmap

The workshop participants avoided the idea of a traditional roadmap which outlines an orderly achievement of milestones in the apparent belief that the research area is too new. Instead, the participants focused on near-term actions that need to be taken into order to cement the HRI community and to accomplish sufficient basic cross-disciplinary research to formulate a roadmap at a later date. A partial roadmap was constructed from the comments.

1) *Immediate: Establishment of HRI Community Infrastructure*: The unanimous first milestone for the study participants was the creation of an HRI community. It was felt that the differences between the robotics, cognitive science, and HCI communities were significant and cross-disciplinary work would not evolve without an infrastructure. (Additional communities in the areas of AI and systems engineering also need to be included.) In particular, researchers need:

Cross-education, such as tutorials at major conferences and at follow up workshops. This is needed to reconcile the vocabularies of the different disciplines, familiarize groups with each other methods, and to forge a consensus.

Established benchmark domains to facilitate entry and communication and evaluation. The need for benchmark domains is echoed throughout the findings. Unfortunately, it can take years to adequately identify and model such a domain. This prohibits researchers with domain expertise from applying their research. Instead, studies of each of the domains could be done and made available so that the community can see issues and opportunities without a high entry cost. This is expected to bootstrap the formation of the community.

Centralized infrastructure, to serve as a repository for the field and to continue to foster the development of the HRI community.

2) *Next Three Years: Metrics, Toolkits, Principles*: While researchers in HRI need to formalize themselves as a sustainable community, they also need to work on generating the basic accessories of a science of HRI: metrics, toolkits, and principles. In particular, it would be desirable for the HRI community to identify "standard" components such as software systems for speech and gesture recognition and hardware systems. Further-

more, there should be a considerable effort devoted to modularization of components so that researchers can “plug-and-play” components rather than being forced to reinvent the wheel. One goal is to create standards so that people can interact with robots like they do with a car or at least establish general protocols for interacting with robots.

3) *Next Five Years: Representation, Cognition, and Control*: Representation, cognition, and control are assumed to be the basic research issues for the HRI community. However, the HRI field is still too new to set milestones or benchmarks within these areas.

D. Recommendations

From the follow up surveys and end discussions at the workshop, it appears that there are two actions that can be taken to bootstrap the coalescence of the robotics, cognitive science, HCI, AI and engineering researchers into an HRI community. These two are:

1) *Sponsor the Development of an HRI Community Infrastructure Through Online Websites and Supporting Workshops/Tutorials for the Next Three Years at Major Discipline-Related Conferences*: The online website should at the very least emulate NASA’s PostDoc system where groups place information, papers, etc. However, without funding, websites often become just a snapshot rather than a continuous effort. A funded website could provide links to papers, plug and play module exchanges, serve as a “dating service” for collaboration, and post sources of funding. It could eventually be expanded to include an online journal.

Many participants called for a series of workshops and tutorials. Tutorials should be given at major discipline related conferences to introduce that community to the parlance of the other communities; the tutorials would serve to supply researchers with the relevant fundamentals from the other fields. In addition there should be at least one workshop aimed at bringing together the entire community.

It was noted that the tutorials and workshops were unlikely to happen or happen in a way to provide continuity without some person or group acting as a manager. The coordination effort spans too many communities and requires too much time to be handled through volunteerism.

2) *Fund the Characterization of One or More Benchmark Applications and Researcher Access to the Application Domain*: A common theme throughout the workshop was the need for well-understood benchmark domains and for access to real users and robots in these domains. The study has identified several possible domains; researchers in these domains should be funded to complete a useful characterization and make it available to the community.

One warmly welcomed suggestion was to have summer camps where researchers would be introduced to the users and the domain. Other suggestions included having faculty summer or exchange programs whereby researchers could visit and work at robotics labs. The participants wanted to see mechanisms for multi-year involvement of researchers, not just a one-time introduction to a domain or robotic system. These efforts require significant funding to cover travel costs and the costs associated with the subject matter experts.