Theory and Evaluation of Human Robot Interactions

Jean Scholtz

National Institute of Standards and Technology

100 Bureau Drive, MS 8940

Gaithersburg, MD 20817

Jean.scholtz@nist.gov

ABSTRACT

Human-robot interaction (HRI) for mobile robots is still in its infancy. Most user interactions with robots have been limited to tele-operation capabilities where the most common interface provided to the user has been the video feed from the robotic platform and some way of directing the path of the robot. For mobile robots with semi-autonomous capabilities, the user is also provided with a means of setting way points. More importantly, most HRI capabilities have been developed by robotics experts for use by robotics experts. As robots increase in capabilities and are able to perform more tasks in an autonomous manner we need to think about the interactions that humans will have with robots and what software architecture and user interface designs can accommodate the human in-the-loop. We also need to design systems that can be used by domain experts but not robotics experts. This paper outlines a theory of human-robot interaction and proposes the interactions and information needed by both humans and robots for the different levels of interaction, including an evaluation methodology based on situational awareness.

1. Introduction

The goal in synergistic cyber forces is to create teams of humans and robots that are efficient and effective and take advantage of the skills of each team member. An important subgoal is to increase the number of robotic platforms that can be handled by individuals. In order to accomplish this goal we need to examine the types of interactions that will be needed between humans and robots, the information that humans and robots need to have desirable interchanges, and to develop the software architectures and interaction architectures to accommodate these needs.

Human-robot interaction is fundamentally different from typical human-computer interaction in several dimensions. [11] note that HRI differs from HCI and Human-machine Interaction (HMI) because it concerns systems which have complex, dynamic

control systems, exhibit autonomy and cognition, and which operate in changing, real-world environments. In addition differences occur in the types of interactions (interaction roles); the physical nature of robots; the number of systems a user may be called to interaction with simultaneously; and the environment in which the interactions occur. Each of these differences is discussed in the ensuing paragraphs.

[20] defines three roles: supervisor, operator, and peer. To expand on these roles slightly I have added a mechanic role and divided the peer role into a bystander and teammate role. Supervisory and teammate roles imply the same relationships between humans and robots as they do when applied to human-human interactions. An operator is needed to work “inside” the robot; adjusting various parameters in the robot’s control mechanism to modify abnormal behavior; to change a given behavior to a more appropriate one; or to take over and tele-operate the robot. The mechanic type of interaction is undertaken when a human needs to adjust physical components of the robot, such as adjusting the camera or adjusting various mechanisms. A bystander does not explicitly interact with a robot but needs some model of robot behavior to understand the consequences of the robot’s actions. For example, will the floor cleaning robot in the workplace sense the presence of a person and stop or must the person move from the robot’s path. Each of these interactions has different tasks and hence, different situational awareness needs.

The second dimension is the physical nature of mobile robots. Robots need some awareness of the physical world in which they move. Robots that can physically move from one location to another as opposed to robot platforms that stay in one location but have mobile components present more interesting challenges. Ground robots encounter more obstacles than unmanned systems in the air and under water.

Therefore we consider the more complicated case of mobile ground robots for the purposes of developing our framework. As robots move about in the real world, they build up a “world model” [1]. The model

the robot platform builds up needs to be conveyed to the human in order to understand decisions made by the robot as the model may not correspond exactly to reality due to the limitations of the robot's sensors and processing algorithms.

A third dimension is the dynamic nature of the robot platform. Typical human-computer interactions assume that computer behavior is for the most part deterministic and that the physical state of the computer does not change in such a way that the human must track. However, robotic platforms have physical sensors that may fail or degrade. While some functionality may be affected, the platform may be able to carryout some limited tasks.

The fourth dimension is the environment in which interactions occur. Platforms to monitor robots may have to function in harsh conditions such as dust, noisy and low-light conditions. Environments may be dynamic as well. Search and rescue robots may encounter more building or tunnel collapses during the operation. In a military environment, explosions may drastically change the environment during the mission. Not only will the robot have to function in these conditions but the user interacting with the robot may be co-located as well (a team member, perhaps). Thus interactions may have to be carried out in noisy, stressful, and confusing conditions.

The fifth dimension is the number of independent systems the user needs to interact with. Typical human-computer interaction assumes one user interacting with one system. Even in collaborative systems we usually consider one user to one system with the added property that this user-computer system is connected to at least one other such system. This allows interaction between users, moderated by the computers, as well as computer – computer interaction. In the case of humans and robots, our ultimate goal is to have a person (at least for a number of the interaction roles we're specified) interacting with a number of heterogeneous robots.

The final dimension is the ability of the robot to perform autonomously for periods of time. While typical desktop computers perform autonomously in that they execute code based on user commands, robots use planning software to alleviate the user from dealing with low level commands and decisions. Thus a robot can go from point A to point B without asking the operator how to deal with each obstacle encountered along the path.

2. A background: human-robot interaction

Human-robot interaction was first associated with teleoperation of factory robotic platforms. Sheridan [21] defines telerobotics as : “direct and continuous

human control of the teleoperator” or “machine that extends a person's sensing and/or manipulating capability to a location remote from that person.” He distinguishes telerobotics or supervisory control of a remote machine from supervisory control of any semi-autonomous system regardless of the distance. Human –computer interaction in Sheridan's view includes telerobotics. Human-computer interaction is the term most commonly used to denote that a computer application and its associated files are the objects being manipulated, not a physical system controlled through the computer.

Human –robot interaction (HRI) goes beyond teleoperation of a remote platform and allows for some set of autonomous behaviors to be carried out by the robot. This could range from a robot responding to extremely precise commands from a human about adjustment of a control arm to a more sophisticated robot system planning and executing a path from a start point to an end point supplied by a user. The concept of human-robot interaction has only become possible in the last decade because of advances in the field of robotics (perception, reasoning, programming) that make semi-autonomous systems feasible. An NSF/DOE IEEE workshop [19, 2] identified issues for human-machine interfaces and intelligent machine assistants. These issues included:

- Efficient ways for a human controller to interact with multiple-semi autonomous machines
- Interfaces and interactions that adapt depending on the functions being performed

Kidd [16] noted that human skill is always required in robotic systems. Kidd maintains that designers should use robot technology to support and enhance skills of the human as opposed to substituting skills of the robots for skills of the human. He argued for developing and using robotic technology such that human skills and abilities become more productive and effective, such as freeing humans from routine or dangerous tasks. He points out that robotic researchers tend to focus on issues that are governed by legislative requirements such as safety. Human-centered design issues have been mostly ignored. Kidd suggests that human-centered design of human-robot interaction needs to look beyond technology issues and to consider issues such as task allocations between people and robots; safety; group structure. These issues need to be considered in the early stages of the technology designs. If they are only considered in the final stages, the issues become secondary and have little impact on design considerations.

Fong, Thorpe, and Bauer [11] note that it is clear that benefits are to be gained if robots and humans work together as partners. But partners must engage in

dialogue, ask questions of each other, and jointly solve problems. They propose a system for collaborative control which provides the best aspects of supervisory control without requiring user intervention within a critical window of time. In collaborative control, the human gives advice but the robot can decide how to use human advice. This is not to say that the robot has the final authority but rather the robot follows a higher level strategy set by the human with some freedom in execution. If the user is able to provide relevant advice, the robot can act on that. However, if the user is not available within the time needed, the robot will use default behaviors to react to the situation. Collaborative control is only possible if the robot is self-aware; has self-reliance and can maintain its own safety; has a dialogue capacity; and is adaptive. Dialogue management and user models are needed to implement collaborative control systems.

Hill [14] notes that it is important that research in HRI include human factors practitioners in multidisciplinary teams. It should also be stressed that HRI includes much more than just a clever interface for the user. To truly develop synergistic teams it is necessary to consider the skills of both humans and robots and to develop the overall system that allows all parties to fully utilize their skills. This is even more challenging given the dynamic nature of robotic platforms today. We need to design HRI in such a way that it is useful today but fully capable of evolving as the capabilities of robots evolve.

Robotics researchers use the term human-robot intervention, often in place of human-robot interaction. For robotic systems that have plan-based capabilities, the term intervention is used when a human needs to modify a plan that has some deficiency or when the robot is currently unable to execute some aspect of a plan. While robots carrying out preplanned behaviors is certainly a desired activity (e.g., clean the kitchen floor, watch the perimeter, check all the rooms on the 3rd floor for X), more closely coupled human-robot teams need to interact spontaneously as well. In this paper I use the term “human-robot interaction” to refer to the overall research area of teams of humans and robots, including intervention on the part of the human or the robot. I use “intervention” to classify instances when the expected actions of the robot are not appropriate given the current situation and the user either revamps a plan ; gives guidance about executing the plan; or gives more specific commands to the robot to modify behavior.

3. Human-computer Interaction

In the introduction I listed six dimensions in which HRI is fundamentally different from traditional human-

computer interactions. A first step in developing a framework for HRI is to determine what, if anything, is applicable from work done in previous HCI research. One model of human-computer interaction is Norman’s seven stages of interaction [18] Norman considers these seven stages:

1. Formulation of the goal – think in high level terms of what it is you want to accomplish.
2. Formulation of the intention – think more specifically about what will satisfy this goal.
3. Specification of the action – determine what actions are necessary to carry out the intention. These actions will then be carried out one at a time.
4. Execution of the action – physically doing the action. In computer terms this would be selecting the commands needed to carryout a specific action.
5. Perception of the system state – the user must then assess what has occurred based on the action specified and execution. In the perception part the user must notice what has happened.
6. Interpretation of the system state – having perceived the system state, the user must now use her knowledge of the system to interpret what has happened.
7. Evaluation of the outcome – the user now compares the system state (as perceived and interpreted by her) to the intention and to decide if progress is being made and what action will be needed next.

These seven stages are iterated until the intention and goal are achieved – or the user decides that the intention or goal has to be modified.

Figure 1: Norman’s HCI Model

Norman defines two issues with these seven stages: the gulf of execution and the gulf of evaluation. The gulf of execution is a mismatch between the user’s intentions and the allowable actions in the system. The gulf of evaluation is a mismatch between the system’s representation and the user’s expectations. These correspond to four critical points where failures can occur. Users can form an inadequate goal or may not

know how to specify a particular action or may not be able to locate an interaction object. These result in a gulf of execution. Inappropriate or misleading feedback from the system may lead the user to an incorrect interpretation of the system state resulting in a gulf of evaluation

4. A theory of human-robot interaction

4.1. Level of interaction scenarios

Some assumptions are necessary first. For our theory of human-robot interaction we are concerned with semiautonomous mobile robots interacting alone and in teams. Sheridan [21] outlines five generic supervisory functions: planning what task to do and how to do it; teaching or programming the computer; monitoring the automatic action to detect failures; intervening to specify a new goal in the event of trouble or to take over control once the desired goal state has been reached; and learning from experience. In our theory we are concerned with support for specifying the actions; monitoring the actions; and intervention. We make the assumption that the robot is already programmed to carry out basic functions and any “reprogramming” happens during intervention. For the initial version of our theory we are not considering learning on the part of the robot or on the part of the user. The following scenario illustrates the HRI roles.

An elder care facility has deployed a number of robots to help in watching and caring for its residents. The supervisor oversees the robots which are distributed throughout the facility and makes sure that the robots are properly functioning and that residents are either being watched or are being cared for – either by a robot or by a human caregiver. A number of human caregivers are experts in robot operation and assist as needed depending on their duties at the time. The operators might use a mobile device, such as a PDA to adjust parameters in the robot software. The elder care facility also employs a mechanic who is called when needed to adjust the physical capabilities of the robot – such as cameras becoming dislodged. The caregiver robots can perform routine tasks such as helping with feeding, handing out supplies to residents, and assisting residents to move between locations in the facility. Watcher robots monitor residents and have the capability to send back continual video feeds but also alert the supervisor or a nearby human caregiver to an emergency situation. In most cases, the human and robot caregivers work as teams. Human caregivers can override preplanned behaviors to ask robots to assist with more critical situations, such as moving residents

to another part of the room if an emergency situation occurred – such as a resident falling. Robots interact with the residents as well as visitors to the facility who may not be aware of their capabilities.

What can we learn from this scenario? First of all, the boundary between the levels of interactions is fuzzy. The supervisor can take the operator role assuming the supervisor has the necessary cycles to do so. This might be more efficient than notifying the designated operator and handing off the problem. The team members can command the robots within the intent of the supervisor. Bystanders who have little or no idea of the capabilities of the robot and who do not have access to computer displays of robot status will have some level of interaction with the robots. All of the different interaction roles can occur at the same time. The same person might assume more than one role at a time or different people could have different interaction roles.

4.2. Models of HRI

What changes to this model of HCI are necessary to describe HRI systems? The following sections contain possible models of interaction for the various HRI roles.

4.2.1. Supervisor Interaction. A supervisor role could be characterized as monitoring and controlling the overall situation. This could mean that a number of robots would be monitored and the supervisor would be evaluating the given situation with respect to a goal that needs to be carried out. For robots that possess planning systems, the goals and intentions have been given to the planning system, and the robot software is generating the actions based on a perception of the real world. The supervisor can step in and specify an action or modify plans. In either case, a formal representation of the goal and intention is necessary so that the supervisor can formulate the effect an intervention will have on the longer term plan.

Figure 2 contains a proposed model for the supervisor- robot interaction. The main loop is the perception/evaluation loop as most of the actions are automatically generated by the robot software. Supervisor interactions at the action and intention level must be supported as well. Note that for multiple robotic systems the supervisor must monitor the status of all platforms. Figure 2 shows that the human-robot interaction for the supervisor is heavily perceptually based, and that interactions need to be supported on both the action and intention level.

4.2.2. Operator Interaction. The operator is called upon to modify internal software or models when the

robot behavior is not acceptable. The operator will deal mainly with interacting at an action level –actions

Figure 2: HRI Model- Supervisor Role

allowed to the operator. It will be necessary to then determine if these actions are being carried out correctly and if the actions are in accordance with the longer term goal. The assumption is that the supervisor role is where the intentions or longer term plan can be formally changed – not at the operator level.

Figure 3: HRI Model - Operator Role

4.2.3. Mechanic Interaction: The mechanic deals with physical interventions, but it is still necessary for the mechanic to determine if the interaction has the desired effect on the behavior. So, the model looks similar to the model for the operator interaction. However, the difference is that while the modifications have been made to the hardware, the behavior testing needs to be initiated in software and observations of both software and hardware behavior are necessary to ensure that the behavior is now correct.

4.2.4. Peer Interaction. Teammates of the robots can give them commands within the larger goal/ intentions, though we follow the same assumption here – that only the supervisor role has the authority to change the larger goal/ intentions. This assumption is based on the time that is needed to alter goals and plans. Even with

good user interfaces, teammates may not have the necessary time to perform these interactions. If they do, they can certainly switch to the supervisory role if appropriate.

Figure 4: HRI Model – Mechanic Role

Figure 5: HRI Model – Peer Role

The model in Figure 5 shows the interaction model proposed for peer interactions. We propose that this interaction needs to occur at a higher level of behavior than the operator interactions allow. Human team members talk to each other in terms of higher level intentions – not in terms of lower level behaviors. Terms such as follow me, make a sharp left turn, wait until I get there would be reasonable units of dialogue between a robot and a human team member in the peer role. In this case, direct observation is probably the perceptual input used for evaluation. In the case that the behavior is not correctly carried out, the peer has the choice of switching to the operator model or handing off the problem to someone more qualified as the operator.

4.2.5. Bystander Role. The final role is that of the bystander. Earlier we posed an interesting question: should a bystander be given a subset of interactions with the robot appropriate to this role? For the

purposes of this model, let's assume that this is true. The bystander might be able to cause the robot to stop by walking in front of the robot, for example.

Figure 6: HRI Model – Bystander Role

In this model, the bystander has only a subset (sub A) of the actions available. She is not able to interact at the goal or intention level. Feedback must be directly observable. The largest challenge here is how to advise the bystander of the capabilities of the robot that are under her control. There will most likely not be a typical display. The research on emotion and social interaction with robots is applicable here [4,5].

4. Situational awareness

Given the HRI models proposed in section 4, one question becomes how to evaluate human-robot interactions. In all the models the perceptual step is quite necessary. And in many of the roles, it is necessary to understand not just the state of the robot system after the action has occurred, it is also critical to understand what the robot state was when the action was given. This helps us understand possible mismatches in behaviors specified versus behaviors actually carried out.

A second issue is the separation of the performance of the HRI system from the performance of the user interaction design and the actual interface. Due to the physical nature of the robots as well as the sophisticated software incorporating perception, learning, and planning, a failure in performance may not be due to an issue with the user interaction but may be attributed to the robot's software system or malfunctions of the robot's sensory system.

Therefore we plan to carryout our HRI evaluations in two stages. We will evaluate the perceptual part of the model separately from the intervention part of the interaction design and we will separate both of those

from the actual performance of the HRI system. The evaluation of the intervention portion will not be discussed in this paper as it will be based on current usability evaluation methodologies. The evaluation of the perceptual part of the model will be based on assessing situational awareness.

However, each of the levels of interaction will require a different perspective and hence different situational awareness. These issues will be discussed in the sections detailing the proposed HRI roles. As background it is necessary to have an understanding of situational awareness, as well as methodologies and measurement tools to assess situational awareness.

Situational awareness [9] is the knowledge of what is going on around you. The implication in this definition is that you understand what information is important to attend to in order to acquire situational awareness. Consider your drive home in the evening. As your drive down the freeway and urban streets there is much information you could attend to. You most likely do not notice if someone has painted their house a new color but you definitely notice if a car parked in front of that house starts to pull out in your path. There are three levels of situational awareness [8] which correspond to various stages of evaluation in Norman's model of HCI.

Level One of situational awareness is basic - the perception of cues. You have to perceive important information in order to be able to proceed. Failures to perceive information can result as shortcomings of a system or they can be due to a user's cognitive failures. In studies of situational awareness in pilots, 76% of SA [15] errors were traced to problems in perception of needed information.

Level Two of situation awareness is the ability to comprehend or to integrate multiple pieces of information and determine the relevance to the goals the user wants to achieve. This corresponds to interpretation and a portion of evaluation in Norman's seven stages.

A person achieves the third level of situational awareness if she is able to forecast future situation events and dynamics based on her perception and comprehension of the present situation. This corresponds to the evaluation and iterative formulation and specification stages of Norman's theory.

Performance and situational awareness, while related, are not directly correlated. It is entirely possible for a person to have achieved level three situational awareness but not perform well. This is evident in Norman's stages of action – other reasons for not achieving the correct execution are certainly possible. Some of these reasons can be attributed to poorly designed systems while others can be attributed to a user's cognitive failures. Direct system

measurements of performance of selected scenarios in context is one way to measure situational awareness but only if it can be shown that performance depends only on situational assessment. One method of direct system measurements to overcome this is to introduce some sort of disruption into the system, such as a completely unrealistic pattern, and measure the amount of time that it takes users to detect the anomaly.

The most common way to measure situational awareness is by direct experimentation using queries [8]. The task is frozen, questions are asked to determine the user's situational assessment at the time, then the task is resumed. The Situation Awareness Global Assessment Technique (SAGAT) tool was developed as a measurement instrument for this methodology [6]. The SAGAT tool uses a goal-directed task analysis to construct a list of the situational awareness requirements for an entire domain or for particular goals and subgoals. Then it is necessary to construct the query in such a way that the operator's response is minimized. For example, if a user were being queried about the status of a particular robot, the query might present the robot by location rather than replying on the user to recall a name or to understand a description. The various options for status could be presented as choices rather than relying on the user to formulate a response that might not include all the variables desired.

5.1. Issues with situational awareness

There are individual differences in situational awareness. Experiments by Gugerty and Tirre [13] show that situational awareness is correlated with working memory; perceptual-motor ability; static visual processing; dynamic visual processing, and temporal processing ability. In addition, studies have shown that the ability to acquire situational awareness decreases with age [3]. These are factors that must be accounted for when doing assessments of situational awareness with respect to interface designs in the human-robot interaction domain.

Operators of fully automated systems often have difficulty in responding to emergency situations. The SAGAT tool has been used to show that there is a decrease of situational awareness with fully automated systems [7]. Goodrich, Olsen, Crandall, and Palmer [12] introduce the concept of neglect to capture the relationship between user attention and robot autonomy. The idea is that a robot's effectiveness decreases as the operator fails to attend to that robot. Neglect can be caused by time delays in remote operations or by increased workload on the part of the operator. As robots become more autonomous the

breadth of the tasks they can handle decreases. This makes them less effective but more tolerant of neglect.

5.2 Situational awareness requirements for HRI roles

As noted early the different roles within HRI require different awareness of the situation. In the following sections we propose some information we hypothesize is appropriate to the various roles. We propose to use several sources for guidance. We will attempt to find a corresponding domain and use successful interaction designs as a first basis. Secondly, we will use subject matter experts (as available) for each role to verify this information. In some instances (particularly the peer and bystander roles) we will have to conduct some experiments to gather the necessary information. Based on this knowledge, we will construct situation awareness assessment tools and user interfaces. Using the situation awareness assessment tool we will produce a baseline metric for a number of situations. HRI researchers will be able to use our user interface and assessment tool to assess their work.

5.2.1. The Supervisory Role. We assume that the supervisory interface is done from a remote location. Our hypothesis is that the supervisor needs the following information:

- an overview of the situation, including progress of multiple platforms
- the mission or task plan
- current behaviors of any robots including deviations that may require intervention
- other roles interacting with the robot(s) under her control, including interactions between robots

A corresponding HCI domain is that of complex monitoring devices [22]. Complex monitoring devices were originally based on displays of physical devices. The original devices were just lights and switches that corresponded to a sensor or actuator. Initially these were displayed on physical panels. When these displays were switched to computer-based displays, a single display was unable to show all the information. This produced a keyhole effect – the notion that a problem was most likely occurring on a display that wasn't currently being viewed.

Another issue in complex monitoring devices is that of having an indication of what "normal" is. This is also true in human-robot interactions where physical capabilities of the system change and the supervisor needs to know the "normal" status of the robot at any given time. Another issue is that single devices may not be the problem but rather relationships between existing devices. Displays should support not only problem driven monitoring but knowledge driven

monitoring when the supervisor actively seeks out information based on the current situation or task.

Due to the amount of information present in complex monitoring devices, users have strategies to reduce cognitive demands. These include reducing noise by turning off meaningless alarms, documenting baseline conditions, creating external cues and reminders for monitoring various components. Computer based displays of complex systems present more flexibility for users to view information in different forms. But there is a tradeoff between the time to manipulate the interface and any performance increase because of this increased flexibility.

We suggest that lessons learned in producing displays for monitoring complex systems can be used as a starting point for supervisory interfaces for HRI. In addition, basic HRI research issues that are not addressed in complex systems include:

- what information is needed to give an overview of teams of robots
- can a robot team world model be created and would it be useful
- what (if any) views of an individual robot world models are useful
- how to give an awareness of other interaction roles occurring
- handoff strategies for assigning interventions to others

Situational awareness indicators will be developed based on a task- analysis of the supervisor's role in a number of scenarios (such as those described earlier in this paper). An initial hypothesis about possible indicators of situational awareness includes:

- which robots have other interactions going on
- which robots are operating in a reduced capability
- the type of task and behaviors that the robots are currently carrying out
- the current status of the mission

5.2.2 . Operator interaction. We make the assumption that this will be either a remote interaction or will occur in an environment in which any addition cognitive demands placed on the user are by the environment are light. We will also assume that the operator has an external device to use as an interface to the robot. The operator must be a skilled user, having knowledge of the robotic architecture and robotic programming. If the robot has teleoperation capabilities the operator could take over control. This is the most conventional role for HRI. Moreover, as the capabilities and roles of robots expand, this role has to be capable of supporting interaction in a more complex situation.

We hypothesize that the operator needs the following information:

- The robot's world model
- The robot's plans
- The current status of any robotic sensors
- Other interactions currently occurring
- Any other jobs that are currently vying for the operator's attention (assuming it is possible to service more than one robot)
- The effects of any adjustments on plans and other interactions
- Mission overview and any timing constraints

Murphy and Rogers [17] note three drawbacks to telesystems in general:

- The need for a high communication bandwidth for operator perception and intervention
- Cognitive fatigue due to repetitive nature of tasks
- Too much data and too many simultaneous activities to monitor.

Murphy and Rogers propose the mode of teleassistance which consists of a basic cooperative assistance architecture, joining sensor fusion effects to support the motor behavior of a fully autonomous robot with a visual interaction assistant that focuses user attention to relevant information using knowledge based techniques.

5.2.3. Mechanic role. The mechanic must be co-located with the robot as these interactions will be focused on the physical nature of the robot platform. The mechanic will need to adjust some physical aspect of the robot and then check a number of behaviors to determine if the problem has been solved. The mechanic needs the following information:

- what behaviors were failing and how
- information pertaining to any settings of mechanical parts and sensors
- software setting associated with behaviors of various sensors

In addition, the mechanic needs a way to take the robot "off-line" for testing behaviors. An issue to address here is the nature of an interface. Should an external device be used or should the robot hardware support access to this information? We have speculated that the automated diagnosis and repair domain might be beneficial to examine for possible approaches. At the present time we have not located literature that has been useful but we plan to conduct some field observations in the near future in this area.

5.2.4 .Peer role. We make the assumption that these are face to face interactions. This is the most controversial type of interaction. Our use of the terms

“peers” and “teammates” is not meant to suggest that humans and robots are equivalent but that each contributes skills to the team according to their ability. The ultimate control rests with the user –the team member or the supervisor. The issue is how the user (in this case, the peer) gets feedback from the robot concerning its understanding of the situation and actions being undertaken. In human-human teams this feedback occurs through communication and direct observation. Current research [4,5] looks at how robots should present information and feedback to its user. Bruce et. al stress that regular people should be able to interpret the information that a robot is giving them and that robots have to behave in socially correct ways to interaction in a useful manner in society. Breazeal and Scassellati [4] use perceptual inputs and classify these as social and non-social stimuli, using sets of behaviors to react to these stimuli.

Earlier work in service robots illustrate some of the issues that must be investigated for successful peer to peer interactions. [10] looked at using command and control vocabularies for mobile, remote robots including natural language interfaces. They found that users needed to know what commands were possible at any one time. This will be challenging if we determine that it is not feasible to have a separate device that can be used as an interface to display additional status from the robots that would be difficult to display via robotic gestures.

We intend to investigate research results in mixed initiative spoken language systems as a basis for communicating an understanding of the robot to the user and vice versa. Our hypothesis about information that the team mate will need include:

- What other interactions are occurring
- What the current status of the robot is
- What the robot’s world model is
- What actions are currently possible for the robot to carry out

Other interesting challenges include the distance from the robot that the team can operate. We use other communication devices to operate human-human teams from a distance. What are the constraints and requirements for robot team members?

5.2.5. Bystander role. This is perhaps the most difficult role for interaction, even though bystanders will have the most limited interactions. As described in our scenarios, a bystander role is principally concerned with co-existing in the same environment as the robot. A bystander might be a victim that the search and rescue robot has discovered in the rubble. The victim would like to be able to discover that the robot has delivered water or air and is reporting her location to the rescue team. Or a bystander might simply be a

driver passing an autonomous vehicle. What is it necessary for that bystander to know? Most drivers in that situation would want some assurance that the vehicle has equivalent skills as the majority of licensed drivers. The interface for bystanders is most likely limited to some form of behavioral indications: a robot “smile” or an action on the part of the robot, such as staying in the correct lane on the highway, that gives the bystander an indication of competence [4,5]. New experiments with robot pets and service robots (such as robot lawn mowers) will also help determine what information is needed to make bystanders comfortable with robots in their environment.

A very limited situation assessment might be possible for the bystander role. We would like to determine if the bystander understands:

- what caused the current behavior of the robot (something in the environment, something the bystander did, external forces)
- what the robot might do next, especially given an action on the part of the bystander
- the range of behaviors that the robot can exhibit
- what, if any, behaviors can be caused by the bystander

5. CONCLUSIONS

We propose that human-robot interactions are of five varieties, each needed different information and being used by different types of users. In our research we will develop a number of scenarios within a specific domain, and do a task-based analysis of these types of human-robot interactions suggested by each scenario. We will then develop both a baseline interface for the various roles and a situational assessment measurement tool. We propose to conduct a number of user experiments and make the results publicly available. Other HRI researchers can then use the same experimental design, varying either the user interfaces or the information available to the users, and compare their results to these baseline results. Our initial work will focus on the supervisory role within a driving domain. A research challenge will be what generalizes between different domains. For example, can we take what we learn in the driving domain and apply this to the search and rescue domain?

Our work in this area is interdisciplinary. Note only must we be concerned with generating the user interface we must ensure that the necessary information is available to the user. This will require coordination with experts in robotic software architectures. We have concentrated on the user and her information needs in this paper. However, to achieve a successful synergistic team, it will be necessary to furnish information about the user to the robot and to create a

dialogue space for team communication. We will start by concentrating on the user aspects of the information but intent to expand our research to include capture and use of user information as well.

Acknowledgements

This work was supported by the DARPA MARS program.

References

- [1] Albus, J., Huang, H., Messina, E., Murphy, K., et al., "4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems," NISTIR 6910, August, 2002.
- [2] Bekey, George, 1996. "Needs for Robotics in Emerging Applications: A Research Agenda", Report for IEEE/RIA Robotics and Intelligent Machines Coordinating Council Workshop on Needs for Robotics in Emerging Applications. <http://www.sandia.gov/isrc/Bekey.paper.html>
- [3] Bolstad, C. and Hess, T. 2000. Situation Awareness and Aging, in (Eds.) Mica R. Endsley and Daniel J. Garland. Situation Awareness Analysis and Measurement, Lawrence Erlbaum Associates, Mahwah, New Jersey.
- [4] Breazeal C. and Scassellati, B. 1999. A context-dependent attention system for a social robot, 1999 International Joint Conference in Artificial Intelligence.
- [5] Bruce, A., Nourbakhsh, I., & Simmons, R. 2001. The Role of Expressiveness and Attention in Human-Robot Interaction, AAAI Fall Symposium, Boston MA, October.
- [6] Endsley, M. R. 1988. Design and evaluation for situation awareness enhancement. In Proceedings of the Human Factors Society 32nd Annual meeting (Vol. 1, pp 97-1010 Santa Monica, CA: Human Factors Society.
- [7] Endsley, M. R. and Kiris, E.O. 1995, The out-of-the-loop performance problem and level of control in automation. Human Factors, 37(2), 381-394.
- [8] Endsley, M. R., 2000. Direct Measurement of situation Awareness: Validity and Use of SAGAT in (Eds) Mica R. Endsley and Daniel J. Garland. Situation Awareness Analysis and Measurement, Lawrence Erlbaum Associates: Mahwah, New Jersey.
- [9] Endsley, M. R., 2000. Theoretical Underpinnings of Situation Awareness: A Critical Review, in (Eds) Mica R. Endsley and Daniel J. Garland. Situation Awareness Analysis and Measurement. Lawrence Erlbaum Associates, Publishers: Mahwah, NJ
- [10] Engelhardt, K. G. and Edwards, R.A. 1992. Human-robot integration for Service robots (pp 315-346) in (eds) Mansour Rahimi and Waldemar Karwowski, Human robot interaction, Taylor and Francis: London.
- [11] Fong, T., Thorpe, C. and Bauer, C. 2001. Collaboration, Dialogue, and Human-robot Interaction, 10th International Symposium of Robotics Research, November, Lorne, Victoria, Australia.
- [12] Goodrich, M., Olsen, D., Crandall, J. and Palmer, T. 2001. Experiments in Adjustable Autonomy, Proceedings of the IJCAI01 Workshop on Autonomy, Delegation, and Control: Interacting with Autonomous Agents.
- [13] Gugerty, L. and Tirre, W. 2000. Individual Differences in Situation Awareness, in (Eds.) Mica R. Endsley and Daniel J. Garland. Situation Awareness Analysis and Measurement, Lawrence Erlbaum Associates: Mahwah, New Jersey.
- [14] Hill, S. and Hallbert, B. 2000. Human Interface Concepts for Autonomous/Distributed Control, Idaho National Engineering and Environmental Laboratory, DARPA ITO Sponsored Research, 2000 Project Summary. <http://www.darpa.mil/ito/psum2000/I932-0.html>
- [15] Jones, D.G. and Endsley, M. R. 1996. Sources of situation awareness errors in aviation. Aviation, Space and Environmental Medicine 67(6), pp. 507-512.
- [16] Kidd, P.T. 1992. Design of human-centred robotic systems in Mansour Rahimi and Waldemar Karwowski (Eds.) Human Robot Interaction. Taylor and Francis: London. 225-241.
- [17] Murphy, R. and Rogers, E. 1996. Cooperative Assistance for Remote Robot Supervision, Presence: volume 5, number 2, Spring, pp. 224-240.
- [18] Norman, D. 1986. Cognitive Engineering in Donald Norman and Stephen Draper (Eds.) User-centered design: new perspectives on human-computer interaction, Erlbaum Associates: Hillsdale, N.J, 31-62.
- [19] NSF/DOE, IEEE Robotics & Automation Society & Robotic Industries Association Workshop on Research Needs in Robotic and Intelligent Machines for Emerging Industrial and Service Applications, Oct 1996, Albuquerque, NM. <http://www.sandia.gov/isrc/Whitepaper/whitepaper.html>.
- [20] Scholtz, J. (2002) Creating Synergistic CyberForces in Alan C. Schultz and Lynne E. Parker (eds.), Multi-Robot Systems: From Swarms to Intelligent Automata. Kluwer.
- [21] Sheridan, Thomas B. 1992. Telerobotics, Automation, and Human Supervisory Control, MIT Press: Cambridge, MA.
- [22] Vincente, K., Roth, E. & Muman, R. 2001. How do operators monitor a complex, dynamic work domain? The impact of control room technology, Journal of Human-computer studies 54, 831-856.