

Applying CSCW and HCI Techniques to Human-Robot Interaction

Jill L. Drury

The MITRE Corporation
Mail Stop K320
202 Burlington Road
Bedford, MA 01730
jldrury@mitre.org

Jean Scholtz

National Institute of
Standards and Technology
100 Bureau Drive, MS 8940
Gaithersburg, MD 20899
jean.scholtz@nist.gov

Holly A. Yanco

Computer Science Department
Univ. of Massachusetts Lowell
One University Avenue
Lowell, MA 01854
holly@cs.uml.edu

ABSTRACT

This paper describes our approach for human-robot interaction (HRI) research and summarizes our progress to date. We have concentrated on HRI in urban search and rescue (USAR) because it is an example of a safety-critical application. We analyzed the performance of robotic teams at two USAR robotics competitions using adaptations of techniques from the human-computer interaction (HCI) field and determined that problems experienced by the operators or robots could be traced to a lack of awareness on the part of the operator of the robots' status, location, or immediate surroundings. To aid analysis, we developed a taxonomy of HRI-related characteristics, evaluation guidelines, a coding scheme that categorizes HRI activities, and a fine-grained definition of HRI awareness based on awareness research from computer-supported cooperative work (CSCW). As a result, we are beginning to determine design guidelines for HRI that are being used in developing next-generation robots at the University of Massachusetts Lowell.

Author Keywords

Human-robot interaction (HRI), Computer-Supported Cooperative Work (CSCW), awareness, urban search and rescue (USAR), human-computer interaction (HCI) evaluation techniques.

ACM Classification Keywords

H.5.2 User Interfaces, Evaluation/Methodology

INTRODUCTION

Much progress has been made in robotics in the last decade. For example, roboticists have worked hard to improve communications between humans and robots (and also between robots), the variety and fidelity of sensors on-board the robot, the ability of the robot to traverse rough

terrain, and the level of autonomy that robots are able to achieve. By comparison, relatively little progress has been made in optimizing the partnership between people and robots through improved techniques for human-robot interaction. To address this gap, our research partnership includes robotics, HCI, and CSCW expertise.

We chose to focus on USAR robots because they are a prime example of a class of safety-critical situations: situations in which a run-time error or failure could result in death, injury, loss of property, or environmental harm [Leveson 1986]. Safety-critical situations, which are usually also time-critical, provide one of the bigger challenges for robot designers due to the vital importance that robots perform exactly as intended and support humans in efficient and error-free operations.

The rest of this paper describes our methodology, analysis frameworks, results, and future work.

METHODOLOGY

There are few opportunities to study USAR operations in real disaster situations. Thus, we have used a strategy based on usability tests and robotics competitions.

We have arranged for typical users of USAR robotics to perform rescue tasks in a mock-up of a disaster situation, taking place in NIST-developed test arenas that simulate a building with various levels of destruction [Jacoff et al, 2000; Jacoff et al, 2001]. Consonant with traditional HCI usability testing, we ask participants to "think aloud" [Ericsson and Simon 1980] as they perform rescue tasks, enabling us to identify those portions of the interface that hinder participants or impede efficiency.

However, to date, most of our opportunities to study HRI have come in conjunction with USAR robotics competitions. These opportunities differed from traditional usability testing in two significant ways. First, the robot developers operated the robots (rather than members of the rescue professions). We viewed performance, therefore, as an upper bound: if the robot developers had problems with a part of the interface, it is likely that any other user would also have difficulties. Second, we were restricted to being silent observers who could not ask the operators to do anything differently during the competition than they would

have already done. To at least partially make up for a lack of “thinking aloud,” our observer performed a quick debriefing of the operators via a short post-run interview to obtain the operators’ assessment of their (and the robots’) performance. In addition, we were given the scoring materials from the competition judges that indicated where victims were found and penalties that were assessed.

HRI Taxonomy

To better understand the different types of HRI, we developed a taxonomy to characterize robotic interaction [Yanco and Drury 2002]. Besides determining the classification categories, we defined values to describe each classification. The list of classification categories and their description is contained in Table 1 (values are omitted here due to space limitations).

ANALYSIS FRAMEWORKS

We feel that some of our more important contributions to HRI are our analysis frameworks, since they may help other researchers, robotic designers and evaluators to better understand when and how HRI can be improved.

We have used three different mechanisms to structure our analyses: a detailed definition of HRI awareness, a coding scheme for HRI activities, and Scholtz’ [2002] evaluation guidelines. Each of these mechanisms is discussed below. All three led us to focus on “critical incidents,” which we defined as cases in which the robot, USAR victims, or environment sustained actual or potential damage or harm.

Classification	Description
Autonomy	% time a robot performs a task on its own
Amount of intervention	% time that a human operator must control a robot
Human-robot ratio	The ratio of operators to robots.
Level of shared interaction	Various combinations of whether the humans and robots act independently or as part of team(s).
Composition of robot teams	Whether teams of multiple robots are homogeneous or heterogeneous.
Available sensors	A list of sensor types available on the robot platform.
Sensor fusion	A list of functions mapping the sensor data to the fused output.
Criticality	The importance of getting the task done correctly in terms of its negative effects should problems occur.
Time	Whether the humans and robots work together in the same time (synchronously) or different times (asynchronously).
Space	Whether the humans and robots work together in the same place (collocated) or in different places (non-collocated).

Table 1: Taxonomy for Human-Robot Interaction [Yanco and Drury 2002]

HRI Awareness

Much research has been performed in the CSCW community to characterize awareness. While there are many definitions of awareness in the CSCW literature (see Drury, Scholtz, and Yanco [2003] for a summary), we started with the definition in Drury [2001], the informal version of which is: awareness in a multi-user computing system is a participants’ understanding of the presence, identities, and activities of another participant. There are two differences between CSCW and robotic systems that affect how awareness can be understood, however. The first difference is the fact that CSCW addresses multiple humans working together, whereas HRI can involve single or multiple humans working with single or multiple robots. The second is that human participants will bring at least a minimum level of free will and cognitive ability to the collaboration that cannot be brought by the robotic participants. Thus the HRI awareness framework must account for all combinations of single and multiple humans and robots, and must accommodate the non-symmetrical nature of the human-robot collaboration. The simplest case of HRI occurs when one human works with one robot.

HRI awareness (base case): Given one human and one robot working on a task together, HRI awareness is the understanding that the human has of the location, activities, status, and surroundings of the robot; and the knowledge that the robot has of the human’s commands necessary to direct its activities and the constraints under which it must operate.

Note that greater or lesser amounts of HRI awareness are needed depending upon the level of autonomy that the robot achieves, so the expectations of awareness need to be tailored for the anticipated level of autonomy. The HRI awareness base case can be generalized to cover multiple humans and robots coordinating in real time on a task.

HRI awareness (general case): Given n humans and m robots working together on a synchronous task, the general case of HRI awareness consists of five components:

- Human-robot: the understanding that the humans have of the locations, identities, activities, status and surroundings of the robots. Further, the understanding of the certainty with which humans know the aforementioned information.
- Human-human: the understanding that the humans have of the locations, identities and activities of their fellow human collaborators.
- Robot-human: the knowledge that the robots have of the humans’ commands necessary to direct their activities and any human-delineated constraints that may require a modified course of action or command noncompliance.

- Robot-robot: the knowledge that the robots have of the commands given to them, if any, by other robots, the tactical plans of the other robots, and the robot-to-robot coordination necessary to dynamically reallocate tasks among robots if necessary.
- Humans' overall mission awareness: the humans' understanding of the overall goals of the joint human-robot activities and the moment-by-moment measurement of the progress obtained against the goals.

In human-robot awareness, “activities” refer to such phenomena as speed and direction of travel and progress towards executing commands. Status information includes battery power levels, the condition of sensors, etc.

Sufficient HRI awareness is needed to ensure smoothly functioning human-robot coordination on a shared task. When insufficient HRI awareness is provided, we say this is an HRI awareness violation:

HRI awareness violation: HRI awareness information that should be provided is not provided.

There are five possible types of HRI awareness violations, corresponding to the five types of HRI awareness defined above. We discussed the results from a USAR competition in terms of types of awareness violations that occurred during critical incidents in Drury, Scholtz and Yanco [2003].

Coding Scheme

To help in analyzing videotapes of the robot competitions and usability test runs, we noted each critical incident and categorized it in terms of the type of HRI awareness violation that occurred (if one was present) and the type of task being attempted at the time of the incident.

Because all cases that we analyzed so far concerned a single operator and one or more robots that did not coordinate with each other, HRI awareness problems consisted solely of human-robot awareness violations. We anticipate that more of the HRI awareness framework will be employed when we analyze more diverse configurations.

We defined five types of tasks relating to critical incidents.

Local navigation: An operator is navigating in constrained or tight situations, and encounters difficulty because of the constraints. An example of a local navigation problem is when the robot slips down a ramp or bumps a wall.

Global navigation: An operator is navigating in all other situations. An example of a global navigation problem is when an operator does not have a clear understanding of the robot's position, potentially

leading to driving the robot out of the arena unintentionally or covering areas already searched.

Obstacle encounter: An operator is working to free the robot from an obstacle; the robot is hindered in moving towards a goal.

Victim identification: An operator is attempting to characterize the state of a victim (e.g., conscious or not, warm or cold, speaking or silent, moving or not moving). An example of a problem occurring during victim identification is inaccurate interpretation of sensor data.

Vehicle state: An operator is attempting to perform USAR tasks despite the fact that the robot is in a degraded state (e.g., it is not stable or upright or its sensors are impaired or broken).

We analyzed data from a USAR competition using this coding scheme; the results are summarized in Scholtz, Young, Drury, and Yanco [in submission].

Scholtz's Guidelines

Scholtz [2002] developed six evaluation guidelines for evaluating HRI. We treated these guidelines as heuristics to be tailored for USAR systems (Nielsen [1993] recommends tailoring heuristics to be appropriate to the systems being evaluated). After tailoring (including combining two of the guidelines into one heuristic), we evaluated the robotic systems in a major USAR competition against the following:

Is sufficient status and robot location information available so that the operator knows the robot is operating correctly and avoiding obstacles?

Is the information coming from the robots presented in a manner that minimizes operator memory load, including the amount of information fusion that needs to be performed in the operators' heads?

Are the means of interaction provided by the interface efficient and effective for the human and the robot (e.g., are shortcuts provided for the human)?

Does the interface support the operator directing the actions of more than one robot simultaneously?

Will the interface design allow for adding more sensors and more autonomy?

A discussion of how we tailored these heuristics, plus our results after applying the heuristics, is contained in Yanco, Drury, and Scholtz [to appear].

RESULTS TO DATE

We found that all critical incidents could be traced to HRI awareness violations. Thus, when we developed a preliminary set of guidelines for designing interfaces for HRI [Yanco, Drury, and Scholtz, to appear], we began with

awareness and also included guidelines to address the other major problems we observed with HRI:

Enhance awareness. Provide a map of where the robot has been. (Operators using systems with maps were more successful in navigating the arena.) Also, provide more spatial information about the robot in the environment; operators must be aware of their robots' immediate surroundings to avoid bumping into obstacles or victims.

Lower cognitive load. Provide fused sensor information to avoid making the user fuse the data mentally.

Increase efficiency. Provide user interfaces that support multiple robots in a single display/window. In general, minimize the use of multiple windows. With additional sensor fusion, more information could be displayed in a single window, which is more efficient for users than having to switch between windows.

Provide help in choosing robot modality. Provide the operator assistance in determining the level of robotic autonomy that would be most appropriate for a given situation.

DISCUSSION AND FUTURE WORK

One of the primary goals of our further research is to expand and refine our set of design guidelines. We have taken the design guidelines developed so far and are in the process of applying them to new robots and interfaces being developed at the University of Massachusetts Lowell [Hestand and Yanco, in submission].

We found coding to be very difficult at times. Our first attempt at coding (not described in this paper) involved accounting for every second of human/robot activities; we found that the detailed data did not yield as many insights as hoped. In contrast, the scheme described in this paper concentrated on characterizing anomalous behavior, analogous to an HCI expert concentrating on users' problems operating interfaces during usability testing. We anticipate that the coding scheme will likely evolve further.

We plan to expand our use of HCI analytical and inspection evaluation techniques. For example, we anticipate performing a Goals, Operators, Methods, Selection rules (GOMS) analysis of several robotic systems.

Few of the robots studied so far include much autonomy. We plan to investigate HRI under varying levels of autonomy, especially via usability testing.

As we evaluate systems that include multiple humans and robots that communicate with each other, we plan to more fully exercise the HRI awareness framework and determine whether it should evolve.

We also plan to refine the taxonomy. By characterizing the robotic system in a useful way, we hope to be able to use

the taxonomy to roughly predict the likely level of efficiency and cognitive load.

ACKNOWLEDGMENTS

The work was supported in part by NSF IIS-0308186, NIST 70NANB3H1116, and the DARPA MARS program.

REFERENCES

1. Drury, J. (2001). Extending Usability Inspection Evaluation Techniques for Synchronous Collaborative Computing Applications. Sc.D. Thesis, University of Massachusetts Lowell, Department of Computer Science.
2. Drury, J., J. Scholtz and H.A. Yanco (2003). Awareness in Human-Robot Interactions. In Proceedings of the IEEE Conference on Systems, Man and Cybernetics, Arlington, VA, October 2003.
3. Ericsson, K. A. and H. A. Simon (1980). "Verbal reports as data." *Psychological Review* **87**: 215 - 251.
4. Hestand, D. and H.A. Yanco. Layered Sensor Modalities for Improved Feature Detection. Submitted to ICRA-04.
5. Jacoff, A., E. Messina, and J. Evans (2001). "A reference test course for autonomous mobile robots." In *Proceedings of the SPIE-AeroSense Conference*, Orlando, FL, April 2001.
6. Jacoff, A., E. Messina, and J. Evans (2000). "A standard test course for urban search and rescue robots." In *Proceedings of the Performance Metrics for Intelligent Systems Workshop*, August 2000.
7. Leveson, N.G. (1986). "Software safety: why, what and how." *ACM Computing Surveys* **18**(2): 125 - 162, June 1986.
8. Nielsen, J. (1993). Usability Engineering. Chestnut Hill, MA, AP Professional Press.
9. Scholtz, J. (2002). "Evaluation methods for human-system performance of intelligent systems". In *Proceedings of the 2002 Performance Metrics for Intelligent Systems (PerMIS) Workshop*.
10. Scholtz, J., J. Young, J.L. Drury, and H.A. Yanco. Evaluation of Human-Robot Interaction Awareness in Search and Rescue. Submitted to 2004 IEEE International Conference on Robotics and Automation.
11. Yanco, H.A. and J.L. Drury (2002). "A taxonomy for human-robot interaction." In *Proceedings of the AAAI Fall Symposium on Human-Robot Interaction*, November 2002, Falmouth, MA. AAAI Technical Report FS-02-03, pp. 111 - 119.
12. Yanco, H.A., J.L. Drury and J. Scholtz. Beyond Usability Evaluation: Analysis of Human-Robot Interaction at a Major Robotics Competition. To appear in the *Journal of Human-Computer Interaction*.