

Awareness in Human-Robot Interactions*

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 provides a set of definitions that form a framework for describing the types of awareness that humans have of robot activities and the knowledge that robots have of the commands given them by humans. As a case study, we applied this human-robot interaction (HRI) awareness framework to our analysis of the HRI approaches used at an urban search and rescue competition. We determined that most of the critical incidents (e.g., damage done by robots to the test arena) were directly attributable to lack of one or more kinds of HRI awareness.*

Keywords: Awareness, human-robot interaction, search and rescue, critical incident analysis, human-computer interaction.

1 Introduction

Computer-Supported Cooperative Work (CSCW) is computer-assisted coordinated activity carried out by groups of collaborating individuals [8]. CSCW software, also called groupware, “is distinguished from normal software by the basic assumption that it makes: groupware makes the user aware that he is part of a group, while most other software seeks to hide and protect users from each other” [15]. We maintain that the human-robot interface is akin to groupware in the sense that humans must use the interface to orchestrate joint human/robot activities. Further, the humans must be aware of the robots’ status and activities via the interface in the cases where visual contact cannot be maintained with the robot. Given these connections to CSCW and groupware, we have mined the CSCW literature for insights that can be applied to human-robot interaction (HRI).

In the CSCW literature, information that collaborators have about each other in coordinated activities is commonly called awareness information. It helps them know who else is working in a shared workspace and what the others are doing. Designed to emulate the kinds of

non-verbal cues that people get when they collaborate face-to-face in the same physical location, awareness information is important for effective collaboration and coordination.

In this paper, we present a framework for understanding awareness in HRI, and use this framework to analyze the HRI performance of four different robotic systems. We believe that this framework can be used by researchers developing methods to evaluate awareness support in HRI.

2 Related work on awareness

There are many definitions of awareness, such as those listed in Table 1. There is no standard definition of awareness yet in the CSCW field and we are unaware of any definitions of awareness specifically tailored for HRI. Understanding of the different types of awareness associated with computer-based systems is still evolving.

The definitions of awareness summarized in Table 1 address awareness in general, as well as awareness specific to tasks, a shared workspace, or the larger environment in which the collaborative activities take place. The common thread among the definitions is the understanding that the participants have of each other in the CSCW environment.

Most of the definitions in Table 1 (adapted from [6]) are somewhat informal (e.g., “where people know roughly what other people are doing” [2]). Our framework adapts and expands one of the more precise definitions of awareness [5] for use in describing awareness in HRI.

3 HRI awareness framework

There are two differences between CSCW and robotic systems that significantly affect how awareness can be analyzed. The first is the fact that CSCW addresses

Table 1. Definitions of Awareness in the CSCW Literature

Awareness term	Definition	Source
awareness	an understanding of the activities of others, which provides a context for your own activities	Dourish and Bellotti [4]
awareness	given two participants p_1 and p_2 who are collaborating via a synchronous collaborative application, awareness is the understanding that p_1 has of the identity and activities of p_2	Drury [5]
concept awareness	the participants' understanding of how their tasks will be completed	Gutwin et al. [11]
conversational awareness	who is communicating with whom	Vertegaal et al. [18]
group-structural awareness	knowledge about such things as people's roles and responsibilities, their positions on an issue, their status, and group processes	Gutwin et al. [9]
informal awareness	the general sense of who is around and what others are up to	Gutwin et al. [9]
peripheral awareness	showing people's location in the global context	Gutwin et al. [10]
peripheral awareness	where people know roughly what others are doing	Baecker et al. [2]
situation awareness	the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future	Endsley [7]
social awareness	the understanding that participants' have about the social connections within their group	Gutwin et al. [11]
social awareness	information about the presence and activities of people in a shared environment	Prinz [16]
(not named by authors; our term is spacial awareness)	the more an object is within your focus, the more aware you are of it; the more an object is within your nimbus, the more aware it is of you	Benford and Fahlen [3]
task awareness	the participants' understanding of how their tasks will be completed	Gutwin et al. [11]
task-oriented awareness	awareness focused on activities performed to achieve a shared task	Prinz [16]
workspace awareness	the up-to-the-minute knowledge of other participants' interactions with the shared workspace	Gutwin et al. [11]
workspace awareness	who is working on what	Vertegaal et al. [18]

multiple humans working with a CSCW application, whereas HRI can involve single or multiple humans working with single or multiple robots. The second is that, while robots can be thought of as participants in the collaborative activities, human participants will bring some 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. By calling out distinct awareness needs for the human and the robot, this "base case" makes the non-symmetrical awareness relationship clear.

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.

Obviously, greater or lesser amounts of HRI awareness are needed depending upon the level of autonomy that the robot is expected to achieve, so the expectations of awareness need to be tailored for the expected level of robot autonomy and the roles played by the human collaborators. Scholtz [17] defines human roles in the context of robotic systems as supervisor, operator, mechanic, teammate, and bystander. The HRI awareness framework focuses on the operator: the person most directly controlling the robot's activities.

The base case can be generalized to multiple humans and robots coordinating in real time on a task. Due to the non-symmetrical nature of HRI awareness, four distinct cases need to be defined. We refer to the awareness that the human has of each robot ("human-robot"), the robot has of each human ("robot-human"), the human has of other human(s) ("human-human"), and each robot has of the other robot(s) ("robot-robot").

Finally, due to the need for the human(s) to coordinate the efforts of multiple humans, multiple robots, or both, a fifth type of awareness was defined to encompass the humans' overall understanding of the joint goals and activities. The resulting definitions follow.

HRI awareness (general case): Given n humans and m robots working together on a synchronous task, 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 robots' knowledge of the humans' commands needed to direct activities and any human-delineated constraints that may require command noncompliance or a modified course of action.

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 measurement of the moment-by-moment 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. Examples of status information are battery power levels and the condition of sensors. "Surroundings" refer to both the changing and unchanging parts of the robot's physical environment. Note that we speak of humans having understanding but the robots having "knowledge."

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.

These specific concepts of HRI awareness have not previously been applied to HRI. To understand their utility in analyzing HRI performance, we gathered data at the American Association of Artificial Intelligence (AAAI) 2002 Robot Rescue Competition and analyzed the performance of four different teams in terms of HRI awareness violations. (For the sake of brevity, we will often drop the "HRI" from "HRI awareness" and "HRI awareness violation" and speak of "awareness" and "awareness violations.")

4 Applying the awareness framework

The search and rescue domain was chosen because it is a prime example of a safety-critical situation (defined as a situation where a run-time error or failure could result in death, injury, loss of property, or environmental harm [14]). Safety-criticality imposes a requirement for error-

free operation and is also often time-critical, resulting in a special need for efficient, intuitive HRI. We focused on the effectiveness of techniques for making human operators aware of pertinent information regarding the robot and its environment.

The goal of the AAAI-2002 Robot Rescue Competition was to find and accurately map the locations of victims in a simulated urban disaster situation. The robots competed in the Reference Test Arenas for Autonomous Mobile Robots developed by the National Institute of Standards and Technology [12, 13].

4.1 Methodology

The competition uses rules and a scoring algorithm developed by a joint rules committee consisting of domain experts and researchers from the RoboCup and AAAI communities [1]. The scoring algorithm was designed to address several issues that arise in real urban search and rescue situations, including the number of people required to operate the robots (fewer rescue personnel needed to control robots), the percentage of victims found, the number of robots that find unique victims (leading to quicker search times), and the accuracy of victim reporting (it is best to be as localized as possible). There are also penalties for bumping into victims or the environment. We used the competition scoring as an objective measure of how well each team performed, and compared this performance to the types and severity of HRI awareness violations observed during the competition.

All teams voluntarily agreed to participate in our study. (Note that, although we collected data from the entire competition, we restricted our data analysis to that of the top four teams since those were the only teams to find victims.) We observed the operator(s) of each team's robot(s) during three 15 minute runs of the competition. The operators were videotaped while operating the interface and the interface screens were also recorded via videotape. Further, the robots were videotaped with cameras placed in various locations around the arena. We were silent observers, not asking the operators from the team to do anything differently during the competition run than they would have already done. At the conclusion of each run, an observer performed a quick debriefing of the operator via a post-run interview. In addition to collecting data from the team operators, we were able to collect data from a search and rescue expert: a fire chief was a judge for the competition and agreed to use the robots as well. He was tested two of the robot systems.

The resulting data consisted chiefly of videotapes, competition scoring sheets, maps of robot paths, questionnaire/debriefing information, and researcher observation notes. To make the most of the videotaped information, we developed a coding scheme to capture the number and duration of occurrences of various types of

activities observed. Our scheme consists of a two-level hierarchy of codes: header codes capture the high-level events and primitive codes capture low-level activities. The header codes were defined as: identifying a victim, robot logistics (e.g., undocking small robots from a larger robot), failures (hardware, software, or communications), and navigation/monitoring (directing the robot or observing its autonomous motion when the other header codes do not apply). Three primitive codes were defined: monitoring (watching the robot when it is in an autonomous mode), teleoperation (“driving” the robot), and user interface manipulation (interacting with the interface to control the robot).

4.2 Overview of user interfaces

We studied the interfaces and performance of four teams, denoted Team A, B, C, and D for anonymity.

4.2.1 Team A

Team A developed a heterogeneous robot team of five robots, one iRobot ATRV-Mini and four Sony AIBOs, for the primary purpose of research in computer vision and multi-agent systems. All robots were teleoperated serially. The four AIBOs were mounted on a rack at the back of the ATRV-Mini. The AIBOs needed to be undocked to start their usage and redocked after they were used if the operator wanted to take them with the larger robot. Team A developed two custom user interfaces, which were created for use by the developers: one for the ATRV-Mini and another for the AIBOs.

The user interface for the ATRV-Mini had multiple windows: a video image taken by the robot, a map constructed by the robot using the SICK laser scanner and odometry, the raw laser scan information presented as lines showing distance from the robot, and a window with eight radio buttons to allow the user to switch camera views. The operator drove the robot using keys on the keyboard to move forward, backward, right and left.

The user interface for the AIBOs had a window with the video image sent from the robot. The operator controlled the robots using buttons on the GUI or with the keyboard.

4.2.2 Team B

Team B had developed their robot system for use in hazardous environments. The control of the robot was selected from teleoperated (operator does all control), safe (operator drives robot, but sensors are used to prevent user from driving into obstacles), shared (robot and operator both control parts of the driving), and autonomous (robot drives itself). Communication between the user interface and the robot was achieved through a proprietary, low-bandwidth communication protocol over 900 MHz radio.

The custom user interface was developed for expert

users and displayed on a touch screen. In the upper left corner was the video feed from the robot. Tapping the sides of the window moved the camera left, right, up or down. Tapping the center of the window recentered the camera. (During the competition, the window was not finished, so the video was displayed on a separate monitor, but the blank window was still tapped to move the camera.) The robot was equipped with a color video camera and a thermal camera. The operator could switch the video feed between these two cameras. In the lower left corner was a window displaying sensor information such as battery level, heading, and tilt of the robot. In the lower right corner, a sensor map was displayed, showing filled red areas to indicate blocked directions.

The robot was controlled through a combination of a joystick and the touch screen. To the right of the sensor map, there were buttons for various operations modes: autonomous, shared (a semiautonomous mode in which the operator can “guide” the robot in a direction but the robot does the navigation and obstacle avoidance), safe (in which the user controls the navigation of the robot, but the robot uses its sensors to prevent the user from driving into obstacles), and teleoperation (the human controller is totally responsible for directing the robot). Typically, the operator would click on one of the four mode buttons, then start to use the joystick to drive the robot. When the operator wished to take a closer look at something, perhaps a victim or an obstacle, he would stop driving and click on the video window to pan the camera. For victim identification, the operator would switch over to the thermal camera for verification.

4.2.3 Team C

Team C used two identical robots, RWI Magellan Pros. Communication between the user interface and robots was achieved with an RF modem. The robots had a mixed level of autonomy: they could be fully teleoperated or the robots could provide obstacle avoidance. The robots were operated serially, but could run simultaneously. The primary command of the robots was achieved by giving them relative coordinates to move towards. The robots then autonomously moved to that location using reactive obstacle avoidance. This allowed for the perception that the operator moved both robots “at once,” even though he controlled them serially. It was the operator’s trust in the autonomy of the robots that allowed this type of operation.

A custom interface was developed for a “sophisticated user” (according to the developers). Team C started their run 1 using a graphical user interface, but switched back to a text-based interface when there were command latency problems with the GUI. In the GUI, the screen was split into two down the middle. Each side was an interface to one of the two robots. The top window for each robot displayed a current video image from the robot and the bottom window displayed map information.

The text-based interface had 14 text windows and 4 graphic windows, half for each of the robots. The 7 text windows were for the following: the IPC (interprocess communication) server, the navigation module, the vision module, the mapping module, the navigation command line, for starting and monitoring the visual display, and for starting and monitoring the map display. The two graphic windows were for displaying the camera image and the map image. The computer ran an enlarged desktop during the competition, and the operator sometimes needed to switch to another part of the desktop (effectively switching to another screen) for other pieces of the interface. The robots were controlled with the keyboard.

4.2.4 Team D

Team D developed two custom built robots for search and rescue, one wheeled and one tracked, with the same sensing and operating capabilities. The robots were teleoperated serially. A wireless modem was used to communicate between the user interface and the robots.

Team D developed a custom user interface on two screens. One monitor displayed the video feed from the robot that was currently being operated. The other monitor had a pre-entered map of the arena, on which the operator would place marks to represent the locations of victims that were found. The robots were driven with keyboard controls.

4.3 Results

We examined critical incidents occurring during the competition to determine their causes, and then analyzed the incidents and other results in terms of the five different types of HRI awareness. We define critical incidents as anomalous situations in which the operator or robot encounters a problem.

4.3.1 Critical incidents

All teams in our study experienced critical incidents. In this section, we present a selection of critical incidents.

Team A deployed small dog-like robots (AIBOs) off of docks on the larger, ATRV-Mini robot. The dogs were particularly useful for getting into small spaces that were difficult for the ATRV-Mini to maneuver in. Once, however, an AIBO fell off, then became trapped under fallen Plexiglas, yet the operator was did not know this.

At one point when a search and rescue expert was using Team B's robot, he was frustrated because the robot would not travel forward into an apparently clear space as seen on video while in safe mode. He put the robot into teleoperation mode and viewed the video to guide his choice of commands. He then drove through a Plexiglas panel, which he could not see using video. However, the sonar sensors on the robot were picking up the obstacle

and had indicated the blockage on the sensor map, which was located on a different screen than the video monitor.

During run 1, the Team B operator moved the robot's video camera off-center to perform a victim identification. After the victim identification, he let the robot maneuver itself out of the tight area in shared mode, forgetting to re-center the camera. When he switched back to safe mode, he thought his camera was pointing forward, when it was really pointed 90 degrees to the left. This resulted in the operator accidentally driving the robot out of the arena into the crowd, and bumping into a wall trying to get back into the arena. It also resulted in substantial operator confusion (we recorded quotes such as, "it's really, really hard," "I got disoriented," "hmmmmn," and "oh, no!"). During run 3, Team B's operator did not have good visibility into the areas behind the robot, making it difficult for him to maneuver it out of narrow spaces ("this is very difficult"). After the last run, Team B's operator commented that he had not bumped into anything, four bumping penalties were assessed by the judges.

Team C started a run using a graphical user interface (GUI), but within two minutes, the operator determined that there was too much lag time between command issuance and response. As a result, he shut down the GUI windows and brought up nine windows that formed an earlier version of the interface (the debugging version). It took a little over a minute and a half for the operator to shut down the GUI and bring up all the windows for the earlier interface version. To operate the robot, the operator needed to shuffle through the nine windows to view different types of information and enter commands in several of the windows.

4.3.2 Analysis using the HRI awareness framework

In this section, we examine each of the five types of HRI awareness: human-human, robot-robot, human-robot, robot-human, and overall mission awareness.

The competition did not afford an opportunity to view human-human awareness, since one human directed the robot(s) in each case. Although Teams A, C, and D fielded multiple robots, the robots did not communicate with each other, leading to limited robot-robot awareness. The robots with video cameras could "see" each other when they were within each other's fields of view but the robots did not act on that information except to treat the other robots as additional obstacles to avoid. This leaves the three remaining cases of general HRI awareness.

Once again, robot-human awareness is the knowledge that the robots have of the humans' commands and any human-originated constraints that may require a modified course of action or command noncompliance. An example of the latter condition is when the robot is given a

command to continue straight for six feet yet the robot senses a wall of the arena three feet ahead. This situation indicates a modified course of action: continuing straight for two feet and then stopping. The fact that there were multiple instances of robots running into walls or victims during the competition indicates insufficient robot-human awareness of the constraints imposed by the environment.

We did not note any cases in which the robots did not have sufficient awareness of the direct commands given them by the human operators; however, there were many times that the communications link was lost and the robots had to cease operations until communication was resumed. Additionally, Team C tried to use the GUI in run 1, but the communication to the robot lagged enough to hinder robot-human awareness; once Team C started the text-based interface, this awareness problem ceased.

Many instances of human-robot awareness violations were noted. There are five parts to human-robot awareness: humans' understanding of the identities, locations, surroundings, status, and activities of the robots.

Human-robot awareness of a robot's identity would be violated if an operator could not identify which robot was which. We did not note any of this type of awareness violation. When there were multiple robots, the interfaces were duplicated and contained the identity of the robot.

We did note many instances of awareness violations regarding location and surroundings. In several cases (e.g., Team B during run 3 and Team D overall), there was not enough awareness of the area immediately behind the robot, causing the robot to bump obstacles when backing up or maneuvering the robots in small spaces. Even when moving forward, several operators (including Team B during run 3) hit walls and were not aware of doing so. The operator was not aware of the robot's size in relationship to the space through which the robot was attempting to navigate. Team A's trapped dog also constituted an awareness violation of this type.

When the domain expert was using Team A's system, we observed that he would use the video from the larger robot, the ATRV-Mini, to watch the progress of the AIBO. This provided a solution to the human-robot awareness of robot surroundings, yet required the larger robot to stay within sight of the smaller robot.

We had postulated that the teams with the highest levels of awareness of their surroundings and locations would have the richest sets of different types of sensory data, and the most effective ways of fusing this data. An objective means of measuring human-robot location and surroundings awareness violations is via counting the number of bumping penalties incurred. Table 2 shows the teams' ranking regarding bumping, with Team C having the fewest bumping penalties and Team D having the

most. Table 2 also shows the types of sensors used by the robots and whether the sensor information had been combined.

Table 2. Bumping Performance versus Sensor Information

Rank	Team	Sensor Types Provided	Sensor Fusion
1	C	Video, sonar, infrared	Sonar, infrared fused into overhead map
2	A	Video, laser ranging [primary robot]	Odometry, laser ranging fused into laser map
3	B	Raw video, raw thermal imaging, sonar, infrared, bump, laser ranging	Infrared, bump, laser ranging, and sonar fused into sensor map
4	D	Video	None

According to our hypothesis, Team B should have placed first; instead it placed third. The reason for the lower-than-expected ranking is because Team B also experienced an awareness violation of a different sort: human-robot awareness of robot status. The fact that Team B's interface did not provide any reminders that the video camera was pointing off-center meant that this lack of awareness of robot state caused him to run into more obstacles and find fewer victims.

The last type of human-robot awareness pertains to the robots' activities. In general, if the operator did not have a good idea of the robots' locations or surroundings, he or she also did not have a good idea of the robots' activities (especially if "bumping into a wall" could be considered to be an activity). In one case, the Team A operator knew where his ATRV-Mini robot was located but didn't realize that an AIBO had fallen off of the ATRV-Mini.

Overall mission awareness has two parts: the humans' understanding of the overall goals and the moment-by-moment measurement of the progress obtained against the goals. In all cases, the operators were quite aware of the goal to find as many victims as possible with the fewest penalties possible. Because this was such a clear-cut task, the operators were able to measure their progress in terms of how many victims they had found so far.

In terms of progress exploring the arena, most teams were able to construct reasonable maps showing the portion of the arena that had been explored. Teams A and C had automated this process and displayed maps created by their robots in their interface. We see in table 2 above that these two teams had the fewest number of bumping penalties assessed.

Team B and D did not have the robot create maps of the arena for the operator. Both of these teams experienced more difficulty navigating, receiving a greater number of bumping penalties. For Team B, the operator had an additional awareness problem as he was not able to comprehend how much of the arena had been covered because of the previously described problem with the angle that the camera was pointed.

5 Discussion

All critical incidents were due to some type of awareness violation. In fact, they were primarily due to a lack of human-robot awareness of location and surroundings. Because of the makeup and technical capabilities of the human-robot teams, we did not have opportunity to apply the human-human nor robot-robot awareness criteria. We feel the latter two types of HRI awareness will be more useful as robotic systems become more complex and the robots communicate among themselves. It will also be harder to perceive robot-robot awareness violations by observation, since the robots will not frown or express confusion as humans do.

HRI awareness is a useful concept for evaluating human-robot interfaces. By watching humans and robots interact and classifying instances of awareness violations, the community can generate guidelines for information needs in human-robot interfaces.

References

- [1] AAI/RoboCup Rules Committee, "Rules for the RoboCup and AAI 2002 Robot Rescue Competition," <http://www.csee.usf.edu/robotics/USAR/RobotRescue2002/>
- [2] R. M. Baecker, D. Nastos, I. R. Posner and K. L. Mawby, "The user-centered iterative design of collaborative writing software," Proc. InterCHI '93, Amsterdam, The Netherlands, ACM, 1993.
- [3] S. Benford, S. and L. Fahlen, "A spatial model of interaction in large virtual environments," Proc. 3rd Euro. Conf. on CSCW, Milan, Italy, 1993.
- [4] P. Dourish and V. Bellotti, "Awareness and coordination in shared workspaces," Proc. CSCW '92, Toronto, Canada, 1992.
- [5] J. L. Drury, "Extending usability inspection evaluation techniques for synchronous collaborative computing applications," Sc.D. Dissertation, Computer Science Dept, University of Massachusetts Lowell, 2001.
- [6] J. L. Drury and M. G. Williams, "A framework for role-based specification and evaluation of awareness support in synchronous collaborative applications," Proc. IEEE Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE) 2002, Pittsburgh, PA, 2002.
- [7] M. R. Endsley, "Design and evaluation for situation awareness enhancement," Proc Human Factors Society 32nd Annual Meeting, Santa Monica, CA, 1988.
- [8] I. Greif, *Computer-Supported Cooperative Work: A Book of Readings*, Morgan Kaufmann, 1988.
- [9] C. Gutwin, S. Greenberg and M. Roseman, *Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation*, University of Calgary, 1996.
- [10] C. Gutwin, S. Greenberg and M. Roseman, "Workspace awareness support with radar views," Proc. CHI 96 Conference on Human Factors in Computing Systems, Vancouver, British Columbia, Canada, 1996.
- [11] C. Gutwin, G. Stark and S. Greenberg, "Support for workspace awareness in educational groupware," Proc. Comp. Supported Collaborative Learning (CSCL), 1995.
- [12] A. Jacoff, E. Messina, and J. Evans, "A reference test course for autonomous mobile robots," Proc. SPIE-AeroSense Conference, Orlando, Florida, April 2001.
- [13] A. Jacoff, E. Messina, and J. Evans, "A standard test course for urban search and rescue robots," Proc. Performance Metrics for Intelligent Systems Workshop, August 2000.
- [14] N. G. Leveson, "Software safety: why, what and how," *ACM Computing Surveys*, Vol 18, No. 2, pp. 125 – 162, June 1986.
- [15] K. Lynch, J. Snyder and D. Vogel, "The Arizona analyst information system: supporting collaborative research on international technology trends," In *Multi-User Interfaces and Applications*, S. Gibbs and A. Verrijn-Stuart, Eds., North-Holland, 1990.
- [16] W. Prinz, "NESSIE: an awareness environment for cooperative settings," Proc. 6th Euro. Conf. on Computer Supported Cooperative Work, Copenhagen, 1999.
- [17] J. Scholtz, "Human-robot interactions: creating synergistic cyber forces," In *Multi-Robot Systems: From Swarms to Intelligent Automata (Proceedings from the 2002 NRL Workshop on Multi-Robot Systems)*, A. C. Schultz and L. E. Parker, Eds., Kluwer Academic Publishers, 2002.
- [18] R. Vertegaal, B. Velichkovsky and G. van der Veer, "Catching the eye: management of joint attention in cooperative work," *SIGCHI Bulletin*, Vol 29, No. 4, 1997.