Up from the Rubble: Lessons Learned about HRI from Search and Rescue

Robin R. Murphy & Jennifer L. Burke
University of South Florida

The Center for Robot-Assisted Search and Rescue has collected data at three responses (World Trade Center, Hurricane Charley, and the La Conchita mudslide) and nine high fidelity field exercises. Our results can be distilled into four lessons. First, building situation awareness, not autonomous navigation, is the major bottleneck in robot autonomy. Most of the robotics literature assumes a single operator single robot (SOSR), while our work shows that two operators working together are nine times more likely to find a victim. Second, human-robot interaction should not be thought of how to control the robot but rather how a team of experts can exploit the robot as an active information source. The third lesson is that team members use shared visual information to build shared mental models and facilitate team coordination. This suggests that high bandwidth, reliable communications will be necessary for effective teamwork. Fourth, victims and rescuers in close proximity to the robots respond to the robot socially. We conclude with observations about the general challenges in human-robot interaction.

INTRODUCTION

A 2001 DARPA/NSF Study on Human-Robot Interaction concluded that search and rescue is a good domain for studying human-robot interaction (Burke, Murphy, Rogers, Lumelsky, & Scholtz, 2004). Urban search and rescue (US&R) is the emergency response function that deals with collapsed man-made structures. Rescue robotics is a *field application*, since the robots a) work in an unpredictable environment and b) serve as an extension of a human(s) in these situations. As noted in Murphy (2004), rescue robotics is interesting because humans cannot conduct search and rescue in certain structures without robots, the robot operators are "average" users who deploy robots as one type of tool rather than use only robots, the robots mediate the presence of human operators and observers (e.g., structural engineers, medical personnel, safety officer, etc.), and the use of robots must fit into an existing task and personnel hierarchy.

The Center for Robot-Assisted Search and Rescue (CRASAR) has collected HRI data at three actual responses (World Trade Center, 2001; Hurricane Charley, 2004; and the La Conchita mudslide, 2005) and nine high fidelity field exercises. The World Trade Center disaster was the first use of rescue robots and produced two analyses; one of the overall robot performance (Micire, 2002) and another of the human-robot interaction challenges (Casper, 2002). Ground robots were deployed to Hurricane Charley but not used due to the differences between a concentrated multi-story building collapse and a geographically distributed disaster with few unstructured voids. The Hurricane Charley response illustrated the need for other robot modalities (air, water) and issues in network communications and information technology (Murphy & Stover, 2005) Advances in robot platforms and operator interfaces were evaluated in two crushed houses at the La Conchita, California, Mudslide (Murphy & Stover, 2006) Prior to WTC disaster, CRASAR had conducted a field study in July 2001 with Florida Task Force 3 in a partially demolished building (Casper & Murphy, 2002) Following the WTC disaster, CRASAR conducted studies with emergency responders during US&R training exercises in a collapsed warehouse in Miami, Florida (2001) (Burke, Murphy, Coovert, & Riddle, 2004), a demolished housing project in Bridgeport, CT (2002) (Burke & Murphy, 2004) and dormitory in Oklahoma City, OK (2003), and extensive rubble piles at NASA Ames Research Center (2004) and New Jersey Task Force 1 (2005). While those six studies focused on search activities, three additional studies collected data on remote medical assistance: participation in the ShadowBowl (2003) (Gage, Murphy, Rasmussen, & Minten, 2004), a survey of emergency medical personnel (Murphy, Riddle, & Rasmussen, 2004), and a controlled study of robot-assisted medical reachback (Riddle, Murphy, & Burke, 2005)

In terms of conducting HRI studies of robotic field applications under realistic conditions and actual operators, we are aware of only the work of Jones and Hinds (2002) with SWAT teams beside our own. Scholtz and Yanco (Scholtz, Young, Drury, & Yanco, 2004; Yanco, Drury, & Scholtz, 2004) have conducted numerous studies in US&R during the RoboCup Rescue and AAAI Mobile Robot competitions. These studies were conducted with students operating research prototypes in a highly simplified indoor test course without the constraints of the standard operating procedures and command hierarchy and focused on human-computer interfaces.

LESSONS LEARNED

This paper distills our collective experience into four lessons and is intended to introduce the HFES community to the major challenges for HRI in rescue robots. We expect that these challenges will generalize to any robot system whose primary function is to provide humans with a *remote presence* (e.g., military, bomb squad robots) versus robots being used in direct contact with humans (e.g., rehabilitation, etc.) or working around humans (e.g., robot vacuum cleaners).

Lesson 1: Situation Awareness is the Major Bottleneck

Situation awareness (SA), not autonomous navigation, is the major bottleneck in robot autonomy. This lesson was extracted from several studies, beginning with a fundamental analysis of the tasks. The current state of the practice in rescue robotics is the use of small teleoperated robots due to their man-packability, simplistic interfaces and control units (Murphy, 2004). The goal of most robotics researchers to create new robots that can function so autonomously that a single operator can control multiple robots (SOMR). In order to design such levels of autonomy, CRASAR examined data from the WTC and conducted domain studies to provide an understanding of the task and where decisions are being made (Casper & Murphy, 2003)

The analysis produced the surprising observation that the mission portion of the task (i.e., searching for a victim in visual clutter, assessing structural conditions) was far beyond the capabilities of autonomous perception. This contradicted the assumption that autonomous navigation in rough, unstructured terrains is the "hardest part," especially given advances in autonomous navigation for large robots that appear to be waiting only for miniaturization of sensors in order to be applicable. The domain studies indicated that humans would be directly involved with the robots, if only to supplement the lack of perceptual autonomy.

The observation that humans would have to be involved led us to examine how humans and robots form situation awareness, in order to produce decision-aids and the appropriate level of semi-autonomy. These studies, described below, show that robot operators have difficulty building and maintaining the lowest levels of SA and that they can compensate for this by communicating with a partner. Furthermore, the studies confirmed that the mission portion of the task takes almost as long as the navigation portion; so even if the navigation portion is fully autonomous, the human could not be eliminated.

Burke et al (2004) investigated human–robot interaction during a 16-hr, high-fidelity urban search and rescue disaster response drill with teleoperated robots. In this field study, operator situation awareness and technical search team interaction were examined using communication analysis. We analyzed situation awareness, team communication, and the interaction of these constructs using a systematic coding scheme designed for this research. The findings indicate that operators spent significantly more time gathering information about the state of the robot and the state of the environment than they did navigating the robot. Operators had difficulty integrating the robot's view into their understanding of the search and rescue site. They compensated for this lack of situation awareness by communicating with team members at the site, attempting to gather information that would provide a more complete mental model of the site. They also worked with team members to develop search strategies.

These results were replicated and extended in a subsequent field study, where over 60% of robot operator communications were related to building and maintaining SA, while only 28% of robot operator communications pertained to activities using SA. The study was conducted in 2002 in Bridgeport, Connecticut, where 28 multi-operator single robot (MOSR) teams were videotaped as they teleoperated a rescue robot through an apartment in a collapsed building in search of a victim mannequin (Burke & Murphy, 2004). Team communication analysis was conducted using the Robot-Assisted Search and Rescue Coding Scheme (RASAR-CS), in conjunction with the administration of three instruments measuring situation awareness and task performance. The major findings were that 1) rescuer teams with high SA operators were 9 times as likely to find victims than rescue teams with low SA operators, and scored 26% higher on ratings of task performance; 2) team communication related to the tasks, system and environment helps the robot operator create a better situation model, and improves the quality of the human-robot interaction taking place.

Our work shows that two humans working together are nine times more likely to find a victim. The effectiveness of two humans per robot does not appear to be due to the challenges of teleoperating the robot for navigation, but rather the difficulty in building and maintaining situation awareness. This is consistent with the prior literature. Known issues in teleoperation include time delay, sensing and display difficulties, (diminished depth perception, camera viewpoint, lack of proprioceptive feedback), communication bandwidth, operation safety & errors, and operator training ((Draper & Blair, 1996; Draper, Handel, Hood, & Kring, 1991; Massimino & Sheridan, 1994; Sheridan, 1993). All of these issues create challenges in building and maintaining situation awareness. Indeed, the robot is only being actively teleoperated 51% of the time, the remainder of the time the robot is stationary as the humans try to determine what they are looking at.

The results also suggest that major advances are needed in sensors and sensor interpretation to facilitate lower level SA activities so that the operator will more rapidly have a higher level SA; and that ways (e.g., training, software agents) must be found to facilitate appropriate communication to support productive team processes.

Figure 1. Rescue teams with high SA operators were 9 times as likely to successfully locate the victim. (Operator pair in white, CRASAR observers in blue, robot underground.)

Lesson 2: It’s the Information, Not the Robot

Human-robot interaction should not be thought of in terms of how to control the robot, but rather as how a team of experts can exploit the robot as an active information source. Search and rescue is an information process which involves both a hierarchy (search team, task force, incident commander) that filters information according to task (Casper & Murphy, 2003) and distributed users (search, medical, hazmat, structural collapse, safety) which use the same data from a robot in different ways following the information flow model developed in (Murphy, 2004). This slows the response process down. In the future, given the increasing availability of communications, some of the users will not be on-site, increasing the challenge of distributed teamwork. The distributed users have competing goals, yet must work as a team.

The World Trade Center rescue response provided an unfortunate opportunity to study the human robot-interactions during a real unstaged rescue for the first time (Casper & Murphy, 2003). A post-hoc analysis was performed on the data collected during the response which resulted in seventeen findings on the impact of the environment and conditions on the skills displayed and needed by robots and humans, the details of the Urban Search and Rescue (USAR) task, the social informatics in the USAR domain, and what information is communicated at what time. Robot information was described as a one to many mapping, with temporal and abstraction hierarchies. Robots provided information to only the operators, who were expected to distribute information to the technical search leader, task force leader and incident commander. During the response, this was done manually where robot specialist teams would verbally inform the search leader of findings. Often it was 12 hours before information on victims made it to the right authority. The response further illustrated that robot information is not simply distributed by broadcasting: not all members within the task force needed the same information at the same time. For instance, a structural specialist requires structural information from the robot, while the search team leader expects victim and void information. This suggests that robot information will have to be packaged to meet the needs of the users and that contention for control of the robot between multiple “information consumers” who each view the robot as an extension of themselves must be managed.

Murphy (2004) provided a more structured view of this process with a domain theory of the search activity. The domain theory consists of two parts: (1) a workflow model identifying the major tasks, actions, and roles in robot-assisted search (e.g., a workflow model) and (2) a general information flow model of how data from the robot is fused by various team members into information and knowledge. As the author noted, the information flow is hierarchical and compartmentalized, yet could be profoundly changed by distributed communications. The information flow model also captures the types of situation awareness needed by each agent in the rescue robot system.

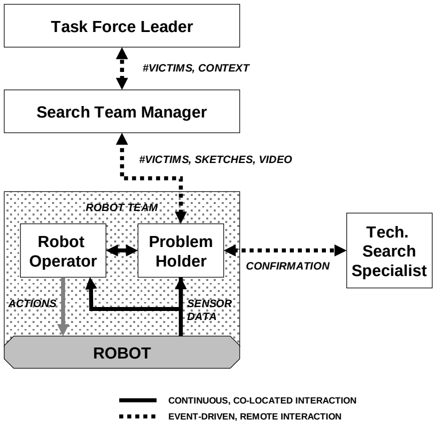

Figure 2. Information flow model of US&R from (Murphy, 2004).

Lesson 3: Shared Visual Information Builds Understanding

Previous lessons indicate that SA is the bottleneck, that teamwork is essential to building SA, and that many of the “information consumers” may be geographically distributed and have different tasks. The lesson derived from additional studies focusing on how these distributed teams can cooperate in Robot-assisted medical reachback (RAMR) show that they use shared visual information to build shared mental models and facilitate team coordination; the shared visual presence provided by the robot’s view may be the key to reducing the need for complex mental models of the robots.

RAMR involves remote medical personnel conducting operator- and robot- mediated victim assessment and triage decision making for US&R. A simulated medical reachback exercise was developed to examine RAMR (Riddle et al., 2005). The study showed that human-robot interaction in medical reachback can be decomposed into 5 distinct sequential phases, and the situation awareness requirements are phase specific. The team members used shared visual information to: build shared mental models and to facilitate team coordination; coordinate team activities through targeted and non-targeted communication; increase the efficiency of team communication; and perform the task. Videotapes of reachback exercises along with responses to interview questions were used to explore issues related to how providers perform victim assessment via the robot and operator.

Key findings suggest 1) it is critical for providers and operators to maintain a shared visual space – this is central for developing mental models and facilitating team coordination, and 2) the role of problem holder dynamically changes through the scenario suggesting benefits to adding an information facilitator role. Communication analysis revealed that across the RAMR task shared visual information was used approximately 50% of the time to facilitate the development of shared mental models (statements reflected the team members understanding of the task and environment), and again approximately 50% of the time to facilitate team coordination activities (statements directing or initiating navigation).

Changes in coordination as the scenario progresses from an initial orient phase to victim assessment indicate a shift in problem holder. During orientation, operators initiated 32% of coordination activities with providers initiating 68%. However, in the victim assessment phase, provider initiation increases to 81%. Even under the controlled laboratory conditions of the reachback exercise, the distributed team members made errors. Our experience is that even more errors will be made in the field given the cognitive and emotional stress and fatigue associated with a real incident.

Lesson 4: Social Interaction Occurs Where You Least Expect It

Social interaction between a human and a robot is often ignored, as researchers concentrate on fundamental HCI issues. However, rescuers working in close proximity to the non-anthropomorphic robots respond to the robot socially as seen in the study in Fincannon, Barnes, Murphy, & Riddle (2004). This study investigates data collected from operating an Inuktun robot in an Urban Search and Rescue (USAR) confined space training exercise task at Virginia Beach Training Center. The robot was being used by the safety officer and the task force leader to assess progress by a team breaching debris in a highly confined space. Data was collected from the coding of a video 57 minutes and 10 seconds in length. The video had no sound so all analysis is based on the video feed. Indicators of communication, gestures, physical interaction with the robot, and robot movements were analyzed. The findings indicate that the robot emerges as virtual presence for the support of the team outside of the confined space and that the team members respond socially to the robot. Rescuers working along side a robot in two different venues both maintained eye to robot contact and observed "personal space" etiquette. This confirms that research showing that people project human social behavior onto computers (Nass & Moon, 2000) transfers to humans and robots. This unexpected social interaction with robots which are not anthropomorphic or intended to be affective suggests this will be an important area of research.

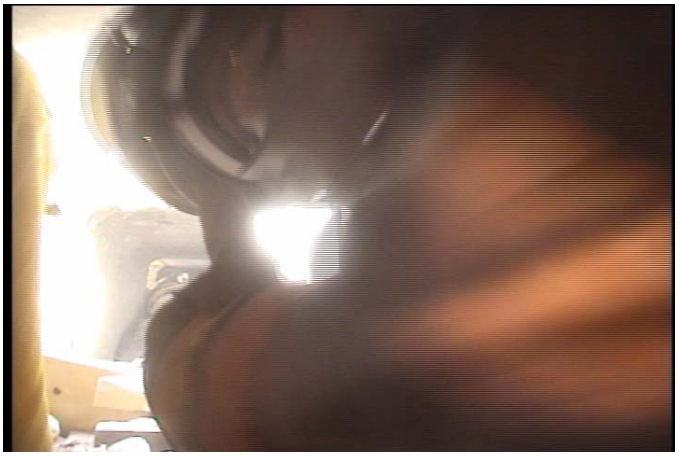

Figure 3. A rescuer turns and gestures to the robot to come forward as he speaks.

Additional Observations

In addition to our studies, we have compiled a set of "folk wisdom" about the use of robots and how that impacts human-robot interaction. To begin with, is our experience that every response is totally different and causes unforeseen problems or opportunities. We have never gone to an actual response and used the equipment the way we thought we would. At the World Trade Center disaster, robots thought to be too small to be viable were used in 7 out of the 8 runs during the rescue phase (Casper & Murphy, 2003). At Hurricane Charley, the ground robots were useless though there was a clear need for unmanned aerial and underwater vehicles (Murphy & Stover, 2005) The La Conchita mudslide response illustrated the need for diagnostic interfaces (Murphy & Stover, 2006) Furthermore, rescue robotics applications are continually emerging; as seen in Fincannon et al (2004) even at the most casual of field exercises has fostered new uses. Therefore it is hard to construct accurate, detailed models of the tasks and interactions. We speculate that focusing on detailed models may even be counterproductive, because it distracts from the real issues. Instead, we believe that the emphasis should be 1) on modeling how people successfully adapt their decision-making and team processes and 2) how to rapidly incorporate new technology into an event and capture novel uses.

SUMMARY

In this paper we have attempted to distill from the results of numerous field studies 4 important lessons learned in human-robot interaction. First, situation awareness is the major bottle neck in robot autonomy, not autonomous navigation. Results from field studies with 33 operators show that the robot is stationary half the time, as operators try to understand what is going on around them by communicating about the task, system and environment. Second, human-robot interaction should not be thought of in terms of how to control the robot, but rather as how a team of experts can exploit the robot as an active information source. Search and rescue is an information process which involves both a hierarchy that filters information according to task and distributed users which use the same data from a robot in different ways (a process, we think, that will generalize to other robot-assisted tasks). Third, the shared visual presence provided by the robot's view may be the key to reducing the need for complex mental models of the robots. Team members use shared visual information to build shared mental models and facilitate team coordination; the use of mobile robots as a shared visual presence is posited as a means of addressing the communication and coordination challenges for distributed team members both onsite during an incident response, and also for others removed from the incident site. Lastly, victims and rescuers in close proximity to the non-anthropormorphic robots respond to the robot socially as seen in the study in (Fincannon, 2004). Teleoperation applications have traditionally been where the robot was acting in areas where there were no people; here, robots are operating side by side with humans, which could change the nature of the task and the demands on the robot operator. Training for robot operators may have to address expressing social interaction cues as well as executing the task.

ACKNOWLEDGEMENTS

These efforts were supported in part by grants from the DARPA CSEE program, NSF CISE-RR program, and SAIC.

REFERENCES

Burke, J. L., & Murphy, R. R. (2004). *Situation Awareness and Task Performance in Robot-Assisted Technical Search: Bujold Goes to Bridgeport* (No. CRASAR-TR2004-23). Tampa, FL: Center for Robot-Assisted Search and Rescue.

Burke, J. L., Murphy, R. R., Coovert, M. D., & Riddle, D. L. (2004). Moonlight in Miami: A field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *Human-Computer Interaction, 19*(1-2), 85-116.

Burke, J. L., Murphy, R. R., Rogers, E., Lumelsky, V. J., & Scholtz, J. (2004). Final report for the DARPA/NSF interdisciplinary study on human-robot interaction. *Ieee Transactions on Systems Man and Cybernetics Part C-Applications and Reviews, 34*(2), 103-112.

Casper, J. (2002). *Human-robot interactions during the robot-assisted Urban Search and Rescue reponse at the World Trade Center.* Unpublished Master's, University of South Florida, Tampa, FL.

Casper, J., & Murphy, R. R. (2003). Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center. *Ieee Transactions on Systems Man and Cybernetics Part B-Cybernetics, 33*(3), 367-385.

Casper, J. L., & Murphy, R. R. (2002). *Workflow study on human-robot interaction in USAR.* Paper presented at the IEEE International Conference on Robotics and Automation.

Draper, J. V., & Blair, L. M. (1996). Workload, flow, and telepresence during teleoperation. *Proceedings. 1996 IEEE International Conference on Robotics and Automation (Cat. No.96CH35857)*, 1030-1035 Volume 1032.

Draper, J. V., Handel, S., Hood, C. C., & Kring, C. T. (1991). *Three experiments with stereoscopic television: when it works and why.* Paper presented at the Systems, Man, and Cybernetics, 1991. 'Decision Aiding for Complex Systems, Conference Proceedings., 1991 IEEE International Conference on.

Fincannon, T., Barnes, L., Murphy, R. R., & Riddle, D. L. (2004, September 28 - October 2, 2004). *Talking and Gesturing to a Robot: Emergent Social Interaction in Rescue Robots.* Paper presented at the IEEE International Conference on Intelligent Robots and Systems (IROS), Sendai, Japan.

Gage, A., Murphy, R. R., Rasmussen, E., & Minten, B. (2004). Shadowbowl 2003: Lessons Learned from a Reachback Exercise with Rescue Robots. *Robotics & Automation Magazine, IEEE, 11*(3), 62-69.

Jones, H., & Hinds, P. (2002). *Extreme work groups: Using SWAT teams as a model for coordinating distributed robots.* Paper presented at the CSCW 2002 Conference on Computer Supported Cooperative Work, New York.

Massimino, M. J., & Sheridan, T. B. (1994). Teleoperator performance with varying force and visual feedback. *Human Factors, 36*(1), 145-157.

Micire, M. (2002). *Analysis of the Robotic-Assisted Search and Rescue Response to the World Trade Center Disaster.* Unpublished Master's, University of South Florida, Tampa, FL.

Murphy, R., Riddle, D., & Rasmussen, E. (2004). *Robot-Assisted Medical Reachback: A Survey of How Medical Personnel Expect to Interact with Rescue Robots.* Paper presented at the 13th Annual IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN 2004), Kurashiki, Okayama Japan.

Murphy, R., & Stover, S. (2006, Feb. 12-15, 2006). *Rescue Robot Performance at 2005 La Conchita Mudslides.* Paper presented at the American Nuclear Society 2006 Sharing Solutions for Emergencies and Hazardous Environments Conference, Salt Lake City, UT.

Murphy, R. R. (2004). Human-robot interaction in rescue robotics. *Ieee Transactions on Systems Man and Cybernetics Part C-Applications and Reviews, 34*(2), 138-153.

Murphy, R. R., & Stover, S. (2005, June 28-30, 2005). *Lessons Learned on the Uses of Unmanned Vehicles from the 2004 Florida Hurricane Season.* Paper presented at the AUVSI 2005 Conference on Unmanned Systems North America, Baltimore, MD.

Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues, 56*(1), 81-103.

Riddle, D. L., Murphy, R. R., & Burke, J. L. (2005, August 13-15, 2005). *Robot-Assisted Medical Reachback: Using Shared Visual Information.* Paper presented at the 14th IEEE International Workshop on Robot and Human Interactive Communication, Nashville, TN.

Scholtz, J., Young, J., Drury, J. L., & Yanco, H. A. (2004). *Evaluation of human-robot interaction awareness in search and rescue.* Paper presented at the Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on.

Sheridan, T. (1993). Space teleoperation through time delay: review and prognosis. *Ieee Transactions on Robotics and Automation, 9*(5), 592-606.

Yanco, H. A., Drury, J. L., & Scholtz, J. (2004). Beyond usability evaluation: Analysis of human-robot interaction at a major robotics competition. *Human-Computer Interaction, 19*(1-2), 117-149.