

Human-Robot Interaction in Rescue Robotics

Robin R. Murphy, *Member, IEEE*

(*Invited Paper*)

Abstract—Rescue robotics has been suggested by a recent DARPA/NSF study as an application domain for the research in human-robot integration (HRI). This article provides a short tutorial on how robots are currently used in urban search and rescue (USAR) and discusses the HRI issues encountered over the past eight years. A domain theory of the search activity is formulated. 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. The information flow model also captures the types of situation awareness needed by each agent in the rescue robot system. The article presents a synopsis of the major HRI issues in reducing the number of humans it takes to control a robot, maintaining performance with geographically distributed teams with intermittent communications, and encouraging acceptance within the existing social structure.

I. INTRODUCTION

Urban search and rescue (USAR) is the emergency response function which deals with the collapse of man-made structures. The World Trade Center (WTC) disaster [1], [2], [3] was the first known use of mobile robots for USAR. The WTC disaster demonstrated that small robots which can fit inside a backpack have a unique capability to collect useful data in USAR situations. Robots can enter voids too small or deep for a person, and can begin surveying larger voids that people are not permitted to enter until a fire has been put out or the structure has been reinforced, a process that can take over eight hours. They can carry cameras, thermal imagers, hazardous material detectors, and medical payloads into the interior of a rubble pile far beyond where a boroscope can reach. Rescue robotics has been identified by the National Research Council's study "Making the Nation Safer: The Role of Science and Technology in Countering Terrorism"[4] and the 2003 Computing Research Association's Grand Research Challenges in Information Systems [5] as a critical technology.

While USAR poses many as yet unsolved challenges in mobility, sensing, and artificial intelligence, our experience suggests that the biggest obstacles in rescue robotics stem from a limited understanding of Human-Robot Interaction (HRI). This experience spans over eight years as both as a researcher and as a technical search specialist, with studies before the World Trade Center disaster as a member of the state Florida Task Force 3 USAR response team[6], during as leader one of

four groups of roboticists operating under the direction of the Center for Robot-Assisted Search and Rescue (CRASAR) [1], [2], [3], and since then as part of an internationally recognized response support team.[7], [6], [8]

Rescue robots is a near-ideal application for studying HRI. Robots are just now being developed for search and rescue, providing an opportunity to observe the impact of HRI on technology adoption in a real life domain. The rescue enterprise involves a diverse set of people, from ordinary citizens that are victims of a disaster to highly trained rescue professionals on through robot specialists. Humans have to communicate directly with the robots, either as operators or as victims, but humans may be consumers of robot information without having any prior knowledge of how a rescue robot works or even awareness of the source of the information. The complexity of relationships and tasks in a rescue enterprise pose challenges for current methods of cognitive, task, and social modeling. Teamwork pervades the rescue enterprise with a buddy system for safety. But rescue robots particularly require teamwork since it takes two humans to operate one robot. Advances in robot autonomy and distributed network communications technologies will place even more demands on the humans team members. Not surprisingly, rescue robotics has recently been cited by the DARPA/NSF study on human-robot interaction [9] as one of only two "Grand Challenge" applications (with the AAI Grand Challenge of a robot attending a conference and delivering presentations as the second).

This article summarizes the HRI issues in rescue robotics and presents a preliminary domain theory of the visual technical search task, which is the most mature task for rescue robots at this time. It expected that this will serve both the HRI and the robotics communities as an introduction to rescue robotics and will provide a foundation for additional HRI research. The article first reviews the HRI literature on rescue robotics and other field domains, in particular space exploration, and SWAT teams, in Sec. II. HRI research in field domains has been very limited and it is clear from the review that new methodologies for modeling team processes and measuring performance are needed. Sec. III paints a broad picture of the rescue enterprise. It describes the overall ecology of a USAR application: the people that are involved, the impact of the environment on them, the robots and the cognitive and perceptual demands they place on the human, and the general rescue sequence of activities. The ecology is essential for grounding the domain theory which is presented in Sec. IV. The domain theory for search consists of two models. One is a workflow model of the tasks and actions robots currently perform as part of the search activity. The workflow model gives an understanding of what the humans and robots are doing, but does not necessarily shed insight on the cognition or team processes involved. Those are

Manuscript received March 5, 2003, revised January 14, 2004. This work was supported under grants from DARPA (CSF Seedling Program N66001-01-1-8917, IPTO N66001-03-8921), the Center for Disaster Management and Humanitarian Assistance, and SAIC.

R.Murphy is with the Center for Robot-Assisted Search and Rescue at the University of South Florida, Tampa, FL 33620; email: murphy@csee.usf.edu.

captured by the second model, a model of how information is propagated and transformed through the organizational hierarchy. Particular attention is paid to how raw image data is transformed into the situation awareness (SA) needed for robot navigation and for search. These two models will likely change with advances in robotics, artificial intelligence, and network communications; the types of changes and overall impact on HRI are discussed in Sec. V. The article concludes in Sec. VI that the most critical HRI issues presented by rescue robotics are cooperative perception and distributed team work in the presence of unreliable communications.

II. RELATED AND PREVIOUS WORK

Human-robot interaction has only been recognized as a research topic in the past five years. As such, there is only a small corpus of literature on models, applications, and methods. Since this article is focused on USAR as an application domain, this section first covers HRI research in other field applications, then presents a chronological narrative of our research in rescue robotics.

A. HRI in other Field Domains

Although human interactions with robots span almost every robotic endeavor [10], [11], HRI research has been conducted primarily in six application domains: entertainment [12], [13], [14], [15] (see Breazeal in this issue), museum docents [16], [17], [18], [19], [20], personal assistants [21], [22] (see Hüttenrauch, Green, Norman, Oestreicher, and Eklundh and Lisetti in this issue), health care [23], [24], [25], space exploration [26], [27] (see Clancey in this issue), police SWAT teams [28], military robotics [29] (see Endo, MacKenzie, and Arkin and Skubic, Perzanowski, Blisard, Schultz, Adams, Bugajska, and Brock in this issue), and rescue robotics [7], [6], [2], [30]. Of these application domains, we consider only space exploration, SWAT, military, and rescue robotics as *field applications*.

Field application domains have two relevant characteristics. First, the robots are subject to unpredictable environmental effects that possibly impair platform and perceptual capabilities. In a field application, there is the possibility that at any minute the robot will fall into a hole or a rainstorm will suddenly obscure the video camera. Office environments, in contrast, are benign. They usually have constant lighting, environmental conditions, and a smooth terrain favorable to wheeled traffic. One outcome of the severe operational environment is that field robots almost always involve some degree of teleoperation to cope with current limits on artificial intelligence.

Second, in field applications, robots are primarily extensions of humans; the robots are often intended to remove a human from harms way. Field robots primarily interact with their physically distant operator and other humans; humans are generally not co-located with the robot and are often *bystanders*. [31] The lack of humans next to the robots is in contrast to entertainment, museum docents, health care, and personal assistant applications where the primary users are humans which are co-located. This means the “human-centered aspect” is focused on the operator or those “behind”

the robot not those physically co-located with the robot, or “in front.”

The NASA Robonaut program is dedicated to building an autonomous humanoid robot to work side-by-side with astronauts, but the project is currently focusing on the mechanics and control, with no clear HRI research results available at this time. [32], [26] NASA is also sponsoring a EVA Robotic Assistant project, where a mobile robot carries and helps deploys devices. [27] The HRI motivation for the EVA Robotic Assistant project is discussed in the article by Clancey in this issue. In the EVA case, the robot is automating an existing activity or task, whereas in rescue, robots are being used for novel tasks previously impossible for humans.

The domain of SWAT teams identified by Jones et al, is the nearest to rescue robots. [33], [28] In particular, Jones and Hinds observed police SWAT teams in training exercises, then created a Correspondence Agent was created to assist the operator in building global awareness, and to send commands to distributed robots using their own frame of reference. The time-criticality of SWAT is similar to USAR as well as the novel introduction of robotics.

Military applications of unmanned aerial, ground, underwater, and surface vehicles have traditionally focused on the development of the platform, with less emphasis on HRI. A notable exception is the unmanned aerial vehicle domain, which draws on a history of cockpit automation. [34], [35] A call for expedited research in human-robot interaction for US Army unmanned ground vehicles was made in the National Research Council’s report, Technology Development for Army Unmanned Ground Vehicles. [4] One of the most longest running projects addressing HRI in ground robotics has been Arkin’s Mission Lab intelligent interface for specifying missions for robot teams. [29] A formal study of Mission Lab appears in this issue.

The theoretical aspects of rescue robotics are being explored by many researchers through the RoboCup and AAAI rescue robot competitions (see [36], [37], [38], [39]). These competitions rely on the NIST standard reference course for urban search and rescue which is a highly simplified indoor arena, and stresses the relationship of the robot to the operator(s). [40] As such they provide a useful preliminary introduction to USAR, but, as argued in [41], are not representative of the challenges to platform agility or sensing, nor the larger HRI relationships. However, the value of the competition for HRI research should not be dismissed, as seen by the analysis by Drury, Yanco and Scholtz identifying poor HRI as the source for robot collisions in the most recent event. [42] Research is being in non-HRI aspects of USAR, such as platforms (see [43]) or multi-agency (see [44], [45]), though generally in simulation.

B. Rescue Robotics at CRASAR

The core of work in HRI specifically for rescue robotics has been under the direction of the author, starting with research in the aftermath of the Oklahoma City (OKC) bombing in 1995. Robots were not used at the OKC response, but a graduate student, John Blitch, participated and took notes as

to how robots might have been applied. The initial post-OKC efforts took two directions: the development of an expert system for determining which existing robots are useful for what situations (see [46]) and exploration of marsupial (mother-daughter) class of robots ([47]) as a solution to several platform deficiencies. The work by Blitch in [46] provided some of the motivation for the DARPA Tactical Mobile Robot (TMR) program [48], which produced prototype small robots for military operations in urban terrains. Robots and personnel for WTC response were drawn from the TMR program.

In 1998, work turned to direct field studies with Hillsborough County Fire Rescue in Tampa, Florida, leading to the inclusion of the author and some grad students as members of Florida Task Force 3 (the Tampa Bay state regional USAR team). The work produced both assessments of platform needs ([49], [50]) and a workflow study of search[6], which contributed to the model presented in Sec. IV. The study established methods of minimally disruptive data collection and identified tasks within the visual technical search task for a downed fire fighter. It concluded that two humans are needed for visual search with a teleoperated robots, because operator kept missing obvious signs of victims while a supervisor who was not controlling the robot was able to easily notice the downed fireman. This study supports the hypothesis in the article by Woods et al in this issue that there is a *robot operator* and a *problem holder* (in this case, the supervisor).

The WTC response was the next study and provided over 10 hours of videotapes of the robot's eye view in the rubble pile at the WTC. These tapes were examined and reported on in two MS thesis ([3], [1]). Casper's thesis focused solely on the human-robot interaction aspects and a summary appears in [2]. The thesis found that the robots were sensor-impooverished, requiring the human operator to use only the visual channel for all information extraction. The operators consisted of experienced robot operators who made videotapes of the robot's eye view. The tapes were then reviewed on-site by experience federal rescue professionals. All of the operators and rescuers ere already physically and cognitively fatigued, exacerbated by poor interfaces and lack of functional presence. The post-WTC analysis showed that the robots were being operated with some sort of operator error at least 18.9% of the time, and that remains of victims were missed, even by rescue professionals.

Immediately after the WTC, CRASAR formed a formal response support team. The response team participated in a highly realistic 16-hour USAR exercise for rescue professions held in collapsed structures in Miami, Florida.[7] Data was collected of rescue workers using the robots, following the procedure established in [6], but analyzed using conversational analysis with a coding scheme developed especially for the purpose (RASAR-CCS).[7] The analysis confirmed that two operators are currently needed to drive and look, primarily due to deficits in situation awareness.

The research into building situation awareness in search and rescue continues to evolve. The Miami field study has since been followed up with two, as yet unpublished, studies. One study was conducted in Dec, 2002, in Bridgeport, Connecticut, and the other in April, 2003, in Oklahoma City, Oklahoma. Both were similar to Miami but with more users and improve-

ments to the data collection and analysis methodologies.

None of the three above studies considered the impact of distributed teams and communications. A separate study as part of the 2003 ShadowBowl emergency response exercise effort did address this.[8] In this case, a robot team was pre-deployed in San Diego for the Super Bowl. During the Super Bowl, a mass casualty incident was simulated, and data from arrays of sensors and the robot were sent to medical experts throughout the US over the internet. Members of the ShadowBowl team in Tampa, Florida, attempted to view and interpret data sent from the robot in San Diego. The event identified major shortcomings in current thinking about reach back, and highlighted that the raw sensor data from the robot was insufficient for situation awareness.[8] It concluded that the data coming from the remote site should be labeled (possibly using via XML), that all data should be sent to a centralized proxy server outside of the Hot Zone where it can maintained. Besides the expected lack of situation awareness, the report noted that there were team communication breakdowns between the San Diego and Tampa teams. Humans need to make sure they are communicating not just sending data between computers. The work in distributed teams for the medical activity of search and rescue is continuing.

III. ECOLOGY

As highlighted in the DARPA/NSF study on HRI [9], traditional topics in HRI are task, environment, and social modeling. Understanding these areas establishes a basis for measuring performance of team members. Social relations-where social interaction is needed or expected, how robots can work with multiple humans, and task trading- are also targeted.

This section addresses these topics by broadly describing the *ecology* of a USAR response: the environment, the robots, the people, and the task. The ecology is needed for environmental and social modeling, and it provides a framework for understanding the specific task and cognitive models captured by domain theory of the technical search activity in Sec. IV. The section begins with an overview of the physical environment of a rubble pile, illustrating how the physics of the rubble greatly constrains robots, limits user interfaces, and imposes high cognitive demands on humans. Since rescue robots are quite different from laboratory robots, the models and capabilities of robots actually used in the field by CRASAR are discussed next. As will be seen in Sec. IV, the limited capabilities of the robots place huge demands on human operators. Next, the section covers the various people involved in search and rescue, their organization, their backgrounds, and their attitudes toward technology. This provides a foundation for creating the social models of how human-robot teams will interact and how novel technologies will be accepted. However, understanding the components of a search and rescue incident is not the same as understanding the sequence and interactions in a response. Therefore, the section concludes with an overview of a USAR response.

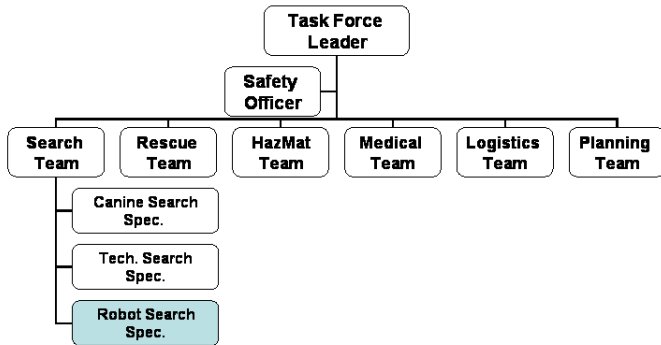


Fig. 1. The six activities of urban search and rescue with a possible insertion of robots shown in gray.

A. People

Cognitive and social modeling requires an understanding of the people involved in the rescue enterprise. People involved with the rescue system fall into two categories. The first category are the *rescue professionals*, while the second is the *survivors* themselves, who are beyond the scope of this article. For simplicity, the description of rescue professionals and their roles will follow the organization of a US Federal Emergency Response Agency (FEMA) Urban Search and Rescue Task Force. It should be noted that state task forces and rescue professionals from other countries use a different, but qualitatively equivalent, organization.

This section focuses on the rescue professionals in general, excluding extra-mural robot operators and other technology support personnel. This is because there are no robots embedded in the US federal or state task forces, with the exception of CRASAR's host team, Florida Task Force 3. As of publication time, CRASAR is the only known response support team in the world which has rescue robots and operators with USAR training.

Fig. 1 shows that a federal task force is divided into six functional teams: *Search*, *Rescue* (or *Extrication*), *Medical*, and *HazMat* (assessing and mitigating the presence of hazardous materials or conditions), *Logistics*, and *Planning* (assessing the structural condition of the site, etc.) Each team is lead by a Manager and staffed with Specialists. If robots were used for search, they would report to the Search Team Manager. Altogether, a federal task force consists of over 100 responders.

The majority of rescue professionals are either part, or have been part, of a municipal fire rescue department. Fire rescue is a profession, usually regulated by a union. The large majority are male and have varying levels of education, usually at least two years of college and advanced emergency medical training. As part of a union, firefighters may be suspicious that technology will negatively impact their job or is intended to eliminate their job. Since rescuers often do not

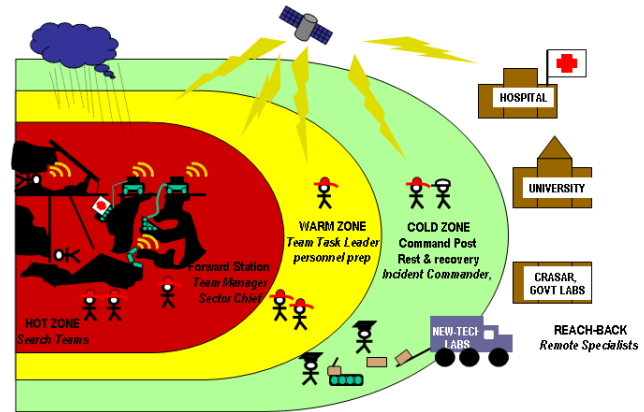


Fig. 2. The three zones within the physical and working environment.

have direct control over purchases, they may have had negative experiences with inferior technology purchased by government officials based on cost and vendor influence.

While the majority of rescue workers come from a fire department, a significant number of team members are *civilian specialists*. Often these civilians are former firefighters or have been closely associated with the profession because of working in an emergency response industry. However, civilians, with the exception of canine specialists, are required to take all the basic training with the response team and so fit within its group dynamic. While canine teams are a standard part of a search, most fire departments do not have full-time dog handlers. Instead, canine teams are generally female civilians with widely varying backgrounds. Most do not have a fire rescue background and are generally not required to take the basic rescue training. The lack of common experiences and different world views often engenders noticeable resentment. The rescue robots at the WTC response were organizationally treated as canine teams. Since the WTC, CRASAR has focused on robots being accepted as tools, not a separate function, in order to avoid any possible resentment and delays in technology acceptance.

B. Physical Environment

The USAR physical and working environment consists of three zones (see Fig. 2). The *Hot Zone* is the area of actual devastation. Access to the Hot Zone is tightly controlled and the number of people permitted inside is tightly regulated (if possible). The *Warm Zone* is the immediate surrounding area where rescue workers assemble prior to entering the Hot Zone and prepare their equipment. After they exit the Hot Zone, rescuers decontaminate themselves and their equipment (the rubble pile is contaminated with sewage from toilets, body fluids, decaying foodstuffs, etc.) in the Warm Zone. The Warm Zone is also an assembly area for an emergency evacuation. The *Cold Zone* surrounds the Warm Zone and is also restricted to emergency workers. The Cold Zone is the location of the Incident Command headquarters, the press and media liaison

area, the area for victim's families to await news in privacy, and where workers eat, rest, and repair/re-charge equipment. Remote resources may be (rarely) accessed through telecommunications, including phones or the internet, in a process called *reach-back*. The primary working environment is the Hot Zone, and within the Hot Zone the primary attribute of interest are voids, or openings into the rubble.

1) *Voids*: In urban structures, access to the interior of the rubble pile is through *voids*, or openings. Voids are important to search and rescue because they are conduits to where trapped victims may be. However, they pose technical and safety risks to human exploration and extrication activities. As a result, voids are where rescue robots offer the greatest benefits over current technologies.

While there has been no formal study of rubble characteristics for mobile robots, CRASAR divides voids into three general types. *Semi-structured* voids are entries into structures which still resemble buildings, as in Fig. 3a. These entries are usually through stairwells, doors, or exposed (or breached) openings and afford human entry. *Confined space* voids are those where a person could crawl in but would pose great risk and would require significant safety precautions. These are voids where Occupational Safety and Health Administration Agency (OSHA) regulations requires that a structural expert approve the void as being reasonably safe, and then the worker can enter only with a safety line, supplied air, and a rapid intervention team standing by to retrieve him if something happened. *Sub-human confined space* is too small for a human to physically enter. Fig. 3c. shows an example of a robot entering a sewer pipe at the WTC; this is an example of a regularly shaped void. Most sub-human confined space voids are irregular with protruding rubble and construction material, obscuring the cameras and making it difficult to navigate.

Rescue robots are most extensively used for confined and sub-human confined space voids. These voids are primarily vertical, with drops of one to two stories, and have irregularly shaped cross-sections. In general, voids are fairly rare; the surface rubble is usually dense and many confined space voids do not appear until some rubble has been removed. It should be noted that these voids may still be on fire or have depleted oxygen; in these cases, the rescuers still need to pass through them in hopes of reaching survivable areas within the interior.

2) *Hot Zone*: While voids are the primary area of interest, it is important to understand the impact of the Hot Zone environment in general on human performance. The Hot Zone or rubble pile is challenging to work in for many reasons. First, it is dangerous for humans. The rubble pile is a collapsed or damaged urban structure which must be treated as if on the verge of further collapse (e.g., a secondary collapse). The rubble contains exposed metal surfaces and protruding concrete reinforcement bars; a fall could result in severe injury or death. Dangling overhead structures (dubbed "widow makers") could collapse at any time, killing anyone below. There may be natural gas leaks or live electric lines in the rubble.

Second, the Hot Zone is physiologically demanding. All work is conducted outdoors, subjecting a worker to local weather conditions. Personal protection equipment (PPE), such



a.



b.



c.

Fig. 3. Example voids: a.) semi-structured at FLTF-3 exercise, b.) confined space at WTC response, and c.) sub-human confined space at WTC response.



Fig. 4. WTC Tower 2 as an example of the physical challenges imposed by rubble: 0.5 x 0.3 m void opening searched by robot is shown within the box.

as a hard hat, safety goggles, steel-toed and steel-shanked boots, must be worn at all times. Safety glasses restrict vision (and interfere with heads-up displays). Safety boots reduce mobility and flexibility and add to personal discomfort. It is not uncommon for a site to be very dusty, requiring respiratory protection. A respirator is not only uncomfortable to wear, but requires the wearer to exert more energy to breathe. Thus, a person tires more rapidly. There is usually an unpleasant stench along with the presence of rats and other vermin as well as due to rotting foodstuffs or bodies. It is common for responders to go the first 48-52 hours without sleep, relying only on infrequent naps of less than 3 hours.

There is often no convenient place to stand or place bulky equipment, forcing humans to carry and support heavy robot equipment. Equipment is either transported in backpacks or by ropes; rescue workers generally keep both hands free to maintain a safe, three point crawl on the rubble. As an example of how the physiological demands add up, consider the following example. Fig. 4 shows an area of WTC Tower 2 that was searched by a robot. The robot was not used to search the exterior of the rubble pile; surface victims do not require any technology to find. Instead, the robot was used to search the void created by a hollow structural member. Two members of the CRASAR rescue robot team each carried a robot in a backpack weighing approximately 70 pounds. They had to climb down a straight ladder approximately 10 meters down into a field of debris, then climb up a nearly vertically slope of sharp metal and loose debris to reach a single void. There was no flat surface to set up the robot control unit.

The Hot Zone is also cognitively demanding. Perception is extremely difficult. The interior of the rubble pile is dark and dusty. Almost everything is covered in gray dust from pulverized cement cinder blocks, sheet rock, and ceiling tiles. The rubble is disorganized and deconstructed, so there are very few apparent size cues to help estimate depth. The robots are low to the ground, posing unnatural viewing angles for a human.

Another stressor is the lack of reliable communications. Workers are spread out over a large area and communicate over radios. Unfortunately, urban structures often create radio (and wireless internet) interference, and cell phone towers may be down or saturated. As a result teams and team members often have to work highly autonomously with little feedback from others, adding to a sense of isolation and pressure.

There is a clear time pressure in USAR, as the mortality rate exponentially increases after 48 hours. The task of interpreting the robot's images for signs of survivors, state of the collapse, etc., is known to be cognitively demanding. While the task itself is cognitively fatiguing, more stress is added through the emotional demands of working on a task that is literally life and death. Missing a victim could mean death; USAR is a domain that has high consequences for error.

In addition, the activities occur at unpredictable times, adding to cognitive fatigue. Casper [1] documents common work flows at the WTC response where robotic teams would wait on stand-by for hours before being called up for a seven-minute deployment. In addition, once on the rubble pile, the rescuers may have to evacuate on a moment's notice. In one case at the WTC, a robot team was deployed on the pile seven times during one 12-hour shift and evacuated each time before reaching the intended void. The "hurry up and wait" nature of technical search contributes to cognitive fatigue and general stress.

C. Robots

No robot is currently made specifically for USAR. Models used by CRASAR are adapted from explosive ordinance disposal (EOD), military operations in urban terrain (MOUT), and ventilation duct inspection applications. Because there are many types of voids that need to be explored, there is no one size robot or mobility configuration that is appropriate for all situations. CRASAR has identified three sizes of rescue robots: *man-packable micro*, *man-packable mini*, and *man-portable*. Examples of robots in each of these size classes are shown in Fig. 5. All the robots require an *operator control unit* and are teleoperated. CRASAR will only field a robot if it has two-way audio and a color camera, is man-packable, and is reasonably water resistant for decontamination. All of the robots currently fielded are *polymorphic*: the shape of the effectors can be dynamically changed to fit the environment.

Man-packable robots are robots that can be carried in one or two backpacks. This is of critical importance as rescuers must keep their hands free while moving about in the rubble. Man-packable robots are sub-divided into *micro* and *mini* classes. Mini robots can typically enter the same spaces as a small human, while micro-robots are smaller, often on the order of the size of a shoebox, and are well-suited for sub-human confined spaces. Micro-robots are usually tethered (to reduce size) and can operate for up to 12 hours on a motorcycle battery. At this time, CRASAR deploys variants of the Inuktun micro-VGTV robot, shown in the center of Fig. 5. This robot carries a color camera and has two-way audio. The camera and audio output can be recorded on a camcorder which serves as the monitor for the second operator. Mini-robots are often wireless and operate from 2 to 10 hours,



Fig. 5. Example of robots brought to WTC response. They range from man-packable micro size (in center) to man-portable (on left).

depending on the terrain, demands, etc. The Foster-Miller Solem EOD mini-robot was used on the rubble pile at the WTC and a prototype of the iRobot Packbot was used to explore collaterally damaged buildings.

Man-portable robots such as the Foster-Miller Talon EOD robot, SPAWAR Urbot, and the iRobot ATRV general outdoor robot are robots which are light enough to be carried by one or two people and small enough to be transported in a truck or all-terrain vehicle. They can only be used on the surface of the rubble pile or in extremely favorable semi-structured voids. The Foster-Miller Talon was lowered into stairwells at the WTC to sample air quality and to conduct a structural examination of the basement. Robots have also been shown to be able to drag victims to the Warm Zone, reducing the physical fatigue of rescuers.

D. Overview of USAR Response

The organization chart in Fig. 1 does not represent the sequence of activities in a USAR response, which is important for understanding human-robot interactions. The particulars of an incident are highly dependent on the situation and the judgment of the Incident Commander. One possible sequence is presented here to aid the reader. ESF9 activities are carried out by the local fire department's USAR team (if any). If the incident exceeds the local department's ability to respond to it, it is declared a disaster and aid is often requested from a state, federal, or international task force.

When an incident occurs, the first responders assume control of the site and establish the Hot, Warm, and Cold Zones (Fig. 2). It is reasonable to assume that the local fire department's search and rescue team will be alone for the first four hours, and will be the prime agents in establishing control over the site and assessing the situation. The first hurdle for rapid entry of teams is that humans and dogs cannot enter the Hot Zone until it is deemed reasonably safe. In order to declare an area safe enough for USAR workers, the Incident Commander must establish control of the site (confirm that natural gas

and electricity is cut off, remove civilians volunteers and looters, provide protection from a secondary terrorist attack or an aftershock) and assess the site for additional safety risks. USAR teams often are not permitted to enter an area with active hazards, such as fires started by natural gas leaks. This may introduce further delays. The need for structural and hazard materials inspection prior to entry often create the bottleneck for this phase, as there are never enough specialists to safely survey and understand the extent of the damage. This reconnaissance and survey activity is generally done prior to the arrival of the task force. The larger the incident, the longer it takes to establish control. Indeed, it may be 10 hours from the event before USAR teams can be deployed on the rubble or inside buildings.

After the assessment, the Incident Commander and staff may divide the Hot Zones into areas to be handled by individual task forces, depending on the size of the incident. The Task Force Leader will then deploy the individual teams. The first activity is to search the Hot Zone. This is done by the Search Team, often incorporating specialists from the Planning and HazMat Teams to ensure personal safety. Areas and buildings are divided into sectors and findings are localized topologically. GPS is used, but is often unreliable in "urban canyons" created by clusters of office buildings and does not work within buildings and within rubble.

If survivors are found, the Task Force Leader would work with the Planning, Medical, and Logistics Teams to create a plan. The Rescue team would be dispatched to extricate the victim, based on the resources, the probability of successful extrication, the state of the survivor, and optimal use of resources. Lack of access to deeply trapped victims typically forces the Medical team to remain on standby until the final stages of extrication. FEMA statistics show that it takes a rescue team of ten members between 4 to 10 hours to extract a single survivor, depending on how deeply trapped the victim is.[51]

After finding a survivor, the Search team is likely to continue searching, so it is possible to have teams from all branches in the field simultaneously, particularly in the first few days of a mass-casualty incident. They may not necessarily work in close proximity, both to maximize coverage and to reduce risk, since extrication may cause a secondary collapse in adjacent areas of the rubble.

Teams work in 12-hour shifts and the work is often interrupted. At the Oklahoma City bombing response everyone left the Hot Zone every 6 hours to permit the deployment of sensitive acoustic sensors and to allow the safety officer to confirm that the areas of effort in the Hot Zone was still reasonably safe. This can cause some duress to a trapped survivor, who may be left alone for over an hour. (Note that a robot would not be subject to the evacuation and could remain as a "rescue buddy" for a trapped victim.) Intermittent emergency evacuations due to sudden risks are frequent as well. At the end of the 12-hour shift, team managers and task force leaders have a scheduled meeting with the Incident Command staff and the next shift.

The first three days to one week of a response are focused on searching for survivors, which is called the *rescue phase*. The

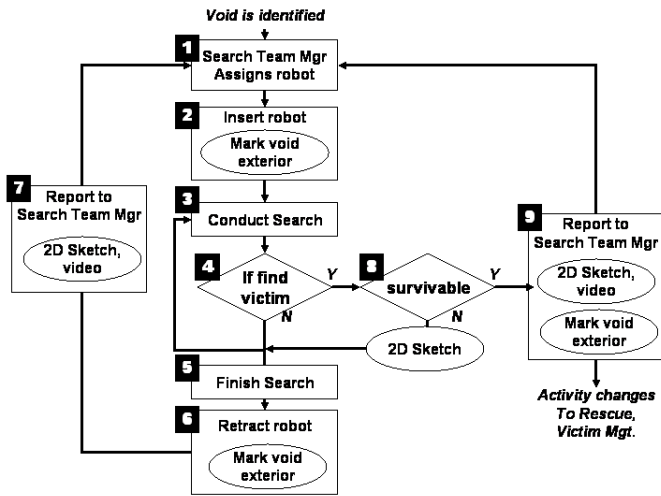


Fig. 6. Work flow model of the technical search activity in urban search and rescue.

intent is to search all the Hot Zone within 48 hours to find all survivors and then remove them before they die. Depending on the situation, the response will change from a rescue (saving lives) focus to a more methodical *recovery phase* (recovering bodies). Time is no longer an issue in the recovery phase, as the likelihood of survivors is considered negligible.

IV. DOMAIN THEORY

A domain theory is needed in order to understand the HRI issues associated for an activity. Based on our studies and experience, this section presents a domain theory of robot-assisted search. The domain theory is composed into two models. In order to understand the search activity itself, a Work Flow Model of the search activity is presented, which concentrates on the jobs and roles of each member of the robot team. The specific robot team tasks and actions within the Work Flow Model are discussed in detail. In order to understand the relationships between all the members of the search activity, a model of the Information Flow between these agents is presented. This Information Flow Model also attempts to characterize the content of situation awareness at each step in the information hierarchy. The domain theory is for the current state of human-robot interaction in robot-assisted search and rescue. Future advances in robotics and network communications may change the domain in the near future. These advances and possible ramifications are also discussed in this section.

A. Work Flow Model of Search Activity

Figure 6 shows the robot-assisted search activity. Rectangles represent steps in the activity, ovals represent information expressed for communication to others, and diamonds represent alternatives or decision-making points in the process. Components of the model are number 1-9 to facilitate discussion below.

Starting at the top of Fig. 6, when a void is identified during reconnaissance or appears during excavation, the Search Team

Manager decides which resource will be used to search it. The choice of a robot over another resource depends on the situation. The Search Team Manager may consult with robot specialists to select the best robot type for the expected conditions, may rely on their own knowledge, or be restricted by robot availability. The personal risk to the search team may be a factor as well. Since the robot team is likely to be composed of civilian robot specialists, they may not be sufficiently trained and equipped to work in a high-risk environment. In that case, the Search Team Manager may decide not to deploy the robot or to deploy the robot with a less robot-experienced team of rescue professionals. The robot team will most likely be accompanied by an experienced technical search specialist who would act as an escort, safety officer, and local decision-maker. The technical search specialist or other designated person might be responsible for several robot teams or other assets working voids in the same area, so it cannot be assumed that the robot team will have constant access to an expert.

The basic sequence is the linear flow from the 1 SEARCH TEAM MGR ASSIGNS ROBOT step to the 6 RETRACT ROBOT step. Upon arrival at the targeted void, the team marks the exterior of the void with the symbology for a search in progress (oval) and inserts the robot. The search continues until the robot encounters a victim (step 4, diamond), which triggers the perceptual decision-making sequence (steps 8-9), or cannot progress farther, at which point the search is complete (step 5). When the robot has explored the space, it is retracted (step 6) and the exterior of the void is marked with additional symbology for a completed search. A report is then made to the Search Team Manager verbally, then any sketches and video would be logged when the robot team returned to the base of operations (step 7).

1) *Representations of Knowledge*: As seen by the number of ovals in Fig. 6, an integral part of the search activity is transforming data into knowledge and expressing that knowledge as information for use by other members of the rescue enterprise. Information is expressed as either

- external markings on the void or
- as verbal or written reports.

Markings on the exterior of the void contain the minimal set of information needed for other teams if they enter the area. If another search team arrives at the void by navigational error or a faulty plan, they will be able to immediately determine that the void has been searched, to what degree, and when. If a rescue team arrives and no search team is present, perhaps due to an evacuation, they will be able to determine that they are at a void with survivors and have some awareness of structural hazards. Note that the markings do not convey map information or what tool was used (human, search-cam, audio, robot). Because USAR workers have uniform training, they have a shared model of how the search was conducted. The search pattern, a right wall-following algorithm if the void has branches, is standardized across tools. Therefore, rescuers can reproduce the search if the original search team cannot be contacted or the sketches are unavailable. The markings and the common training provide a measure of redundancy and reliability in the USAR operation. It also means that rescue

robots must “play by the rules” and not violate this shared task model.

The second form of search information is what is passed on through the ESF9 decision-making hierarchy as reports. This information is either verbal (if uninteresting) or visual (sketches and video or pictures.) Note that the most important information is visual and is unique. Given the deconstructed nature of a USAR incident and that each collapse is different in its own way, is it unlikely that visual information can be effectively reduced to verbiage. The adage “a picture is worth 1,000 words” is an understatement in USAR.

2) *Robot Search Tasks and Actions*: The Work Flow Model flowchart in Fig. 6 provides a broad overview of the search activity. This section details the specific robot tasks and actions associated with the primary task: CONDUCT SEARCH (step 3). Although technical search with robots is still in its infancy, a preliminary search strategy has been developed by CRASAR and has been taught to over 100 rescue workers. Based on the behavior observed in human investigation of voids by FL-TF-3 [6] and in conjunction with training with the USMC Chemical Biological Incident Response Force, the CRASAR robotic search strategy follows a sequence of four tasks with the mnemonic *LOVR: Localize, Observe general surroundings, look specifically for Victims, Report*.

Using this strategy, the operator directs the robot in the void in an estimated direction and distance (for example, move it forward one meter). In confined space voids, there is often no choice of directions, but if there is a choice, a right-wall following algorithm is used, consistent with standard human search practice. The two humans then observe the new set of visible landmarks from the new position and survivors. This *observation* aids with *localization* and detection of any noteworthy feature of the void such as unstable ceilings. The robot then is directed to repeat the observation of the area in all directions with an infrared thermal camera or other modalities in order to find victims. The task script ends with the primary robot operator making a verbal *report* to his/her partner who is also responsible for sketching the path of the robot and layout. This serves to reduce operator disorientation and ensure that the search is complete.

LOVR is repeated every 2 to 5 minutes, based on the complexity of the void. The robot may be directed further between stops in a straight, visually uniform area but may travel much shorter distances in more perceptually demanding situations.

Within these four tasks, there are many actions. For the purposes of this article, actions will be divided into two categories: those which are currently performed by the primary *robot operator*, and those performed by the secondary robot operator and observer, or the *problem holder*.

The actions performed by the *robot operator* are navigation-related and depend on the terrain and the type of robot. In order to explore the void, the robot must navigate forward, climbing over or around obstacles as necessary and without falling into a hole or turning over. The actions may involve changing the shape of the robot. Shape-shifting, or polymorphic, robots have significant mobility advantages, but are particularly difficult to control. Observing and finding victims, which are nominally

cognitive actions, require the robot to conduct a series of motions for surveying an area (look up, look down, look left, look right). These motions depend on the robot. For example, some robots have a tilt mechanism, while some, such as the iRobot Packbot and Inuktun microTracks, can turn in place, and others, such as the Inuktun micro-VGTV, have to move forward in order to turn. The differences in platforms present the opportunity for mistakes as operators apply motions suitable for one robot but not the one being used.

Problem holder actions are primarily recognition- and memory-related. The problem holder is expected to detect victims, and maintain an understanding of the relevant state of the world as opposed to the state of the robot. Detection is largely based on affordances or cues: heat is a cue of a victim (or ignition source), color (that a live victim has moved, thereby knocking off some of the gray dust typically covering a void), motion, sound, and non-random areas and textures (gestalt or perceptual organization). The problem holder also serves as the memory for the system: making maps, maintaining situational awareness, controlling viewpoints, etc. It should be emphasized that that path planning is not generally used for confined spaces because there are few directional choices and a right wall-following algorithm is the standard. Instead, path planning is often helping the operator remember navigational hazards on the way back (“watch out for the low hanging wires, it might snag if you back up”).

3) *Decision-making*: As shown in the diamonds for steps 4-8-9 on the flowchart in Fig. 6, the real decision-making occurs when a victim is found. If a possible victim is found, then it must be determined whether the victim is either survivable (alive) or not (remains), shown in step 8. Live victims are often easy to detect but remains may be difficult to distinguish from an unconscious survivor. Therefore, unless it is clear that there is a live victim, the team will get verification from the technical search specialist. If remains are confirmed, a pen and paper sketch of the location of the remains is made, but the search then continues.

These decisions share two attributes. First, the need for a decision is perceptually triggered (e.g., perceive signs of a victim). Second, the result is whether to gather more perceptual data (e.g., signs of life) and, if not, how to transform the perception into information that will be communicated to other members of the rescue enterprise (e.g., marking and a sketch). The decisions are not about the robot, they are about perceptual processing. The key cognitive decision for the human to make is whether the data presented is sufficient to determine whether a victim is present and survivable, not which way to move the robot. As a result of the decision on evidential sufficiency, the robot team may call in another expert to reduce the uncertainty or move the robot.

If the team encounters a survivable victim, the search activity is disrupted and dynamically changes to rescue and victim management activities (step 9). A verbal report is made immediately over the radio to the search team manager and the robot stays with the survivor. This effectively moves the robot team from the search team activity to a victim management activity, which is beyond the scope of this paper. The exterior

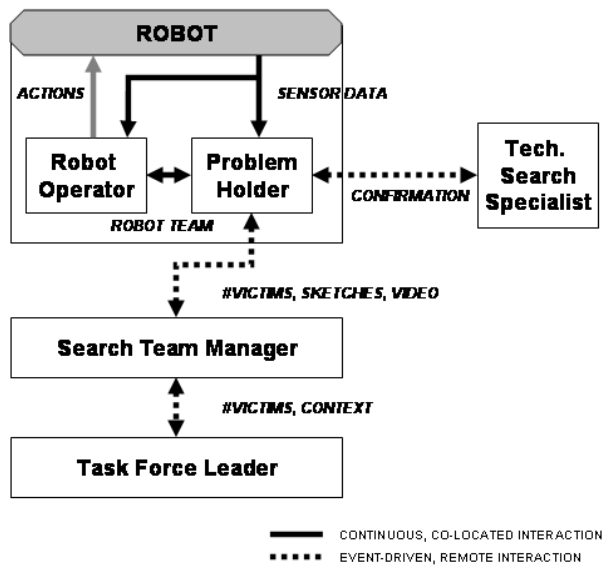


Fig. 7. Model of the transformation of data to knowledge in the search activity.

marking is updated to reflect the presence of a survivor in case the robot team has to be evacuated. In the meantime, a pen and paper sketch of the location of the survivor in the rubble and the surroundings is made for immediate distribution. The sketches may have annotations about structural and other hazards opportunistically observed.

Though not shown in the flowchart, the Search Team Manager makes a report to the Task Force Leader and a pre-plan for extrication and medical support activities is begun. In practice, the Search Team Manager and Task Force Leader would most likely physically go to the void and examine the robot information, playing back the video data to confirm the sketches and extracting more information.

B. Model of Information Flow in Search Activity

The Work Flow Model captures when information is communicated (the ovals in Fig. 6), Information Flow Model for Search in Figure 7 captures what information is being communicated and why. The model shows the propagation of data and its transformation, through the immediate decision-making hierarchy in the Search Activity, starting with the robot data at the top and progressing down to the Task Force Leader. The black lines represent the flow of information, with dotted lines indicating remote communication demands (e.g., having to contact the Search Team Manager over the radio) as opposed to direct communication between two members standing next to each other.

The source of data is the robot team. The robot supplies raw sensor data which is processed into knowledge about the robot's state for use by the robot operator and knowledge about information collected for use by the problem holder. The processing is cooperative, as shown by [7], and the human team members interact continuously. The robot operator uses the data locally to direct navigational actions for the robot.

The problem holder uses the data to detect victims and to help the robot operator navigate. The primary output of the robot team is whether there is a survivable victim, and if so, what is the estimated location and status of the victim, and state of the surroundings. The problem holder is responsible for interpreting the data, gathering the knowledge about the search outcome, expressing it as consumable information in the form of sketches and video clips, and transferring it to the Search Team Manager.

Information communications from the robot team through the layers of the hierarchy is event-driven. The two events are when 1) a victim is detected and 2) a void has been completely searched. The information from the robot team is transformed by the Search Team Manager and Task Force Leader into knowledge about the overall rescue enterprise.

The scope of the robot team is a particular void, while the scope of the Search Team Manager and the Task Force Leader is the entire Hot Zone (or assigned sector). The Search Team Manager fuses information from all searchers and packages it for consumption by the Task Force Leader. The Task Force Leader is the primary decision-maker, and this level is where the information originating with the robot team has the biggest impact. The Task Force Leader uses the robot information, combined with knowledge about resources, to adjust the overall operations. In the case of finding a survivor, the bulk of the visual information presented to the Search Team Manager will be eventually passed on to the Task Force Leader and the rescue and planning teams for the victim extrication.

Although we have not studied the propagation of information in the hierarchy beyond the robot team, Fig. 7 highlights that the information flow is generally one-way, flowing up from robot data to increasing levels of abstraction for the decision-makers in the hierarchy. There does not appear to be a great deal of communication from higher levels in the hierarchy to lower levels. Instead of information flow, a better metaphor might be information packaging and parcel routing. This one-way flow minimizes communication and unnecessary distribution (although radio communications are broadcast), eliminating communication except when a survivable victim has been discovered or a void has been searched and team is available again. Information is truly propagated on a "need-to-know" basis. This is simple, reliable, and reduces information overloading and distractions to decision-makers who are under severe physiological and cognitive fatigue.

Fig. 7 suggests that the structure of information propagation does not optimally support information flow based on associated priority and value. While the one-way flow reduces distractions, it slows down the propagation of information. If a survivable victim is found, it is literally a matter of life and death to push that information through the layers decision-making hierarchy quickly. Opportunities to exploit less imperative imperative information, such as the internal structure of the rubble pile, are not facilitated.

1) *Raw Sensor Data*: At this time, fieldable rescue robots, like most military robots, are teleoperated. Thus, they supply only raw sensor data and are not a cognitive member of the rescue team. The raw data can be categorized into:

- the robot’s internal state (pose, health),
- the robot’s relationship to the environment (going down a slope, stuck on a rock),
- the layout of the environment (range to obstacles, dimensions, general hazards, topology), and
- presence of victims.

Most robots provide internal state information through on-board sensors. The robot’s relationship to the environment is limited by sensor technology and costs. Most fieldable robots now carry tilt sensors so that it is possible to determine that the robot is precariously positioned, but most do not provide any “feeling” that they are stuck or bumping against something out of view. The robot’s relationship to environment, the layout of the environment, and the presence of victims is usually extracted from a single video camera. This puts a large burden on the humans to navigate, make maps, and interpret the scene from a video sequence. If available, a thermal camera may be also used for victim detection.

2) *Transformation into Team Situation Awareness*: We view situation awareness (SA) as the immediately relevant knowledge associated with a particular robot-assisted search. Specifically, the team interprets the raw sensor data in order to gain situation awareness

- about the environment and robot in order to enable safe and complete navigation, and
- about the contents of the void and the state of the search in order to communicate the findings to other member of the rescue enterprise.

The situation awareness being developed is primarily about spatial relationships between objects and how that impacts robot navigation and coverage of the void.

In order to discuss SA and how it is developed, it is helpful to review what intrinsic knowledge is provided by the humans to supplement the raw data provided via robot. The robot operator brings to the team an *a priori* understanding of how the robot is controlled, what it can do, and how to diagnose and recover from failures. The operator can also predict the outcome of actions on the environment and data gathered (e.g., that turning up headlights to improve the view will be compensated for by auto-gain). This *a priori* knowledge is specific to robot models, so a robot operator may only be qualified to control a particular style of robot.

Ideally, the problem holder brings to the team an understanding of the activity and how to collect information needed by decision-makers. In particular, the problem holder is familiar with the visual cues of victims (color, heat, motion) and hazards (such as where to look for structural cracks) and general recognition and scene interpretation. The problem holder may be a technical search specialist with little or no robot training.

Level 1 SA (perception) is constructed from raw sensor data by the humans under unfavorable conditions using their *a priori* knowledge. As noted in Sec. III-B, the video data is likely to be poorly illuminated. The use of a single camera produces keyhole effects, reduced depth perception, and a lack of functional presence which interfere with correct perception. The different sensors do not naturally fuse— they do not have

the same field of view or resolution. As captured by LOVR, video data appears better for navigation and general scene interpretation, while thermal imagery is favored for victim detection.

Level 2 SA (comprehension) is concerned with the identification of key features in the data and interpreting them. Different roles for the two human robot team members become even more apparent. The robot operator is concerned with comprehending the environment for navigation and understanding the state of the robot. The key feature for comprehending the navigational context is 3D depth: not just independently of the robot (seeing if there’s an obstacle or a drop off up ahead), but also in relation to the robot (e.g., if a camera or sensor is going to snag on something hanging from overhead or just to the side of the vehicle).

In contrast to the relatively repetitive and reactive nature of navigation, the problem holder is performing a great deal of data fusion. The problem holder is concerned with taking raw data, creating knowledge about the surroundings from the data, then transforming that knowledge into an informative map with victim and hazard identification. The lack of reliable GPS and miniature localization sensors forces the human to estimate distance traveled and sizes of clearings. This leads back to the same problems with 3D depth reconstruction experienced in navigation. The victim and hazard recognition tasks are beyond the current reach of computer vision. Some hazard information, such as ambient temperature, is text, which may create a problem swapping from an image to looking at a number.

Level 3 SA (projection) for the robot operator for search is concerned with developing the *knowledge* needed to safely navigate and to adequately cover the area. As noted earlier, confined spaces do not require a great deal of path planning. Navigation is more of a matter of anticipating navigational hazards such as obstacles or tight turns which might tangle the tether or safety rope. While this does not require projection on a global scale (localization, mapping, and path planning), it does require a local synthesis of information from the environment, the robot, and the operator’s *a priori* knowledge. How to get over an obstacle may be more challenging than which turn to take.

Level 3 SA for the problem holder may appear to be less demanding than Level 2, but is more difficult. It deals with verifying the presence and state of a victim, matters of evidential sufficiency. This requires projection of *new viewpoints* (e.g., will moving it to a new position increase the evidence for a victim or hazard?) and *additional viewers* (e.g., let’s get someone with more experience/sleep over here to look). On one hand, the projection of viewpoints and viewers suggests that the problem holders developing of SA does not require knowledge of the robot’s capabilities. A directive such as “try to get a view from the left side” is referenced to an object in an image that is being shared by both the problem holder and the robot operator, and does not explicitly require robot knowledge. However, in practice, humans in both roles get easily frustrated if there isn’t a shared understanding that the robot can’t look that high, can’t move any farther to the left because of tether limitations, or the camera doesn’t zoom in.

While our experience is that cooperation between humans is a practical necessity for search, the importance of cross-training and familiarity with robots is not clear. On the CRASAR team, the robot operator and problem holder almost always have the same level of robot operations competency. The subjects in our studies [7] had the same introductory robot training. However, in theory, the problem holder does not have to have robot knowledge. Regardless, the same two humans would not necessarily remain a team during an incident or even a work shift. Consider that as seen in Fig. 6, the robot's activity may suddenly change from search to rescue or medical. At that point a technical search specialist serving as the problem holder might be replaced with an emergency medical technician. The need for cooperation is made even more complex by the unknown effects of individual backgrounds and the likelihood of dynamic changes to the activity and team membership.

Timing regarding how SA is developed is also interesting. According to results from [7], a team spends 49% of a search activity developing SA. During that time, the robot is motionless and the team members are actively collaborating and discussing what they think they see. Essentially the humans move the robot forward less than a meter, then have to stop and think about what they are looking at. We have captured this as a part of our LOVR strategy, where the team stops to conduct a localization observation for navigational SA, then conducts a problem holder observation for victims, then is explicitly encourage to talk aloud and reach consensus by reporting.

V. HRI ISSUES

The information model highlights several opportunities for artificial intelligence and distributed network systems to improve robot-assisted search and rescue. As advances in AI and networking become available, they will generate new HRI issues and opportunities for HRI research to contribute to successful systems. These opportunities can be grouped into three broad areas: reducing the number of humans it takes to control a robot, propagating information through geographically distributed teams with intermittent communications, and encouraging acceptance within the existing social structure. It should be noted that rescuers work in pairs for safety reasons, and that reducing the ratio may be less important than facilitating the teamwork between the operator and problem holder.

A. Improving Human to Robot Ratio

One of the most striking aspects of robot-assisted search and rescue is the 2 : 1 human to robot ratio. This places two people at risk per tool, rather than one. There are numerous ways technology could reduce this ratio.

One approach is to distribute one member of the robot team out of the Hot Zone. The robot operator is not a good candidate because someone will have to carry, set up, and repack the robot in the Hot Zone. Also, as shown in [52], fieldable robots have a high probability of failure. A robot expert is needed to repair the robot in the field.

The problem holder is a better candidate to move out of the Hot Zone because that role does not have to be physically co-located. But, as shown in [7], the problem holder is a fully engaged team member, communicating almost constantly with the robot operator to build SA. Other research [53] has shown that teams are more effective if they are physically co-located. And no matter what, the distributed communications network at an incident will not be perfect. Communications will be lost. Bandwidth may be reduced so that video arrives with a significant time lag, lower resolution, or as still images. Each of these problems is expected to negatively impact team performance unless somehow compensated by artificial intelligence. The realities of networks make it hard to move the problem holder out of the Hot Zone without fundamental research in distributed team decision-making processes with unreliable communications.

Another approach to reducing the human to robot ratio is to transfer portions of both roles to the robot so that one person in the field can perform it. The robot operator's role may be greatly reduced as advances in of range sensors and AI algorithms for 3D mapping demonstrated on large robot become available for mini- and micro- classes. However, the robot operator's role as expert problem solver is unlikely to be automated. The types of robot failures, their symptoms and recovery methods, are only now beginning to be investigated by the robotics community. Many failures are related to the complexity of the environment, and there is no known work formally categorizing the different types of rubble and voids in terms of how robots fail. The high likelihood of failures, combined with the known limitations of humans swapping into a human-in-the-loop control mode for a highly autonomous system, suggests the need for HRI research on cooperative failure diagnosis and recovery.

The problem holder's role could be reduced with the advent of 3D mapping, but it is unlikely to be eliminated since it involves interpreting scenes. 3D mapping would guarantee search coverage and replace the sketch information with a "fly through" model. The problem holder may wish to manually abstract the map into a static 2D cross-section and annotate with hazards or other notes. Image processing techniques, such as image enhancement and the use of software agents to cue the problem holder based on affordances, can certainly help the problem holder in searching the scene for victims and hazards but the state of computer vision still far below what is needed. An important question is how will humans react to computer assistance? Will knowing that a software agent is also examining the issue make a human less diligent? Therefore, a fundamental HRI research issue is how to create effective cooperative human-robot perceptual systems.

B. Improving Information Flow

Distributed communications networks offer the potential to both re-locate a robot team member and immediately propagate information to all members of the rescue enterprise. This could significantly alter the existing model of Information Flow, as shown in Fig. 7, which is one-way with inherent time delays in the propagation of information. This poses fundamental research questions in distributed team performance

and decision-making. It also reinforces the need for systems-level consideration of the visualization of information being propagated to diverse users and the general human-computer interfaces (HCI).

The advent of distributed communications networks could greatly reduce time delays in information propagation. Note that information propagation outside of the robot team is currently disruptive and time-consuming. A person must be contacted, they must interrupt their task, and often information (either a piece of paper or a video playback device) must be physically carried to the information consumer or the consumer come to the source. With effective communications networks, the robot team could dramatically speed up information propagation to the Search Task Manager and the Task Force Leader, as well as begin preparing for the next task. Consider the robot team pushing data and information onto a server that can be immediately accessed by the planning team and structural and medical specialists, leading to a smoother, and more rapid, transition from a search activity to rescue and medical activities. But highly distributed communications undermines the decision-making hierarchy and could create information overload. More research, particularly into methods and metrics, is needed in HRI to help identify the trade-offs between distributed and filtered systems.

The Information Flow Model in Fig. 7 is one way, information flows up the hierarchy in pre-packaged batches. A highly distributed, real-time system makes bi-directional information flow possible, where decision-makers could use the robots as active resources to gather additional information on demand. This would change the locus of control of the robot. In Search, the robot is controlled by the robot team which acts largely independently. Once the activity changes, decision makers may need to actively gather new data as needed. They may wish to direct the robot, causing the robot to go from extensions of the robot team to extensions of any member of the distributed community. This raises the issue of whether the members should have direct access to the robot? What do they need to know about the robot (e.g., mental model) in order to effectively task it? How is contention handled between conflicting requests, such as the structural specialist wants the robot to do X and the medical specialist wants it to do Y?

Distributed communications enable the robot's information to be available to a wide set of consumers. This also poses a scientific visualization challenge: how to manipulate the form of the information into the display best suited for consumer. At this time, each member in the hierarchy is responsible for data fusion and formatting the new information into the form best utilized by the next higher level. With information servers, the use of members of the hierarchy to filter data and create appropriate displays may change. It is hoped that software agents could take over some of these tasks. The use of centralized data servers also introduces challenges in asynchronous processing. What happens when all the data is available, except that from the Search Team Manager? How can information from human intelligence be extracted and fused with the robot-generated information?

Issues in HRI are not limited to cognitive and team process

models; there are significant HCI and human factors issues. The ability to distribute images or information over wireless links throughout the rescue enterprise ignores how that information can be created. With pen and paper, it is fairly easy to annotate visual representations with text. Annotations may be a big short-term improvement in information distribution. For example, it might be helpful for the problem holder to annotate an image with notes, e.g., circling an area and writing "is this a problem?" This requires fundamental research in HRI on the appropriate interface modes.

Another HCI issue is the use of PDAs. PDAs linked with wireless 802.11 protocols are now entering the rescue enterprise. On one hand, they are existing technology that can be exploited and can be converted to "dual use." However, robot-assisted search is a perceptual activity. PDAs may not have sufficient screen resolution for visual tasks and could actually degrade performance, leading to a rapid false negative rather than a slower true identification of a survivor. Again, these devices are dependent on the network communications and some provision needs to be made to transfer the data to a hardcopy that can be physically distributed. An open HRI issue is the impact of cognitive and physiological fatigue, resolution, and imagery modality on perception.

C. Facilitating Social Niche

The social niche of a robot defines its interaction with other members of the rescue enterprise: how robot teams are organized within the team of teams USAR structure, how the robot team collaborates with other teams, and how members of the robot team transition between roles (or outsiders are recruited for confirmation). The social niche is impacted by the diversity of users, the types of decision-making that is expected, the social organization of the larger rescue enterprise, and how much training members have about robots or robot-assisted search and rescue.

An important influence on team performance, and even on technology acceptance, is user diversity. USAR has a spectrum of end-users with differing educational backgrounds, goals and methods, and even history together. It is easy to inadvertently limit the consideration of the impact of diversity to the operator and information consumers in the decision-making hierarchy. After all, the primary human involvement in the search activity is controlling the robot and interpreting the sensor data, or interactions "behind" the robot. Naturalistic social interactions with the robot aren't considered necessary because the humans do not see the robot. However, this example is misleading, because in the rescue and medical activities, the robot will be interacting with a survivor, or someone in "front" of the robot. And the survivor will taken from an even larger population of possible users and may be cognitively impaired due to injury and fear. This raises fundamental issues on the impact of age, background, individual preferences, prior experiences and education on both the use of robots and the acceptance to robots.

One way to smooth out differences in the diversity of members of the rescue enterprise (victims excepted) is to provide training. Unmanned vehicle systems, such as military ground robots and aerial vehicles, often require a year

of specialist training to operate in fairly open spaces. The demands of lower-tech, less reliable rescue robots operating in sub-human confined spaces may actually be higher than that of military aircraft. Yet, rescue workers have significantly less time open for training and work in much less predictable situations. Generally fire rescue training for a new technology or specialist position is on the order of two weeks or less, with one day a year for recertification. This presents a huge challenge in creating both systems which do not require huge amounts of training and in designing effective training procedures.

The effectiveness of training also depends on the locale. Training in a standard rubble pile used for canine searches is very limited compared to training in collapsed structures. However, even federal teams have difficulty overcoming the legal obstacles in using demolished buildings for training, should those buildings happen to be in the area. Simulated courses used for robot competitions are not high fidelity.[54] Human-robot interaction research should help define what is the minimal suitable training facility.

Since robots are a remotely operated technology, one possible approach to the training problem is to provide more fluid training opportunities. One of the limitations on training now is that it is done in standard classroom format with groups of rescuers and an instructor. It may be possible to effectively train rescuers to use robots or at least to keep up their competency over the Internet. The promise of Internet-based training creates other challenges for HRI to research, especially how to measure the effectiveness of such training and whether the student will have confidence (either too much or too little) in their abilities.

The social organization of the rescue enterprise is a consideration in how the technology can be effectively presented and accepted. One FEMA task force leader immediately saw robots as a “union busting” technology; that it would change the nature how fire fighters work and train. While such an extreme organizational change in response to robotics seems a bit far-fetched, rescue robots could change the current paradigm of narrow specialists to one where robot generalists use the same robot with different payloads to support those specialists.

VI. SUMMARY AND CONCLUSIONS

As a Grand Challenge problem for HRI, rescue robotics is an intriguing application domain in which to explore human-robot interaction issues. In this domain robots do not replace human searchers, canines, or existing tools, but instead offer new capabilities. As a result, it represents an opportunity to observe the evolution of creativity, social acceptance, and formation of relationships.

Humans are heavily involved in rescue robotics, from specialists (robot operators) to skilled workers who use the information extracted from the robots (other rescue workers, medical and structural experts) to ordinary people (victims). Rescue robotics is a “team of teams” process. The state of the practice is teleoperation of robots, with a 2 : 1 human to robot ratio. Robot teams of two can better handle the demands of navigating within the highly cluttered interior of a rubble pile

while attempted to identify potential survivors and outperform a single operator.

Although humans are highly involved in the search activity, it is primarily to interpret raw sensor data. The bulk of true decision making is when and how to acquire more perceptual evidence for a victim or hazard. Procedural decision alternatives are codified into procedures. The flow of information is upward through a hierarchy, reducing the need for communication but introducing some delays in propagation of vital information. The advent of advances of robot autonomy and distributed communications have the potential to flatten this hierarchy, but may overload the key decision-makers.

As seen at the WTC and follow-up field studies, the HRI challenges posed by the current state of the practice of rescue robots are already exceeding the human team members’ abilities to cope. Technology-oriented advances without corresponding human-centered advances may lead to rescue robots which may miss survivors. We believe that the key HRI research questions which have the most significant long-term and short term benefit to rescue robotics are:

- 1) How can visual information be propagated to distributed users with diverse information needs, priorities, and deadlines in the presence of communications interruptions and degradations?
- 2) How can intelligent systems be designed to facilitate distributed human-robot teams in the collaborative acquisition of data and transformation into information and information into knowledge for novel, high-stress situations?

The above HRI challenges are similar to those of military operations in urban terrain, police SWAT teams, and battlefield medicine. Robot operators and consumers of the robot-derived information come from a diverse population. They are acting under extreme time pressure, personal risk, distractions, and possible cognitive and physiological fatigue. Rescue workers do not control or predict the pace of work. The terrain is unknown, deconstructed, and difficult to interpret from the robot’s viewpoint, with key features (in the case of USAR, victims and structural hazards) obscured. The robots themselves are subject to subtle failure modes due to the novelty of the terrain, requiring human expertise at unpredictable intervals. The information flow is hierarchical and compartmentalized, yet could be profoundly changed by distributed communications.

Unfortunately, good application domains generally have the disadvantage of high barriers to entry. Rescue robotics is no exception, as fieldable rescue robots are expensive and opportunities to study their use in the field with task forces are rare. Through funding from the National Science Foundation, CRASAR is offering opportunities for scientists to conduct field work in rescue robotics. See www.crasar.org for more details.

Acknowledgments

The work reported in this article has been funded in part by the DARPA SCF Seedling Program, the DARPA CSEE Program, the Global Center for Disaster Management and

Humanitarian Assistance, and donations from SAIC. The author would like to especially thank Dawn Riddle for her helpful comments and suggestions on the preparation of this article, as well as Jenny Burke and Thomas Fincannon; Chief Ron Rogers of Hillsborough County Fire Rescue Department, Florida Task Force 3, and Gary Giddens and Rescue Training Associates for their expertise; and Sam Stover and the CRASAR response team for their continued efforts.

REFERENCES

- [1] J. Casper, "Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center," Ph.D. dissertation, MS thesis, Department of Computer Science and Engineering, University of South Florida, 2002.
- [2] J. Casper and R. Murphy, "Human-robot interaction during the robot-assisted urban search and rescue effort at the world trade center," *IEEE Transactions on Systems, Man and Cybernetics Part B*, vol. 33, no. 3, pp. 367–385, 2003.
- [3] M. Micire, "Analysis of robotic platforms used at the world trade center disaster," Ph.D. dissertation, MS thesis, Department of Computer Science and Engineering, University of South Florida, 2002.
- [4] C. on Science and N. R. C. Technology for Countering Terrorism, *Making the Nation Safer: The Role of Science and Technology in Countering Terrorism*. Washington, DC: National Academies Press, 2002.
- [5] *Grand Research Challenges in Information Systems*. Washington, DC: Computing Research Association, 2003.
- [6] J. Casper and R. Murphy, "Workflow study on human-robot interaction in usar," in *ICRA 2002*, 2002, pp. 1997–2003.
- [7] J. Burke, R. Murphy, M. Covert, and D. Riddle, "Moonlight in miami: An ethnographic study of human-robot interaction in usar," *Human-Computer Interaction, special issue on Human-Robot Interaction*, p. to appear, 2004.
- [8] A. Gage, R. Murphy, and B. Minten, "Shadowbowl 2003: Lessons learned from a reach-back exercise with rescue robots," *IEEE Robotics and Automation Magazine*, to appear.
- [9] E. Rogers, R. Murphy, and J. Burke, "Nsf/darpa study on human-robot interaction," *IEEE Transactions on Systems, Man and Cybernetics, special issue on Human-Robot Interaction*, to appear.
- [10] A. Agah and K. Tanie, "Taxonomy of research on human interactions with intelligent systems," in *Proceedings of 1999 IEEE International Conference on Systems, Man, and Cybernetics*, vol. 6, 1999, pp. 965–970.
- [11] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, vol. 42, no. 3–4, pp. 143–166, 2003.
- [12] R. C. Arkin, M. Fujita, T. Takagi, and R. Hasegawa, "An ethological and emotional basis for human-robot interaction," *Robotics and Autonomous Systems*, vol. 42, no. 3–4, pp. 191–201, 2003.
- [13] C. Breazeal, *Designing Sociable Robots*. The MIT Press, Cambridge, MA, 2002.
- [14] —, "Regulation and entrainment in human-robot interaction," in *International Journal of Robotics Research*, vol. 21, no. 10–11, 2002, pp. 883–902.
- [15] T. Shibata, K. Wada, and K. Tanie, "Tabulation and analysis of questionnaire results of subjective evaluation of seal robot in uk," in *IECON - 2002. 2002 28th Annual Conference of the IEEE Industrial Electronics Society*, vol. 4, 2002, pp. 3074–3079, note: need to get this and refer to it for the methodology part of introduction.
- [16] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz, "Probabilistic algorithms and the interactive museum tour-guide robot minerva," *International Journal of Robotics Research*, vol. 19, no. 11, pp. 972–999, 2000.
- [17] I. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto, "An affective mobile robot educator with a full-time job," *Artificial Intelligence*, vol. 114, no. 1–2, pp. 95–124, 1999.
- [18] W. Burgard, A. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "Experiences with an interactive museum tour-guide robot," *Artificial Intelligence*, vol. 114, no. 1–2, pp. 3–55, 1999.
- [19] "Insect telepresence: using robotic tele-embodiment to bring insects face-to-face with humans," *Autonomous Robots*, vol. 10, no. 2, pp. 149–161, 2001.
- [20] A. D. Brito, J. C. Gamez, D. H. Sosa, M. C. Santana, J. L. Navarro, J. I. Gonzalez, C. G. Artal, J. P. Perez, A. F. Martel, and H. Teiera, "Eldi: an agent based museum robot," *Systems Science*, vol. 27, no. 4, pp. 119–128, 2001.
- [21] K. Severinson-Eklundh and A. G. H. Huttenrauch, "Social and collaborative aspects of interaction with a service robot," *Robotics and Autonomous Systems*, vol. 42, no. 3–4, pp. 223–234, 2003.
- [22] F. Wullschleger and R. Brega, "The paradox of service robots-how passers-by can contribute in solving non-deterministic exceptional conditions encountered by service robots," in *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2002, pp. 1126–1131.
- [23] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun, "Towards robotic assistants in nursing homes: Challenges and results," *Robotics and Autonomous Systems*, vol. 42, no. 3–4, pp. 271–281, 2003.
- [24] M. Hans, B. Graf, and R. Schraft, "Robotic home assistant care-o-bot: past-present-future," in *11th IEEE International Workshop on Robot and Human Interactive Communication (ROMAN 2002)*, 2002, pp. 380–385.
- [25] "Analysis of factors that bring mental effects to elderly people in robot assisted activity," in *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2002, pp. 1152–1157, seal guys.
- [26] R. Ambrose, H. Aldridge, R. Askew, R. Burridge, W. Bluethmann, M. Diffler, C. Lovchik, D. Magruder, and F. Rehnmark, "Robonaut: Nasa's space humanoid," *IEEE Intelligent Systems*, vol. 15, no. 4, pp. 57–63, 2000.
- [27] R. Burridge and J. Graham, "Providing robotic assistance during extra-vehicular activity," in *Proceedings of the SPIE - The International Society for Optical Engineering*, vol. 4573, 2002, pp. 22–33.
- [28] H. Jones, S. Rock, D. Burns, and S. Morris, "Autonomous robots in swat applications: Research, design, and operations challenges," in *Proceedings of the 2002 Symposium for the Association of Unmanned Vehicle Systems International (AUVSI '02)*, 2002.
- [29] R. Arkin, T. Collins, and Y. Endo, "Tactical mobile robot mission specification and execution," 2000.
- [30] S. Hirose and E. Fukushima, "Development of mobile robots for rescue operations," *Advanced Robotics*, vol. 16, no. 6, pp. 509–512, 2002.
- [31] J. Scholtz and S. Bahrami, "Human-robot interaction: development of an evaluation methodology for the bystander role of interaction," in *IEEE International Conference on Systems, Man and Cybernetics*, vol. 4, 2003, pp. 3212–3217.
- [32] M. Diffler, F. Huber, C. Culbert, R. Ambrose, and W. Bluethmann, "Human-robot control strategies for the nasa darpa robonaut," in *Proceedings of 2003 IEEE Aerospace Conference*, vol. 8, 2003, pp. 3939–3947.
- [33] H. Jones and P. Hinds, "Extreme work groups: Using swat teams as a model for coordinating distributed robots," in *ACM 2002 Conference on Computer Supported Cooperative Work (CSCW 2002)*, 2002.
- [34] W. D. D. and N. Sarter, *Cognitive Engineering in the Aviation Domain*. Erlbaum, Hillsdale NJ, 2000, ch. Learning from Automation Surprises and Going Sour Accidents.
- [35] C. Billings, *Aviation Automation: The Search For A Human-Centered Approach*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1996.
- [36] M. Osuka, R. Murphy, and A. Schultz, "Usar competitions for physically situated agents," *IEEE Robotics and Automation Magazine, special issue on Rescue Robotics*, vol. 9, no. 4, pp. 31–40, 2002.
- [37] T. Takahashi and Tadokoro, "Working with robots in disasters," *IEEE Robotics and Automation Magazine*, vol. 9, no. 3, pp. 34–39, 2002.
- [38] S. Tadokoro, H. Kitano, T. Takahashi, I. Noda, H. Matsubara, A. Hinjoh, T. Koto, I. Takeuchi, H. Takahashi, F. Matsuno, M. Hatayama, J. Nobe, and S. Shimada, "The robocup-rescue project: a robotic approach to the disaster mitigation problem," in *IEEE International Conference on Robotics and Automation*, vol. 4, 2000, pp. 4089–4094.
- [39] R. Murphy, J. Casper, and M. Micire, *Potential Tasks and Research Issues of Mobile Robots in RoboCup Rescue*. Springer Verlag, Berlin, 2001, pp. 339–344.
- [40] A. Jacoff, E. Messina, and J. Evans, "A standard test course for urban search and rescue robots," in *Proceedings of the 2000 PerMIS Workshop (NIST SP 970)*, 2001, pp. 253–259.
- [41] R. R. Murphy, J. Casper, M. Micire, and J. Hyams, "Assessment of the nist standard test bed for urban search and rescue," in *NIST Workshop on Performance Metrics for Intelligent Systems*, 2000.
- [42] J. Drury, J. Scholtz, and H. Yanco, "Awareness in human-robot interactions," in *IEEE International Conference on Systems, Man and Cybernetics*, vol. 1, 2003, pp. 912–918, robocup results.
- [43] I. Erkmen, A. Erkmen, F. Matsuno, R. Chatterjee, and T. Kamegawa, "Snake robots to the rescue!" *IEEE Robotics and Automation Magazine*,

- vol. 9, no. 3, pp. 17– 25, 2002, propose snake robots and show simulations.
- [44] R. Dollarhide and A. Agah, “Simulation and control of distributed robot search teams,” *Computers and Electrical Engineering*, vol. 29, no. 5, pp. 625–642, 2003.
- [45] J. Jennings, G. Whelan, and W. Evans, “Cooperative search and rescue with a team of mobile robots,” in *1997 8th International Conference on Advanced Robotics*, 1997, pp. 193–200.
- [46] J. G. Blicht, “Artificial intelligence technologies for robot assisted urban search and rescue,” *Expert Systems with Applications*, vol. 11, no. 2, pp. 109–124, 1996.
- [47] R. Murphy, “Marsupial and shape-shifting robots for urban search and rescue,” *IEEE Intelligent Systems*, vol. 15, no. 3, pp. 14–19, 2000.
- [48] E. Krotkov and J. Blicht, “The defense advanced research project agency (darpa) tactical mobile robotics program,” *International Journal of Robotics Research*, vol. 18, no. 7, pp. 769–776, 1999.
- [49] J. Casper, R. Murphy, and M. Micire, “Issues in intelligent robots for search and rescue,” in *SPIE Ground Vehicle Technology II*, 2000.
- [50] R. Murphy, J. Casper, J. Hyams, M. Micire, and B. Minten, “Mobility and sensing demands in usar,” in *IECON-2K:session on Rescue Engineering*, 2000.
- [51] *Technical Rescue Program Development Manual*. United States Fire Administration, 1996.
- [52] J. Carlson and R. R.R. Murphy, in *ICRA '03. IEEE International Conference on Robotics and Automation*, vol. 1, 2003, pp. 274–281.
- [53] R. Kraut, S. Fussell, and J. Siegel, “Visual information as a conversational resource in collaborative physical tasks,” *Human-Computer Interaction*, vol. 18, pp. 13–49, 2003.
- [54] R. Murphy, J. Blicht, and J. L. Casper, “Robocup/aaai urban search and rescue events: Reality and competition,” *AI Magazine*, pp. 37–42, 2002.

PLACE
PHOTO
HERE

Robin R. Murphy Robin Roberson Murphy (M¹ 1987) received a B.M.E. in mechanical engineering, a M.S. and Ph.D in computer science in 1980, 1989, and 1992, respectively, from Georgia Tech, where she was a Rockwell International Doctoral Fellow. She is a professor in the Computer Science and Engineering Department at the University of South Florida with a joint appointment in Cognitive and Neural Sciences in the Department of Psychology, and also serves as the Director of the Center for Robot-Assisted Search and Rescue. Dr. Murphy is

also the recipient of an NIUSR Eagle Award for her participation at the World Trade Center.

She is the author of over 70 publications in the areas of multi-agent sensor fusion, human-robot interaction, and rescue robotics as well as the textbook, *Introduction to AI Robotics*. Her research is funded by NSF, DARPA, ONR, DOE and industry and includes ground and aerial unmanned systems.

Dr. Murphy is a member of IEEE, AAAI, ACM, and SPIE. From 2000-2002, she served as the secretary of the IEEE Robotics and Automation Society and is the North American co-chair of the Safety, Security and Rescue Robotics technical committee. She co-chaired the first IEEE Workshop on Safety, Security, and Rescue Robotics (2003). Since 2000, she has been a member of the US Air Force Scientific Advisory Board and participated in numerous defense studies and panels on unmanned systems.