AI Magazine Volume 25 Number 2 (2004) (© AAAI)

National Science Foundation Summer Field Institute for Rescue Robots for Research and Response (R4)

Robin R. Murphy

■ Fifteen scientists from six universities and five companies were embedded with a team of search and rescue professionals from the Federal Emergency Management Agency's Indiana Task Force 1 in August 2003 at a demolished building in Lebanon, Indiana. The highly realistic 27-hour exercise enabled participants to identify the prevailing issues in rescue robotics. Perception and situation awareness were deemed the most pressing problems, with a recommendation to focus on human-computer cooperative algorithms because recognition in dense rubble appears far beyond the capabilities of computer vision for the near term. Human-robot interaction was cited as another critical area as well as the general problem of how the robot can maintain communications with the rescuers. The field exercise was part of an ongoing grant from the National Science Foundation to the Center for Robot-Assisted Search and Rescue (CRASAR), and CRASAR is sponsoring similar activities in summer 2004.

The Center for Robot-Assisted Search and Rescue (CRASAR) at the University of South Florida hosted on 2–3 August 2003 the first summer field institute for scientists working, or wanting to work, in the field of rescue robotics. The purpose of the annual summer field institute is to put scientists directly in the field with rescue professionals and fieldable rescue robots to better understand the domain and conduct experiments. The institute is part of the Rescue Robots for Research and Response (R4) Program sponsored by the National Science Foundation's Computer and Information Science and Engineering (CISE) directorate. The R4 program is a three-year grant whose goal is to facilitate information technology research into robot-assisted urban search and rescue (USAR) by providing access to domain experts, meaningful test sites, and expensive specialized equipment. All equipment used for the R4 program is fieldable, so the grant also increases the availability of mobile robots and sensors for an emergency response, such as the World Trade Center disaster.

The specific objective of this field institute was to embed 15 scientists with rescue workers as they went through a complete deploy-search-cleanup cycle or "evolution." FEMA response force Indiana Task Force 1 (INTF-1) hosted the event and served as subject-matter experts with the CRASAR response team and members of the United States Marine Corps Chemical Biological Incident Response Force (CBIRF). INTF-1 arranged for the partial demolition of the old Lebanon, Indiana, town library to create a realistic collapse site (figure 1). Mannequins were placed within the structure before it was demolished. The scientists observed the INTF-1 technical search team as it arrived on site and conducted a reconnaissance of the collapsed building. During the technical search phase, the scientists went in small groups with the rescue workers when they conducted a complete technical search using both traditional search tools and rescue robots supplied by CRASAR. The scientists also got to observe the process by which the search team manager decided whether to use traditional tools, such as acoustic sensors or search cameras, or a robot (figure 2).

Three robots were used: two of the Inuktun MICRO-VGTV class, which was the most commonly deployed robot at the World Trade Center response, and an iRobot PACKBOT, which is being evaluated for use. A fourth robot built by Carnegie-Mellon University (CMU), a snake mounted on a wheeled base, was tested in the rubble but was not part of the deployments. The robots were deployed nine times, averaging about one hour for each deployment. In each case, either INTF-1 or CBIRF would identify or construct a typical search scenario. The search team manager would then request robots over the radio; the team manager would discuss with Robin Murphy, the CRASAR team leader, what robots to use and where; and, finally, three to four scientists would be deployed with the rescue robot squad to perform the search. Scientists witnessed over 10 hours of actual robot activity plus an additional 10 hours of preparation and decontamination time (figure 3). Each scientist went into the rubble at least two times and witnessed the deployment of each brand of robot. The robot operation was taped, and a DVD was created, which is available on request.

The Indiana evolution lasted 27 hours to introduce cognitive and

Figure 1. Lebanon Town Library Partially Demolished under INTF-1's Direction to Simulate an Earthquake.

Note that robots are needed to go into the interior of the collapse, under surfaces that people might walk on.

Figure 2 (below). A Group of Scientists and Graduate Students Listen to INTF-1 Technical Search Specialist Sam Stover Explain about Structural Collapses and the Markings Used to Represent the State of the Search.

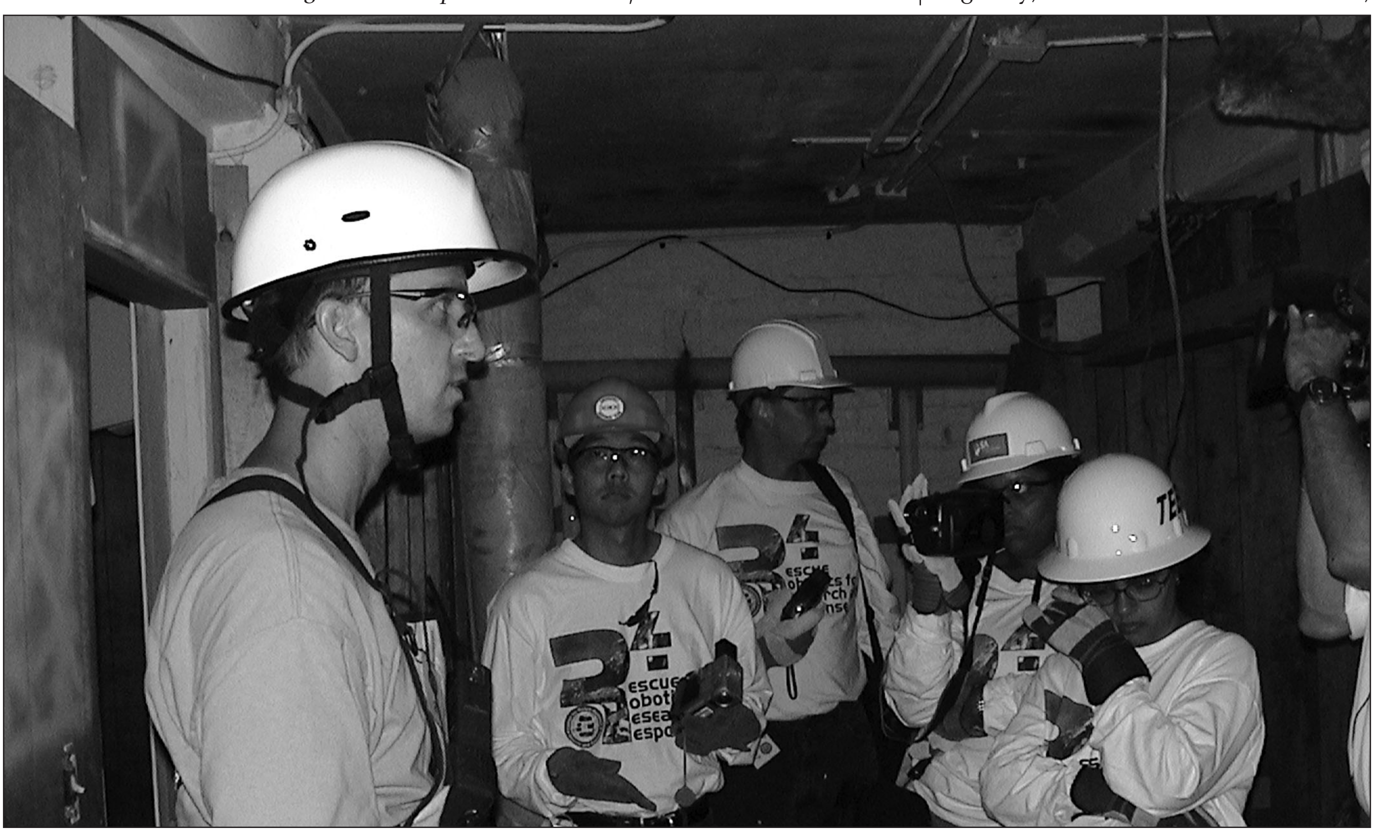

physical fatigue. Embedding under such realistic conditions permitted the participants to gain an ethnographic understanding of rescue robotics, direct access to one type of collapse site, and an introduction to standard operating procedures such as decontaminating the robots that might impact the design of better robots and software. The scientists brought sleeping bags and slept during the single four-hour rest cycle allotted to the rescue workers. After the 27-hour evolution, the scientists conducted a 2-hour outbriefing before catching flights back home.

The field camp was attended by six university researchers and three graduate students from CMU, The Ohio State University, Stanford University, Texas A&M University, and University of Minnesota. Six research scientists from the National Institute of Standards and Technology, iRobot, SA Technologies, Palo Alto Research Corporation, and Time Domain also observed.

The primary contribution of the event was the identification of research issues in physically situated agency, human-robot interaction,

Figure 3. Scientists Participated in All Aspects of the Response, Including Decontaminating Robots from Dirt (and Possible Biological Residue from Ruptured Sewers or Body Fluids).

Decontamination of minisized robots is on the order of an hour, about half the expected run time. A robot or sensor that is not designed to be quickly cleaned off is not fieldable.

and distributed systems. These research issues are in addition to the traditional problems posed by more autonomous robotics in outdoor navigation; thus, they add to the canon of challenges being addressed by the AI community. Independent of the value of identifying new research issues, the majority of respondents said the most informative part of the institute was simply being there, embedded with rescuers and robots under realistic conditions.

Robots, as physically situated agents, require perception to act rationally in the world. The field camp reinforced the challenges associated with two fundamental questions of perception: (1) how to detect victims and unsafe conditions for rescuers in a highly cluttered, unfavorable environment and (2) how to ensure sensor coverage of a volume of space. The collapsed library was unusual in that it posed favorable conditions for survivors, with voids large enough for people to enter and little dust coating survivors and rubble. Even under this best-case scenario, experienced robot operators and scientists missed most of the mannequins placed in even open areas. These errors were presumed to be because of the dark, keyhole effects from narrow-sensor fields of view as well as unnatural viewpoints (the robot's sensors were generally only three inches above the ground plane). The scientists were greatly surprised by the lack of miniature sensors available for USAR. Currently, only video cameras and thermal and gas detectors are commercially available in small enough sizes to be deployed. It was agreed that autonomous recognition of either victims or interesting structural artifacts was far beyond the capability of computer vision for the foreseeable future. A more productive approach might be cooperative sensing, where the human and sensor agents work together. The most simple form would be the robot identifying a region of interest, then cuing the human to examine the image more closely. Unfortunately, the consensus was that robots are sensor impoverished, and manufacturers are not providing adequate user interfaces or application programmer interfaces to permit extensions. Another issue in perception is to ensure that a void has been completely sensed by all sensors. Video and FLIR cameras do not have the same field of view or resolution, yet both contribute to the search activity. The robot might sweep enough so that the video camera (which is used for structural assessment) is exposed to the entire void interior but not enough that the FLIR (which is used for victim detection) has complete coverage. A further challenge for sensor coverage is that it might not be possible for the robot to position the sensors in such a way as to get complete coverage or only through a complex set of motions. Planning will be required.

The ability of a robot to act on the environment presented another set of research challenges. Current robots only move through the environment. Although some robots have arms, these arms are more sensor stalks than arms; they are not capable of lifting more than a few pounds, so manipulation of the environment is not a practical consideration. Au-

Photo courtesy Richard Voyles

Figure 4. A Tethered Inuktun Micro-VGTV Robot Being Lowered from the Upper Story of the Library.

Many voids in a rubble pile are vertical or have a vertical portion; so, robots must often be attached to a safety line even if the robot uses wireless communications .

tonomous navigation remains a problem for several reasons. First, more mobile platforms are needed. The group found that mobility is harder than was expected, especially because robots are most useful operating in the interior of the rubble, not crawling on the surface of the pile. Second, a terrain taxonomy needs to be developed to inform platform designers. Polymorphic vehicles appear to have the advantage, leading to the third problem of how to autonomously adapt the platform configuration to the terrain (figure 4). Fourth, the lack of miniature sensors and the complexity of the environment is preventing the application of simultaneous localization and mapping (SLAM) algorithms that have already been proven for much larger vehicles. SLAM algorithms require accurate range detection, such as with lasers, which are currently too big to be mounted on microclasses and miniclasses of deployable rescue robots. Other range modalities, such as sonar or infrared, do not work in highly confined, narrow, and irregular voids.

The field camp illustrated many of the problems of human-robot interaction (HRI) in search and rescue. A basic constraint is that the human's visual channel is the primary means of perceiving robot control information, constructing situation awareness, and performing object recognition. The scientists found that a search-and-rescue site is very noisy and would interfere with voice-recognition commands. Likewise, the use of gloves makes personal digital assistant–size devices physically difficult to interact with, and safety glasses further distort the low-resolution screens. One important question in HRI is how to present information to the user.

Although the Lebanon exercise focused only on the technical search team, it was clear that information is manually propagated up the incident command hierarchy, and any findings must be verified by another party, which poses challenges for distributed agency. On the one hand, the potential for wireless networks to speed up the flow of information is tremendous. On the other hand, the scientists witnessed the unpredictable communication degradations and dropout that occurred during the use of a wireless robot outside the immediate line of sight. Improvements in wireless technology are expected to reduce some of the problems encountered, but in real life, a robot or a rescuer might just be a millimeter away from being too deep in the rubble for the wireless network to function. Therefore, a key question is how to maintain communications. Proposed approaches include the use of other robots or beacons as *repeaters* (also known as *breadcrumbs*), autonomous behaviors to direct the robot to move until it has reestablished communications, and dynamic resource allocation of bandwidth. Putting data gathered by the robot onto a local network means that a larger set of people can consume the information. For example, remote structural engineers and medical personnel could be looking at the data and advising the robot operator (or robot itself) directly rather than waiting until the end of the shift to review tapes or be faxed reports. However, the unregulated demand for data could slow the overall flow of vital data and reduce overall team effectiveness. The availability of data introduces the larger question of how to maintain distributed communications and appropriate information flow to users with diverse needs and priorities?

The next field outing is on 15 to 17 June with Louisiana State University.[1] The outing is expected to be either on the incident command structure and information flow in a response or mobility and sensor challenges in rubble. Participation is open to scientists, although space is limited.[2]

Note

1. www.crasar.org.

2. For more information, contact Robin Murphy, murphy@csee.usf.edu.

Robin Murphy received a B.M.E. in mechanical engineering and an M.S. and a Ph.D. in computer science in 1980, 1989, and 1992, respectively, from the Georgia Institute of Technology. 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. She is also director of the Center for Robot-Assisted Search and Rescue at the University of South Florida.