AAAI/RoboCup-2001 Urban Search and Rescue Events

Reality and Competition

Robin Murphy, John Blitch, and Jennifer Casper

■ The RoboCup Rescue Physical Agent League Competition was held in the summer of 2001 in conjunction with the AAAI Mobile Robot Competition Urban Search and Rescue event, eerily preceding the September 11 World Trade Center (WTC) disaster. Four teams responded to the WTC disaster through the auspices of the Center for Robot-Assisted Search and Rescue (CRASAR), directed by John Blitch. The four teams were Foster-Miller and iRobot (both robot manufacturers from the Boston area), the United States Navy's Space Warfare Center (SPAWAR) group from San Diego, and the University of South Florida (USF). Blitch, through his position as program manager for the Defense Advanced Research Projects Agency (DARPA) Tactical Mobile Robots Program, was a supporter of the competition; he also served as a member of the rules committee and a judge. USF participated by chairing the rules committee, judging, assisting with the logistics, providing commentary, and demonstrating tethered and wireless robots whenever entrants had to skip around during the competition.

Based on our experiences and history, we were asked to comment on the validity of the competition. The CRASAR collective experience suggests that most of the basic rules of the competition matched reality because the rules accurately reflected deployment scenarios, but the National Institute of Standards and Technology (NIST) Standard Test Course, and hardware or software approaches forwarded by competitors in last summer's event, missed the mark. This article briefly reviews the types of robots and missions used by CRASAR at the WTC site, then discusses the robot-assisted search and rescue effort in terms of lessons for the competition.

CRASAR is a nonprofit National Institute for Urban Search and Rescue (NIUSR) center of excellence created to address confined space access and operations, among other goals. By 9:15 AM on September 11, operating under a standing invitation from fellow NIUSR board member Special Operations Chief Ray Downey for the Fire Department of New York (FDNY), Blitch called on robot manufacturers and field robotics groups to respond with equipment and personnel to the ever-widening disaster. By late afternoon, Blitch and the Foster-Miller and iRobot teams had driven in and met at Stewart Airfield in Newburgh, New York. They then proceeded to the disaster site, gaining access late in the evening. By early morning, the teams had begun to deploy robots in the rubble. USF arrived in Newburgh Wednesday morning after driving straight from Tampa. The United States Navy Space Warfare Center (SPAWAR) joined the effort on Friday, 14 September. The first 11 days concentrated primarily on deploying robots for rescue and recovery operations (finding survivors or victims) with Federal

Emergency Management Agency (FEMA) task force rescue teams, then shifted to structural inspection with city engineers and FDNY personnel. CRASAR personnel were on site through October 2.

The robots used at the World Trade Center (WTC) fall into two categories: (1) tethered and (2) wireless. CRASAR and USF contributed four tethered robots, two Inuktun MICROTRACS and two MICROVGTVs, which were heavily used. (One of USF's tethered robots had been demonstrated at RoboCup 2000 in Melbourne, Australia.) Two other tethered robot models were available but not used. Each of the micro models is about the size of a shoebox, with a video camera, headlights, and two-way audio and controlled and powered remotely through a 100-foot tether. Because of their small size, these robots were extremely effective in entering narrow, winding voids. They were used directly on the rubble pile at WTC 1 and WTC 2 and routinely penetrated 20 to 45 feet in the rubble pile versus the 8-foot limits on traditional search cameras. These teleoperated robots were responsible for finding at least five victims and determining whether voids were safe for human exploration or worth rubble removal for further investigation.

Because of their small size, these robots were extremely effective in entering narrow, winding voids. They were used directly on the rubble pile at WTC 1 and WTC 2 and routinely penetrated 20 to 45 feet in the rubble pile versus the 8-foot limits on traditional search cameras.

Four models of wireless robots were used: one for searching larger voids in the rubble (Foster-Miller's SOLEM) and three for a proof-of-concept search of buildings with collateral damage (Foster-Miller's TALON, iRobot's PACKBOT, and SPAWAR's URBOT). In addition, USF's iRobot URBANS (demonstrated at the 2001 National Conference on Artificial Intelligence) and the ATRV robot were available but not used because of the superiority of PACKBOTS over URBANS and the size of the ATRV. The four wireless models are the size of small carry-on suitcases and carry video; two-way audio; additional sensors, such as FLIR (forward-looking infrared); and effectors or sensor masts. The larger size makes them well suited for searching semistructured environments such as relatively intact buildings. All robots were teleoperated because of the challenging environment and the lack of time to port existing AI software to the new systems. These robots found no victims.

Observations about the Competition

Three of the most frequently disliked aspects of the competition rules turned out to be the most realistic components: (1) the acceptability of teleoperation in an "AI competition," (2) the inclusion of transportation and placement of the robots as part of the competition, and (3) the emphasis on accurately reporting the location of victims. These aspects are discussed in detail here. As a side effect of participating in competitions, USF was well prepared in terms of logistics and able to field a team that was self-sufficient. All USF team members have participated in at least one robot competition, and the adviser, Robin Murphy, has fielded teams to the American Association for Artificial Intelligence (AAAI), Association of Unmanned Vehicle Systems International, or RoboCup competitions annually since 1993. As a result of this prior experience, USF had packing materials for the robots and a standing manifest of tools and sundries to ship for robot support. USF transported all needed equipment to the WTC site without breakage and only had to purchase asbestos-rated respirators to complement their personal safety gear and a few parts from Radio Shack to transfer USF sensors to other CRASAR robots. In addition, previous experience napping on convention center floors during competitions prompted USF to bring sleeping bags and pillows, which were immediately put to use as a shared resource with other CRASAR members. When FEMA provided cots and daily hot meals to workers staged in the Javits Convention Center later in the week, the joke became that the difference between a AAAI robot competition and a disaster site was that a disaster site was more comfortable.

The acceptability of teleoperation has been controversial. One of the main differences between the RoboCup and AAAI competitions has been the role of AI. RoboCup has focused on promoting fieldable technical approaches for urban search and rescue (USAR), of which AI is expected to play a role; AAAI has focused on interesting AI approaches, with their clear relevance to fieldable USAR robots a secondary consideration. It should be noted that despite the on-board computing capability of at least one model of wireless robot, all CRASAR robots at the WTC were teleoperated, which is not to say that AI software did not exist or would not have been useful. Indeed, USF and iRobot had software for older robots that could not be ported in the field in time. It is clear that searching buildings that have collateral damage could be automated, although work in the rubble pile is more demanding. However, regardless of the level of automation, the robots have to work as team members with the rescue workers and others. Rescue workers have to be informed of the findings and most likely cooperatively identify objects and victims. In addition, certain mission payloads presume the involvement of a human; for example, a medical payload was available and Dr. Eric Rasmussen, chief surgeon, U.S. Navy Third Fleet, was on standby from San Diego. Another aspect to consider in automation is sociological: Even if searching the rubble pile could be fully automated, it is unclear whether the technologically conservative rescue community would accept such solutions. Therefore, a mixed-initiative approach is expected to be the most productive from the performance and user-acceptance standpoints; the near-term and the competition rules should not penalize teleoperation.

Transportation of the robots was a major issue. Rescuers and robots had to be staged at Javits Convention Center, miles from Ground Zero. All team members would have to leave the rubble pile and return to the convention center at the end of a 12-hour shift. The nearest any vehicle could get to the rubble pile was generally 2 to 3 blocks away. From the drop-off point, the robots had to be hand carried to a Base of Operations (BOO) at the beginning of a shift, then later to the forward station on request, and then to the rubble pile. The robots were requested for use multiple times during a shift, which equates to a significant time hauling the robots between the BOO and the forward station. The solution was to put the robot and operator control unit in one or more backpacks. The weight of an Inuktun system was between 60 and 80 pounds, depending on the model and size of the tether. Figure 1 shows a void where rescuers had to climb down 40 feet on a ladder, cross the ravine, and then climb 30 feet up to reach a void on the other side. Although this degree of difficulty might seem extreme, stand-off distances of 250 feet are not uncommon in large building collapses, especially when there is the possibility of a secondary collapse. One might make the argument that if the rescue community accepted robots, they could make accommodations for the robots just like for dogs. It is instructive to consider how few accommodations are made for the search and rescue dogs; they sit in a seat on a bus or even an airplane, and their logistics trail is less than a human's. A robot more demanding than a person or dog in terms of transportation and maintenance is unlikely to be effective. In the future, the competition rules should place a maximum weight and footprint on entries to be more realistic and encourage fieldable robots.

Robots that were both more easily transported and required less people were clearly favored by FEMA rescue teams. Complete Inuktun robot systems that could fit into a single backpack and could be run by one person were fielded more frequently than larger robots that required two backpacks and persons. Robots that could not be hand carried or fit into luggage racks on a bus or van simply were not viable because it would take special transportation arrangements. In a situation such as the Ismet, Turkey, earthquake where the site was a city with lots of open space, larger robots could have been used to carry gear and smaller robots from site to site but still could not have entered the small voids of the collapsed structures. The point is that larger robots take more people just to carry them. Aside from the Inuktuns, the remaining CRASAR pool of robots required two people for each robot for portage, not necessarily control. This factor is limiting in getting them on the rubble pile for larger voids. In general, the fire rescue sector chiefs in charge of sections of the rubble pile wanted as few people as possible on the rubble pile or in a structure. Appropriately, the current competition rules penalize entries for more than one person on site.

Once the FEMA task forces and robot crews were at the forward station at the edge of the rubble pile, they often waited before being deployed or returned to the BOO. At the point of deployment, they would have to scramble to climb into the rubble pile and insert the robot. In many cases, the rescue crews were only allowed to investigate voids during crane operators' breaks. Thus, the robot operators had on the order of only 15 to 20 minutes to get in, set up the robot, conduct the search, extract the robot, get off the rubble pile, and stand 250 to 300 feet away. By coincidence, this is the time allocated for the competition runs. The competition should reinforce the impact of logistics and transportation on the mission by requiring the entrants to walk to the course from the front of the building and set up all equipment as part of the competition round.

The normal standoff distance suggests that wireless robots operated from the forward station would be a good solution. However, there are three practical difficulties that need to be addressed. First, the wireless robots were too large to enter most of the voids; only once was a void large enough to send a wireless robot in. Second, the quality of wireless communications is still too uncertain. In one case, it was proposed that a wireless Solem be sent across the street to enter a ground-level void while a crane continued to remove rubble, and personnel were not allowed in. In this case, the options for moving the OCU, even with a directional antenna, to reacquire signal are limited, and the rescue task force was uncom-

> *The competition should reinforce the impact of logistics and transportation on the mission by requiring the entrants to walk to the course from the front of the building and set up all equipment as part of the competition round.*

Figure 1. The Rubble Pile at the World Trade Center.

Figure 2. The Red Section in the Competition.

fortable with fielding a robot that might stop next to the crane and distract the operator. When the crane stopped, and the Solem was deployed directly, it provided good information, but the video feed was intermittent, which would have wrecked havoc on image-processing algorithms. Eventually, the Solem lost communication when returning and had to be abandoned. Third, most of the voids explored were not ground level, and the robots did not have sufficient mobility to climb up and into the void, so it was necessary for a person to enter the rubble, if only for a short while. The lesson is that wireless is imperfect; therefore, approaches that assume sustained operation over long distances or through large amounts of rubble are not practical. The competition should have some method of randomly disrupting wireless communications to be more realistic.

The importance of accurately reporting the location of findings or mapping the robot's path in three dimensions cannot be overemphasized. Clearly in each case, the robot crews and task forces knew the location of the void entrance, but the entrance could be in the middle of four stories of rubble. It was necessary to know where the victims were relative to the entrance to make the decision to extract this area of a rubble pile or to even defer until later and work at an easier, more productive location. Consider the difference between going 30 feet down at a 10-degree slope and 20 feet down at a 45-degree slope. This aspect of the competition rules needs to be strengthened to encourage localization and mapping solutions.

Suggestions for Future Competitions

Although we have listed some observations about the appropriateness of the competition rules, there are some overall suggestions that might help the competition foster fieldable robots and algorithms.

First is to increase the fidelity of the test course to be more commensurate with USAR sites. One of the biggest weaknesses of the course is that it is well lit and open to provide audience viewing. In practice, the interiors of buildings are almost pitch black, and hanging obstacles are present. The yellow section of the National Institute for Standards and Technology course is intended to be representative of a semistructured building, but it is too open and not sufficiently cluttered with furniture and papers. Even though it was intended for the more "AI algorithm–oriented" entrants with nonfieldable robots, the course should introduce more realistic environmental conditions, particularly darkness and high density of clutter. Vision algorithms developed in controlled lighting simply will not transfer otherwise. In fact, the competition should rethink whether semistructured sections are for nonfieldable platforms or better used as representatives of the building search mission. In the latter case, ceilings and signs of structural instability would be critical. Although the orange section (confined space) is more three-dimensional than the yellow section, it is still too open. The stairs are wood, not the more challenging metal flashing, and there is no debris, slippery dust, or water to duplicate the residue of a collapse and sprinkler system release. Neither the orange nor the red section (rubble) has truly vertical and irregular passages or tunnels that lead to dead end (figure 2).

Second, the competition staging could be made realistic. For example, the teams could be isolated from the audience and commentator, truly working in a vacuum. In 2001, the teams could hear the announcers and often caught glimpses of the robot's situation from the large monitors on display to the audience. This approach was used to help the audience see the teams, as well as the course, and keep an entertainment flavor to the competition. Another model though is a TV reality show, where the audience can see simultaneous broadcasts of the robots and the remote team members. As noted earlier, teams should have to carry all their equipment in from a separate area. To introduce the cognitive stress, it would be good to have the rounds occur at random to duplicate the "hurry up and wait" effect of working with rescue teams.

The third suggestion is to offer both qualifying rounds and challenge rounds. Qualifying rounds would allow the organizers to better schedule the public rounds. A reoccurring problem has been teams that develop their entries on site. In the WTC, CRASAR team members had little facilities, time, or sleep to develop new solutions. More importantly, there is no place to adequately test such solutions. A simple programming error could cause rescue workers to reject a robot model, which prevented us from attempting to port software in the field; therefore, the competition should foster a conservative, disciplined development. Challenge rounds would allow teams that do not have the resources to develop a completely integrated system to demonstrate and competitively evaluate a key technology. For example, confined space navigation mapping and stair climbing, when the stairs are not coded and covered with slippery dirt and rubble, are two important mechanical abilities that could be evaluated in dedicated "minicourses."

Finally, it is important to note that few robot operators allowed on the rubble pile had complete safety gear and some form of basic safety awareness training prior to September 11. Only about half of the robot operators on site had these qualifications; fortunately, the USF team did and, as a result, was able to log the most time in the field of any group over the first 11 days of rescue. The lesson here is that even if you develop a useful USAR robot, it won't be used if you aren't certified for USAR safety by an officially recognized organization or if you can't quickly train a certified rescuer on how to use it.

Our experiences suggest a preliminary roadmap of technologies needed for victim detection. This roadmap can be considered as an evolution of increasing intelligence in mobility, sensing, mapping, planning, power management, communications management, and human-robot interfaces.

Roadmap for AI Development

Our experiences suggest a preliminary roadmap of technologies needed for victim detection. This roadmap can be considered as an evolution of increasing intelligence in mobility, sensing, mapping, planning, power management, communications management, and human-robot interfaces. We see eight levels of competence:

Robust teleoperation with basic mixed-initiative capabilities: Robots that can handle rubble and confined spaces are teleoperated. The operator handles all the control, mapping, and planning using the topological wall-following strategies developed by fire rescue workers. Sensor suites should be able to detect the basic affordances of a survivor: heat, motion, sound, and color. To be robust, the robot must have on-board intelligence that allows it to reestablish lost communications. The operator is responsible for estimating the remaining mission time based on power consumption and distance to exit. The user interface is visual and capable of displaying multiple sensors simultaneously.

Intelligent assistance: The next level is for the robot or workstation to actively aid the operator. The operator still directs and plans the robot's actions, but these are carried out under a guarded motion regime. The robot also cues the operator to signs of survivors, aids the operator in constructing and maintaining the topological map and location of victims, and estimates the time left in the mission before the robot must begin egress from the site. The user interface should now support views from other robots (for example, collaborative teleoperation).

Semiautonomous control: At this point, the robot is capable of autonomous execution of portions of the victim-detection script as well as automatic pose control of polymorphic platforms. Sensing is still cooperative with the operator, but the robot ensures that the search of a volume is complete and provides sensor fusion of cues. It also estimates the power availability for performing intensive tasks with margins for returning to the egress site. The interface displays adapt to the context and user preferences.

Victim assessment: Although navigation and victim detection have become fully automated, victim assessment is still cooperative. The robot can now carry and deploy radios or biosensors to leave behind. The user interface is now multimodal and doesn't rely only on visual displays.

Metric map making and planning: At this level, the basic control swaps from topological representations and maps to three-dimensional (3D) localization, metric maps, and optimal searches even in confined space and irregular rubble with nonhomogeneous materials.

Structural assessment: Building on the ability to make 3D metric maps, the robot is also able to add data about the volumes that allow the operator to characterize the structure and make decisions for victim extraction.

Adaptive shoring: In a natural expansion of the assessment task set, robots will selectively brace critical points in void space frameworks to prevent subsequent collapse in the presence of aftershocks, secondary device detonations, or other conditions that might threaten stability of damaged structures or rubble concentrations.

Trapped victim assistance: The continued trend in device miniaturization is expected to enable an infusion of medical technologies with the potential to significantly expand telemedicine capabilities into/onto mobile robot platforms designed for confined space access.[1]

Note

[1] For more information on the robots used at the WTC site, as well as video and stills, go to www.crasar.org.

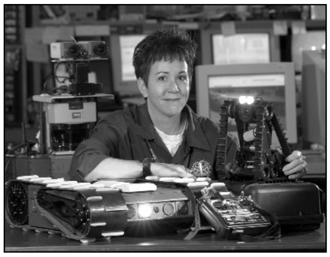

Robin Murphy is an associate professor of computer science and engineering at the University of South Florida in Tampa. She also serves as the deputy director of research for the Center for Robot-Assisted Search and Rescue (CRASAR) and on the Joint Executive Board for the National Institute for Urban Search and Rescue (NIUSR). Murphy is the recipient of the 2001 Eagle Award for her role in the response to the World Trade Center attacks in September 2001. Her e-mail address is murphy@csee.usf.edu.

John Blitch is currently serving as vice president and center director for SAIC's Center for Intelligent Robotics and Unmanned Systems (CIRUS), after retiring from federal service at the Defense Advanced Research Projects Agency (DARPA) where he was program manager for the Tactical Mobile Robots Program. He is also the director of the Center for Robot-Assisted Search and Rescue (CRASAR) and serves on the Joint Executive Board for the National Institute for Urban Search and Rescue (NIUSR). Blitch is the recipient of the 1997 High Lonesome Award for his graduate research in robot-assisted search and rescue and the 2001 Eagle award for his response to the World Trade Center attacks in September 2001. His e-mail address is CRASARHQ@aol.com.

Jennifer Casper is a master's student in computer science and engineering at the University of South Florida, conducting research in human-robot interaction in urban search and rescue (USAR). She is also cochair of the 2002 AAAI Mobile Robot Competition USAR event. Her e-mail address is jcasper@csee.usf.edu.