A Case Study of How Mobile Robot Competitions Promote Future Research

Jennifer Casper, Mark Micire, Jeff Hyams, and Robin Murphy

Computer Science and Engineering, University of South Florida
Tampa, Florida 33620
{jcasper, mmicire, hyams, murphy}@csee.usf.edu
http://www.csee.usf.edu/robotics

Abstract. The purpose of mobile robot competitions is to push the field of research and inspire future work. Our involvement in the Urban Search and Rescue event of the 2000 AAAI Mobile Robot Competition allowed us the opportunity to test our Urban Search and Rescue robot team in a standardized testbed. The result was ideas for future enhancements and research work. This article presents our experience as a case study on how mobile robot competitions promote research.

1 Introduction

One objective of mobile robot competitions is to push the field and inspire future research. The Urban Search and Rescue event of the 2000 AAAI Mobile Robot Competition was the first of its kind. It encouraged involvement from many different types of robotic systems. The competition platform was provided by the National Institute of Standards and Technology (NIST) and contained three different levels to allow for robots of varying capabilities. It is now being used as the basis for the RoboCup/AAAI Rescue Robot League.

The purpose of this paper is to present our experience as a case study on how competitions promote research. Section 2 presents the approach taken for the competition. Next, a description of the competition platform is discussed in Section 3. Section 4 discusses the results obtained from the competition. Enhancements to the current USAR Team and future research topics inspired by the competition are presented in Section 5. This paper does not address the utility of the NIST testbed and our recommendations for scoring USAR competitions; the reader is directed to [6] for such a discussion.

2 Approach

The approach taken for the USAR event of the 2000 AAAI Mobile Robot Competition was novel in two ways. First, the autonomy was implemented in the perception rather than the navigation. The second novel aspect was the idea of collaborative teleoperation. Implementing the autonomy in the perception was done in part by the recommendation of the Hillsborough County Fire Rescue

A. Birk, S. Coradeschi, and S. Tadokoro (Eds.): RoboCup 2001, LNAI 2377, pp. 123–132, 2002.

teams. The result was an Intelligent Assistant Agent which searched for victims based on cues (heat, color, motion, etc.) and informed the operator when a victim was located (or the high potential of a victim). This allowed the operator to concentrate on navigating the robot through the complex environment. The agent was adjustable in that the operator could specify which cues to look for and the ability to notify the operator (beeping in this case). The adjustable autonomy system is discussed in the next section, followed by the description of the Intelligent Assistant Agent. Collaborative teleoperation is presented in 2.3.

2.1 USAR Team

Fig. 1. USF USAR team: Fontana and Taylor.

The USF team used two outdoor robots: 1) a RWI ATRV with sonar, video, and a miniature uncooled FLIR (Fontana) and 2) a customized RWI Urban with a black and white camera, color camera, and sonars (Taylor). This was intended to be a marsupial pair, but the transport mechanism for the team was still under construction at the time of the competition. The USF team used a mixed-initiative or adjustable autonomy approach: each platform was teleoperated for purposes of navigation but the Intelligent Assistant Agent was utilized for autonomous victim detection.

Each operator controlled one robot through the use of dedicated operator control units (OCU), see Figure 2. There were two distinct OCUs. The OCU used for the Urban, Taylor, was designed and implemented by RWI with a proprietary interface using Mobility. It consists of a laptop, battery pack, and a wireless Ethernet in a carrying case. The OCU provided the user with a camera view, battery power levels, speed, inclinometer, compass, pose, and sonar. Taylor was

Fig. 2. The operator station setup during the AAAI conference.

controlled by moving the cross-hairs on the camera view using the touch pad. Her flippers were controlled similarly by cross-hairs on the power level display.

The OCU used for the ATRV, Fontana, was an in-house design using X-windows display widgets. It attempted to simplify the teleoperation control challenges imposed by the RWI OCU as well as permit the display of results from the Intelligent Assistant Agent. It consists of a laptop and wireless Ethernet. The operator was provided with sonar, an intelligent operator system, and keyboard control. The arrow keys were used to drive Fontana.

The robots communicated with the OCUs for the operators through a pair of Breezecom wireless Ethernet nodes and a base unit connected to a hub. The Intelligent Assistant Agent processing resided on Fontana, not on the OCU.

2.2 Intelligent Assistant Agent

The Intelligent Assistant Agent for this work consisted of one sub-agent, a vision agent to aid with victim detection. The vision agent performed cueing and behavioral fusion [8] based on the output of four concurrent detection algorithms: *motion, skin color, color different, and thermal region.* The user interface displayed the extraction results from each applicable algorithm, highlighted in color. Thus, the operator had visual cues for victims. In addition, the operator was notified by beeping when the agent had sufficient confidence in the presence of a victim based on fusion of multiple detectors. All active detectors were tiled side by side along with an un-augmented view from the video camera.

The vision agent can support any number of vision algorithms. The choice of these four algorithms was based on our understanding of the cues used by USAR technical rescue personnel and on the use of low computational affordances. A living victim will always radiate heat, which can be measured by a thermal imager. Heat is a primary affordance of victims in USAR. The only heat sources in the rubble should be people, animals, or indicators of a fire. The region properties of the heat source (e.g., does it look like the profile of a human) are not as critical. The rubble may insulate large portions of the victim, causing a more sophisticated (and computationally expensive) segmentation and matching algorithms to report false negatives (nothing is present).

Another affordance of a victim is color. In a collapsed structure, there is a grey dust from sheet rock or ceiling tiles which covers everything, and as such, most color is washed from the image. In a structure that has experienced a fire, the debris is similarly covered in black soot. Therefore rescuers look for anything that is distinguishable from that grey dust. For instance, a victim may move and shake off some of the dust, or there may be distinguishable colors such as blood.

Notice that the two previous victim detection affordances apply regardless if the victim were conscious or not. If the victim is conscious, an easily detectable affordance is motion, e.g., moving arms and waving for attention as the robot comes near.

Each of the four detection algorithms extracted a single feature that afforded the presence of a victim. The detection algorithms were written in C. Sufficient size of a feature or co-location of cues was interpreted by the vision agent as a possible victim. The operator was notified of the results of each algorithm through a continuously updating display, written in QT. The operator was also given the ability to switch on and off any of the algorithms. For further details of the Intelligent Assistant Agent, see [9].

2.3 Collaborative Teleoperation

The second novel aspect of the approach concerned the robot control used during the competition: *collaborative teleoperation*. Collaborative teleoperation is the use of a secondary robot to provide the operator with an external view of the primary robot. The external view is intended to help performance (e.g., navigate through a tight spot), diagnose and recover from a problem (e.g., high-centering of the robot, loss of a track, etc.), or cooperate (e.g., make an opening for the target robot to go through).

This task was performed by the operators who communicated with one another when help was needed. For instance, Fontana needed Taylor's point of view to assist in navigating around corners of the confined spaces (See Figure 3). Taylor had to be turned 180 degrees to face Fontana. Taylor's view could then be used by the operator to navigate the turn.

3 2000 AAAI Mobile Robot Competition

The USF USAR Team tested the NIST standard test bed for USAR as part of the 2000 AAAI Mobile Robot Competition [6]. The test bed consisted of three sections, each providing a different level of difficulty in order to accommodate most competitors. The easiest section, Yellow, contained mainly hallways, blinds,

Fig. 3. To navigate, Fontana's view and Taylor's view were needed in certain situations (Collaborative Teleoperation).

and covey holes to search through. The intermediate Orange Section provided more challenge with the addition of a second level that was reachable by stairs or ramp. Other challenges included those found in the yellow as well as some added doors. The Red Section was intended to be the most difficult. It contained piles of rubble and dropped floorboards that represented a pancake-like structure. The Orange and Red sections clearly required hardware that was capable of traveling such spaces. The USF Team concentrated on the more realistic Orange and Red sections.

The victims were mannequins and dolls. "Live" victims were adjacent to a heating pad to generate heat. Parts of mannequins were used to represent dismembered body parts and served as a source of false positives. A metronome type of device created a motion signature and a tape recorder playing simulated victims' cries for help were also utilized. The victims placed within the Yellow and Orange sections were surface victims [2]. The Red section victims were either in voids or entombed [2].

The USF Team ran through five events total; preliminary qualification, final competition, and three exhibitions. The ATRV found an average of 3.75 victims per each of the four runs recorded, while the Urban found an average of 4.67 victims. A fifth run was not recorded and no data is available. The data reported here is based on analysis of tape recordings from four of the five events made by USF, since the competition officials failed to record any data (number of identified victims, number of victims in each section, etc.).

4 Data and Results

Table 1 shows the performance of the USF heterogeneous team for victim detection. The data shows that in the best run (was the Final round, respectively),

Run	Time (min)	Total Victims Found	Robot	Victims Found Solo	Victims Found Jointly
Preliminary	32	4	Fontana	1	2
			Taylor	1	2
Final	30	7	Fontana	1	3
			Taylor	3	3
Exhibition #3	not recorded	6	Fontana	0	3
			Taylor	3	3
Exhibition #4	28	6	Fontana	1	4
			Taylor	1	4

Table 1. Results of victim detection in the four recorded runs.

3 of the 7 victims were detected jointly, or cooperatively. In the Exhibition #4 round, 4 of the 6 victims were detected jointly. This was due to two reasons. First, the more agile Urban robot, Taylor, could quickly enter an area and find a sign of survivors, but often could not confirm that the sign was a victim using only a black and white camera with a narrow field of view. Second, Taylor could not confirm whether a victim was "living" or "dead" if the victim was not moving; this had to be disambiguated with the FLIR thermal imager on Fontana. Fontana found at most 1 victim by herself, or solo, because she was larger, less agile, and had to slowly navigate the more restricted spaces of the Orange and Red sections.

Run	Time (min)	Robot	Seconds Spotting
Preliminary	32	Fontana	25
		Taylor	222
Final	30	Fontana	45
		Taylor	160
Exhibition #4	28	Fontana	10
		Taylor	250

Table 2. Results of collaborative teleoperation during the four recorded runs.

Table 2 shows that Fontana spent between 10 and 45 seconds and Taylor spent between 160 and 250 seconds acting as the secondary robot, providing an external view for the primary robot. The table indicates that both robots needed collaborative views at some point during each of the recorded runs (each robot spent some time spotting). As expected, the agile Urban, Taylor, needed much less help navigating, while the larger ATRV, Fontana, had to frequently rely on viewpoints from the Urban to assist the operator in guiding it through confined space. This is apparent by Fontana's low average spotting time of 26.7 seconds and Taylor's higher 210.7 second average spotting time.

5 Research Inspired by the Competition

Our experience with the 2000 AAAI Mobile Robot Competition sparked ideas for future enhancements on the current USAR Team and four research topics pertaining to USAR.

5.1 Enhancements for Robot Control

The collaborative teleoperation results in Table 2 show that between 11 and 16% of the total time of the competition run was spent cooperating (Taylor spotting Fontana or Fontana spotting Taylor). It was quickly realized what an inconvenience it was to have to move the robot to a viewing angle for the sake of assistance. Collaborative teleoperation shouldn't cost valuable search time. However it is necessary, as collaborating resulted in keeping the team progressing through the environment and at least 50% (on average) of the total victims found due to joint effort, see Table 2. Multiple lead robots may be helpful so that one specialized search robot may continue to explore while a designated and less equipped robot assists a larger robot, such as Fontana.

Looking up is commonly overlooked when searching. In the case of the competition, a victim on the second floor of the Orange section was clearly visible but slightly above camera range of the robots searching the first floor. Had the robots had the ability to tilt up, the team may have had a chance to find the victim. Stairwells and open balconies are common in apartment complexes and hotel buildings (urban areas where looking up and down are required).

While controlling the robot, the side sonars were providing displays to the user which were somewhat helpful in determining how close obstacles (e.g., walls) were. Yet between 11 and 16% of the total run time on average was still spent with both of the robots spotting and recovering from wall hand-ups, maneuvering mistakes, and course corrections (see Table 2). A logical addition to the teleoperation control would be an adjustable avoidance system, or "guarded motion", for navigation that would assist in keeping the robot from bumping into wall and other obstacles that might not be very easily judged by the operator.

The average 11-16% of the total run time spent by both robot spotting was needed in situations such as navigating a corner. The side sonars were utilized by the operators to help provide proximity information of the environment for navigating. Collisions, such as running into the wall or door frame, continued to occur. This leads to questioning the performance of the sensors, the sonars in this case, in confined space environments.

5.2 Enhancements for the OCU

The operators' feedback (collected from informal interviewing) pertaining to the OCU interface was negative in the effect that the OCUs were cumbersome. The keyboard and touch pad drive control interfaces for the ATRV and Urban were difficult to use in an environment like the mock-collapse NIST structure. The touch pad on the Urban's OCU disabled the ability to rotate the flippers and the drive forward at the same time. The touch pad pointer could only be on the drive display or the flipper display at any one time. The ability to rotate the flippers and drive forward is very useful in recovering the Urban when it is stuck on top of something (e.g., lodged on a pile of rocks). Another method of interface to the multiple functions of robots for the purpose of controlling is needed.

Another problem with the OCUs was the number of tasks the operator was required to perform at one time. The operator for Fontana was required to drive the robot, pay attention to the Intelligent Assistant Agent, and collaborate with the other operator. The OCU for the ATRV included a display panel for the vision detectors, each of which could be turned on or off by using the mouse to click a button in the X-window display. This diverted the operator's attention from controlling the robot, thus leaving the robot at a standstill, to alter what detectors are used. This results in a loss of time spent searching for victim, not to mention a break in the operator's focus.

The current implementation requires the operator to lean over and view the video display on the other robot's OCU in order to gain a second perspective (e.g., two views were needed to navigate the ATRV around tight corners). This was extremely cumbersome due to the physical separation of the two OCU displays. A method by which an operator may switch between views on the designated OCU is desirable.

5.3 Four Main Research Issues

Four main research issues emerged during the competition and exhibition. These contribute to the development of an adjustable autonomous heterogeneous robot team for urban search and rescue.

Collaborative Teleoperation As previously mentioned, collaborative teleoperation is the use of a secondary robot to provide the operator with an external view of the primary robot. Our experience with fielding USAR robots at the Hillsborough County Fire Rescue (Tampa, FL), SRDR (Miami, FL), Montgomery County Fire Training Academy (Rockville, MD), and George Air Force Base (Victorville, CA) training facilities is that it is very easy to flip or high center a robot in rubble. While it is easy to get the robot into trouble, it is difficult to diagnose the problem or construct a solution using only limited proprioceptive sensors (tilt, pose) or an exteroceptive view of the environment. What is needed is an exproprioceptive view, where the robot's pose is known relative to the external environment. Since the robot cannot see itself, the viewpoint of another agent is helpful. In the case of the competition, a situation occurred in which the view from the ATRV was able to prevent the Urban falling from an elevated position. The operator needed the exproprioceptive view from the ATRV in order to save Taylor as Taylor's viewpoint was worthless in this case. Collaborative teleoperation is a function that could be supported by an Intelligent Assistant Agent.

Persistent Pursuit The main idea of persistent pursuit is to obtain and maintain an ideal vantage point of the daughter or team members in order to provide the operator with multiple usable views. For instance, in the red section of the USAR mock up, the team entered the opening on the east side of the structure. Without much maneuverability, Fontana remained stationary until Taylor began to traverse alongside the back of the rubble pile. It was worth the risk to abandon Taylor momentarily in order to follow the outer wall of the structure in the direction Taylor was heading in hope of acquiring a better vantage point. Plexiglas windows outlining the back of the course provided an opportunity to view Taylor, who continued to rummage around the backside of the pile. A certain amount of persistence was needed in order to find a better point of view. If the Plexiglas hadn't been there, the operator would've retreated to the last known position the Urban was seen at. This concept is useful not only for teleoperation or navigation, but for line of sight communications as well.

Semantic, Systematic, and Opportunistic Search Current search methods include semantic, systematic, and opportunistic [7]. Any one used on its own is insufficient to complete this task. A semantic search involves searching the interesting areas for victims, but does not guarantee a thorough search. The systematic search is a way to more thoroughly investigate a site. An example of a simple systematic search is the right wall follow which is often utilized by the Hillsborough County Rescue teams in low visibility situations. Just utilizing a wall follow might be considered insufficient because it will search the interesting as well as uninteresting areas, thus costing precious time. An opportunistic search is being aware of opportunities that may present themselves (i.e., a victim is discovered on the way to a goal). An optimized search of an USAR site would require the ideal combination of the three search types.

In the case of this competition, the teleoperators were utilizing a right hand wall search and tended to any victims found along the way. The teams hadn't been given any semantic information, such as what kind of building this is, or who was in it. The best solution was to systematically search and hope to opportunistically find victims.

Conflict between Navigation and Sensing There existed moments during the competition in which the navigation and sensing conflicted. For example, Taylor sensed a victim and wanted to get within a certain distance from the victim in order to provide communications but couldn't due to being obstructed. She would have to abandon her view in hopes of finding an alternate route to the victim. In a purely reactive system, her victim find behavior and avoid behavior would conflict causing her to stand still and do nothing. This type of conflict needs to be resolved in order to properly automate victim searching in USAR.

6 Conclusions and Current Work

The purposes of robot competitions are to promote research and push the field forward. The first Urban Search and Rescue Competition was held at the AAAI 2000 Conference in Austin, Texas and intended to do just that. It was successful in promoting ideas and research for future work in USAR in our experience. The competition inspired us to implement and test the ideas we had developed from our work within the USAR field. USAR training (trench rescue, confined space, and structural collapse), collaboration with rescue professionals, and research has provided information as to what is required of potential USAR robots.

The competition was a chance to test and compare USAR systems in a standard and reproducible platform. This is an important step up from testing in a sterile lab, but not as effective as an authentic USAR environment in our experience. The competitions will give the opportunity for the comparison of different USAR systems on a common accepted platform. It will continue to be advantageous to strive to make more authentic USAR platforms for competitions.

Acknowledgments

Portions of this work were funded by DARPA and a travel scholarship from AAAI. The authors would like to thank Mike Larson of IR Cameras for the loan of the Indigo Alpha FLIR, and Chief Ron Rogers and Jeff Hewitt from the Hillsborough County Fire Department for their support. Without them, our work wouldn't be possible.

References

1. National Fire Academy, *Rescue Systems I,* 1993.
2. United States Fire Administration, *Technical Rescue Program Development Manual,* 1996.
3. National Fire Protection Association, *Standard on Operations and Training for Technical Rescue Incidents,* 1999.
4. J. Hyams, M. W. Powell, R. R. Murphy, "Cooperative Navigation of Micro-Rovers Using Color Segmentation," *Autonomous Robots*, Vol. 9, pp. 17-16, 2000.
5. J. L. Casper, M. J. Micire, R. R. Murphy, "Issues in Intelligent Robots for Search and Rescue," *SPIE Ground Vehicle Technology II*, 2000.
6. R. R. Murphy, J. Casper, M. Micire, J. Hyams, "Assessment of the NIST Standard Test Bed for Urban Search and Rescue," *NIST Workshop on Performance Metrics for Intelligent Systems*, 2000.
7. R. R. Murphy, "Biomimetic Search for Urban Search and Rescue," *IROS*, 2000.
8. R. R. Murphy, *Introduction to AI Robotics*, MIT Press, Cambridge, MA, 2000.
9. R. R. Murphy, J. L. Casper, M. J. Micire, J. Hyams, "Mixed-Initiative Control of Multiple Heterogeneous Robots for Urban Search and Rescue," *In review, IEEE Transactions on Robotics and Automation.*
10. K. Osuka, "Rescue Robot Contest in RoboFesta-KANSAI 2001," *Proceedings of 2000 IEEE International Conference on Industrial Electronics, Control and Instrumentation*, pp. 132-137, 2000.