

The Find-Life-on-Mars Event

Reid Simmons

■ The Find-Life-on-Mars event of the 1997 American Association for Artificial Intelligence Mobile Robot Competition and Exhibition featured robots trying to find and collect stationary and moving colored objects in an arena littered with real rocks. The 2-day event had 11 entries participating in both single-robot and multirobot categories, both with and without manipulators. During the event, many of the robots successfully demonstrated object recognition, obstacle avoidance, exploration, and the collection and depositing of objects.

The Mars Pathfinder Mission, featuring the Sojourner rover, was the inspiration for the Find-Life-on-Mars event. The general concept was to have the robots locate, collect, and deliver a variety of “life forms,” including both stationary and moving objects. Technically, the event was designed to highlight mobile manipulation, object recognition, exploration, and obstacle avoidance in a relatively unstructured environment. The event was viewed as the successor to similar collection-type events in past American Association for Artificial Intelligence (AAAI) robot competitions, in particular the office cleanup and clean up the tennis court events (Kortenkamp, Nourbakhsh, and Hinkle 1997; Hinkle, Kortenkamp, and Miller 1996; Simmons 1995). The main differences were to be the environment (rock strewn) and the need for more sophisticated object recognition (differentiating colors and shapes of objects).

The Find-Life-on-Mars event was held in a 30-foot-diameter hexagonal arena, surrounded by meter-tall gray plastic walls. Two thousand five hundred pounds of real rocks (painted black to aid in visual recognition) were distributed around the arena, and a “Mars lander” was placed near the center (figure 1). The rocks were unevenly distributed—one half of the arena was sparsely populated with rocks, and the other half was significantly denser. In addition, in some trials, black paper, representing danger zones, was spread out over part of the floor.

The lander was a square cardboard box with two swinging doors on opposite sides, one painted orange and the other blue. The life forms consisted of balls and cubes of various bright colors and squiggle balls, which can move erratically on their own. The robots were to explore the arena, pick up life forms, and deposit them in the lander. Moving objects were to be placed in one of the lander’s doors; the other door was for the stationary objects.

Points were awarded for picking up objects of a specific type (ball, cube, or moving squiggle ball) and specific color (figure 2). Additional points were awarded for placing an object in the lander. Penalty points were deducted for colliding with rocks, placing an object in the wrong door, and traveling within the danger zones. We also specified penalties for modifying the lander, although no group took advantage of this option. No other modification of the environment was allowed.

One of the main problems confronted by the Find-Life-on-Mars rules committee (Tom Henderson, Doug MacKenzie, and myself) was determining how to handle entries of various capabilities. In particular, some robots did not have manipulation capability but could track and recognize objects. There was also a question of how to compare single-robot and multirobot entries. In the end, we decided to have three categories, which were separately judged and had separate awards: (1) multirobot entries with manipulators, (2) single-robot entries with manipulators, and (3) single-robot entries without manipulators. The nonmanipulator robots had to stop within about six inches of an object (except for the moving balls) and announce the type of object seen, at which point one of the judges would pick it up so that it would not be identified again. Many groups participated in more than one category. For instance, some groups with multiple robots ran one of them in a single-robot category. In all, there were 11 separate entries from 7 groups (University of Arkansas, Brandeis University, Brown University, Colorado School of Mines, Georgia Institute

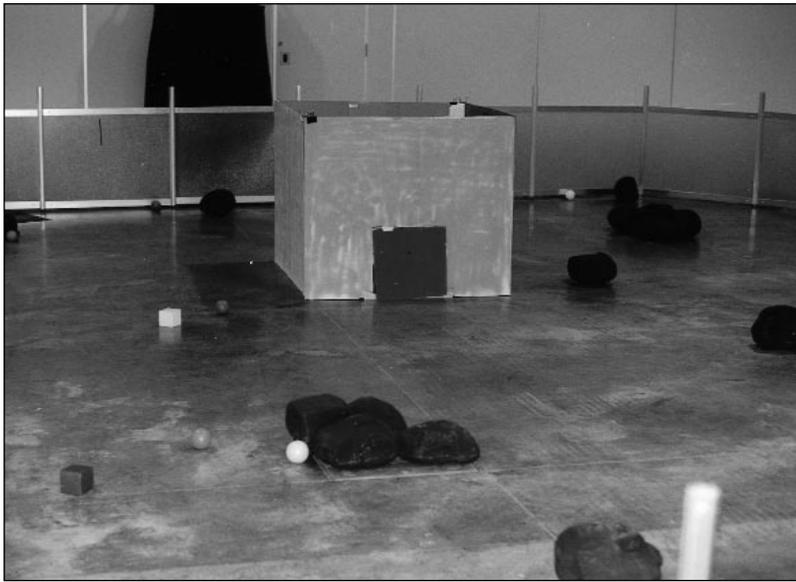


Figure 1. The "Martian" Environment.

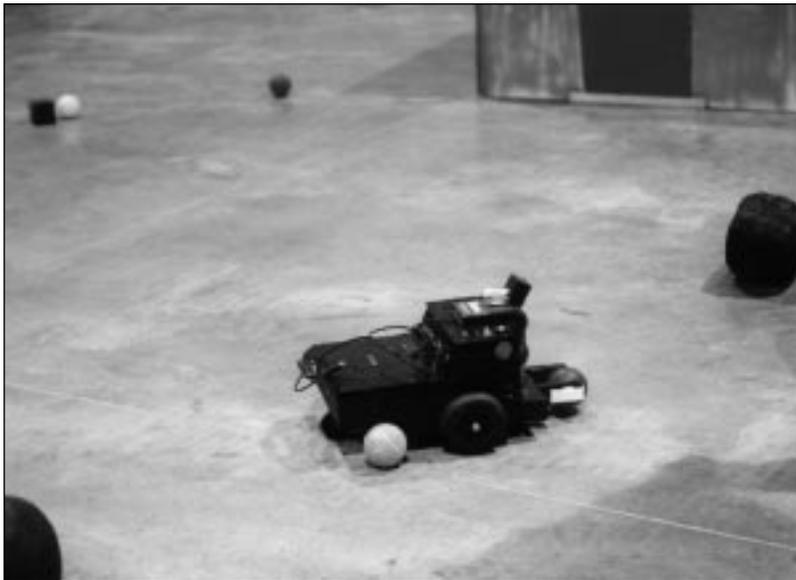


Figure 2. A Robot from Brandeis University Collecting an Unknown "Life Form."

of Technology, McGill University, and University of Minnesota).

The overall Find-Life-on-Mars event was divided into a challenge round, where each entry participated in two separate trials, and a single-trial finals round, which was videotaped and open to the public. Each trial lasted 10 minutes. The idea behind the challenge round was to push the robots to their limits. We had anticipated running the first trial of the challenge round at a moderate degree of difficulty (more objects in the sparse half of the arena;

only one color for each type of object) and then making the second trial harder, with the degree of added difficulty based on how well the robots did in the first round.

Unfortunately, the first trial proved difficult enough for most of the robots. Thus, for the second trial of the challenge round, we made only minimal changes: We added two black paper danger zones near the perimeter of the arena and used two separate colors for each type of object (four different colors in all). Extra points were awarded for picking up the first object of a given color and type (for example, the first red cube got more points than each subsequent red cube). This approach was meant to encourage exploration, but apparently no group tailored its strategy to take advantage of the point differential. In fact, it seems that most groups did not use any type of systematic exploration, preferring to have the robots wander around the arena until they spotted an object, then approaching it directly. In part, this strategy was successful because the arena was relatively small, making it easy for the robots (especially the multirobot entries) to cover the arena fairly thoroughly in 10 minutes using a simple random-walk strategy.

Although the rules and scoring details were fairly well set in place before the competition began, we made one major rule change during the competition itself. The original rules stated that a robot could not be touched, except to remove it from the arena. However, robots became trapped frequently enough that we relaxed this restriction after the first set of runs. Instead, we levied a 40-point penalty whenever a judge was asked to move a robot (we gave the same penalty to groups that needed to straighten bent manipulators during the course of a trial). Needless to say, the judges (Jim Hendler, Sridhar Mahadevan, and myself) were kept busy freeing trapped robots, keeping track of the objects collected by multiple robots, picking up objects in the nonmanipulator trials, and generally staying out of the way of the robots (figure 3)!

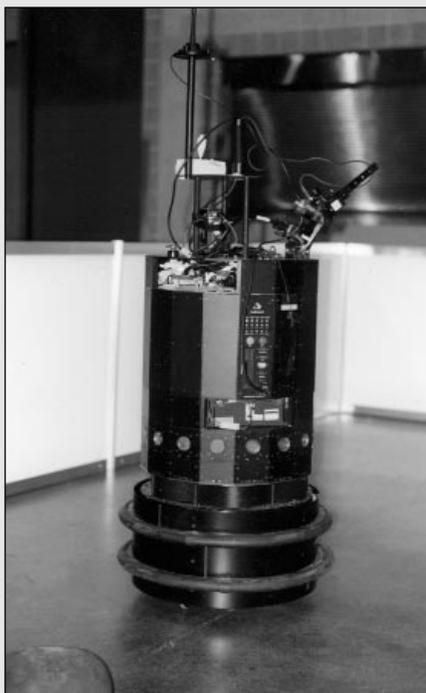
To give participants (and the judges) some time to enjoy the conference, both trials of the challenge round were held on the same day. Because of the large number of runs that had to be held that day, we maintained a strict schedule: Entries were assigned a 20-minute slot, and they had to run their 10-minute trial within the allotted time (however, people were usually willing to switch time slots to help out groups that needed a bit more time to get ready). Of the 22 scheduled runs in the challenge round, only 5 runs were not attempted, mainly because of hardware problems. Every

Profile of a Winner: McGill University

Much of the work at the McGill Mobile Robotics Lab concerns computational problems related to use of sensors: vision, laser, and sonar. As a result, the McGill team entered the nonmanipulator category. The team was made up of four students: Francois Belair, Eric Bourque, Deeptiman Jugessur, and Robert Sim, with myself acting as faculty mentor. The robot they used was a NOMAD 200, a chest-height cylindrical robot with a three-wheeled synchrodrive, a standard ring of 16 sonar sensors, and a single-color camera mounted on a simple pan-and-tilt unit (see figure). Although planning for the competition started early, most of the key software for the competition was custom designed in the 6 to 10 weeks before the conference. Early in the design process, several existing research tools and subsystems were used, but as the system developed, most preexisting code was either heavily modified or replaced. These revisions took place to maximize efficiency and reliability as well as to reduce dependence of subsystems not fully under the control of team members (in the last weeks, the team members worked on an almost 24-hour schedule and were unable to tolerate delays that would have been incurred by having to wait for other people to modify their subsystems). Even at the time, it was apparent that the use of special-purpose modules would limit the reusability of the code and incur other disadvantages, but it became an inevitable necessity.

The design of the McGill entry was very loosely based around the McGill mobile robotics architecture, a software interconnection methodology. The key software modules, instantiated in the form of UNIX, were for planning and collision avoidance, scheduling and error recovery, and user interface and diagnostics. The operation of some these modules is as follows:

The basic visual-perception module dealt with using color to classify and segment the image. It used a multidimensional lookup table to map pixels to specific objects (such as rock, floor, or target). In its initial form, it was able to compute the color distribution of groups of pixels, but this approach proved to be excessive for the



McGill Entry in the Find-Life-on-Mars Event.

types of object that were in the actual environment. The classifier could be trained by showing it samples of objects of interest along with their correct labeling. Once pixels were classified, the vision module proceeded to group them into blobs that could efficiently be described by polygons of limited complexity. The three-dimensional position of these obstacles was computed from knowledge of the camera geometry and a flag ground-plane assumption. These labeled polygons were then transmitted to the mapping and planning module.

The mapping and planning module dealt with the maintenance of a long-term map of the environment, collision avoidance, path planning, and target acquisition. The map was composed of obstacles observed from either sonar or vision. Older objects were gradually removed to account for uncertainty growth as a result of dead-reckoning errors. Path planning was accomplished by computing an obstacle-free convex region about the robot for simple short-range planning, combined with long-range planning that directed the robot within this "safe polygon." Although

several long-range path-planning modules were developed, random motion driven by directly observed targets proved sufficient in the final competition.

The object-recognition module that was used matched already-classified blobs to images of known objects using subspace projection. To improve the performance of the recognizer, the planning module attempted to approach targets to get them into a standard viewing position. This type of technique can be used to recognize complex objects, but it can have difficulty with illumination variations, which proved to be a challenge until the last minute.

One of the greatest tests of the competition was coping with issues of system integration and robustness. In fact, a last-minute communications problem led to the robot occasionally becoming blind, with potentially disastrous consequences. In the last days before the finals, the students worked a straight 40-hour debugging stretch, a frighteningly common behavior at the AAAI competitions. A more subtle issue is the standard question of how to reconcile conflicting input in a complex AI or robotic system. For example, during the preliminaries, the vision module managed to detect a spectator's colored clothing through a semitransparent partition, producing an internal inconsistency that incapacitated the robot. As Rob Sim put it, the robot "chose to go after an audience member's shoes and barreled full speed into the wall, while the unfortunate target leapt back in unmasked terror."

It is fair to say the competition served to illustrate the importance of some underappreciated issues in our lab as well as extensively test some software and hardware modules (some of the work has also played a role in ongoing thesis research). The drive and resourcefulness of my students, and the others in the competition, is a refreshing and wonderful aspect of the event for those who have a direct involvement. I believe the participants, by virtue of their increased appreciation of a range of issues, are much enriched for the experience.

– Gregory Dudek

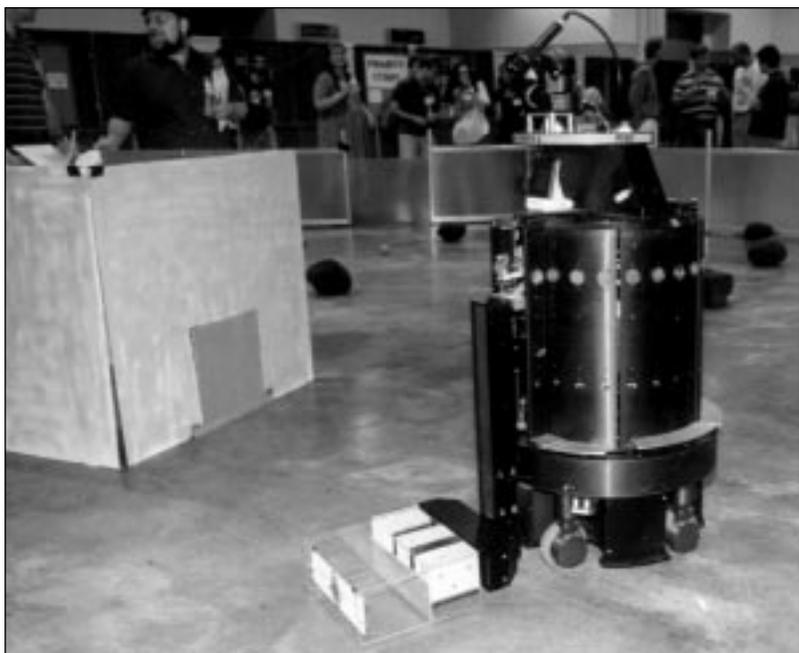


Figure 3. Keeping Out of the Way of Brown University's Robot.

entry competed in at least one of the two trials. This degree of reliability is actually quite an improvement over past competitions. Part of the reason for the increased reliability is that more groups are using commercially available hardware, and part is due to the generally improved state of the art in mobile robotics.

In the single-robot, nonmanipulator category of the challenge round, the entry from McGill University placed first, Brown University was second, and Colorado School of Mines was third (Brandeis and the University of Arkansas participated in this round but were judged not to be sufficiently successful to be awarded a place showing). In the single-robot, manipulator category, the entry from the University of Minnesota placed first, Georgia Institute of Technology was second, and Brown participated but did not place. In the multi-robot category, the order was reversed: Georgia Tech first, University of Minnesota second, with Brandeis participating but not placing.

The results of the challenge round showed that many of the robots needed some fine tuning to be truly competitive. There were the usual hardware problems (for example, grippers not strong enough to withstand the occasional bump against rocks or the lander). Mainly, however, the problems were related to software. In particular, the distribution of rocks proved problematic for many of the entries. For some of the smaller robots that relied on sonar for obstacle avoidance, many

of the rocks were below the height of the sonars. Although the larger robots (figure 3) tended to fare better in detecting the rocks because they used color vision rather than sonar, they often had problems maneuvering between the rocks. To alleviate both these difficulties, we simplified the environment in the final round by piling up many of the rocks and separating them more widely (figure 1).

Another problem that commonly occurred was that a robot would turn toward an object that it noticed, bump into a nearby rock, turn away from the collision, see the object again, and then repeat the process, indefinitely. Although one or two groups used a timeout behavior that would prevent infinite cycles, more often than not a robot would remain stuck until either a team member asked a judge to move it (for a penalty score), or a squiggle ball came along and knocked the object away, at which point the vicious cycle would be broken fortuitously.

Several of the manipulator robots also had difficulty in returning to the lander and depositing the collected objects in it. Often, the problem was that the robot's dead reckoning, especially orientation, was so bad that it completely lost track of the location of the lander. Even when it found the lander, it often had difficulty putting the objects into the relatively small doors. In some runs, a robot would pick up an object almost immediately and then spend most of the remaining 10 minutes futilely trying to deposit it. In addition, robots often spent valuable minutes trying to place an object in the correct door. It turned out that the more successful strategy was to put objects into the closer of the two doors and take the penalty points if it turned out to be the wrong side because the opportunity to collect more objects outweighed the loss of points.

After a day off to attend talks (well, let's be honest, to hack!), the participants reconvened for the finals round, which consisted of a single 10-minute trial for each entry. Changes to the environment included consolidating the rocks (as described previously) and removing the danger zones. As with the second trial of the challenge round, we used two different colors for each object type. We also allotted only 15 minutes per slot because of the pressures of ending in time for the public award presentations and team photographs.

All entries were eligible to compete in the finals round, regardless of performance in the challenge round. Of the 11 separate entries, 8 managed to participate in the finals round. In the single-robot, nonmanipulator category, McGill University was a convincing winner

Profile of a Winner: Georgia Tech

The Georgia Institute of Technology earned three first-place finishes in the Find-Life-on-Mars event at the 1997 American Association for Artificial Intelligence (AAAI) Mobile Robot Competition and Exhibition. Its two robots, LEWIS and CLARK, won the multiagent challenge and finals rounds for robots with manipulators. The student team of robot builders chose a multiagent approach for the reliability and efficiency it offers over single-agent solutions. The advantage of teamwork was demonstrated in competition when CLARK's arm was ripped off midway through the challenge round in an engagement with one of the rock hazards. Fortunately, LEWIS survived and collected enough "life forms" to win the round.

LEWIS and CLARK are Nomadic Technologies NOMAD 150 robots (shown in the photograph). The NOMAD 150 is a 3-wheeled kinematically holonomic vehicle equipped with a separately steerable turret, 16 ultrasonic range sensors, and a ring of rubber bump sensors. Georgia Tech modified the robots to add servodriven grippers and real-time vision. The vision system reports the location of colored life forms, rock hazards, and delivery bins at as much as 30 Hertz. A JAVA-based control system running on a laptop computer communicates with the vision and mechanical control systems using serial protocols.

Control systems for the robots were coded using CLAY, a set of JAVA classes that support sequenced behavior-based control (Balch 1997). Complex behaviors are developed using behavioral primitives called *motor schemas*, independent processes that combine to generate an overall behavior (Arkin 1989). Motor schemas take input from specialized perceptual schemas and generate a movement vector representing the desired direction of travel. The relative importance of each schema is encoded with a gain value. The vectors of active motor schemas are multiplied by their gain values, summed, then normalized and transmitted to the hardware for execution.

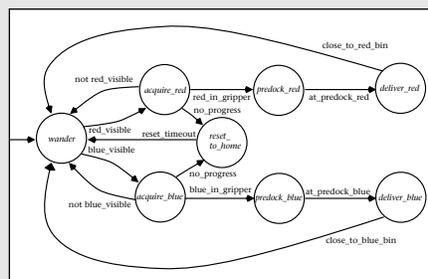
As an example of how behaviors were developed for the Find-Life-on-Mars event, the robots activate the *move_to_red_bin*, *avoid_obstacle*, and *noise* motor schemas for navigation to a red bin. This assemblage



LEWIS and CLARK with Several Members of the Georgia Tech Robot-Building Team.

From left to right: Tucker Balch, Juan Carlos Santamaria, and Tom Collins. (Photo courtesy of Stanley Leary, Georgia Institute of Technology.)

of primitive behaviors moves the robot toward the bin but keeps it from colliding with obstacle hazards. Noise helps move the robot out of any local minima it encounters. In this manner, behaviors were developed for each stage of the task, for example, *wander*, *acquire_red*, *acquire_blue*, *predock_red*, *predock_blue*, *deliver_red*, and *deliver_blue*. The control systems sequence from one behavior to another based on perceptual cues provided by the sensors. The robots begin their search for Martians using the *wander* behavior. When a red life form is detected, the robots transition to the *acquire_red* behavior and, after grasping the object, switch to *predock_red*. The *predock* behavior draws the robot to a position in front of the delivery bin, and *deliver* is used to finally place the object in the bin. A similar sequence is provided for blue life forms. The overall sequence of behaviors is illustrated in the figure that follows.



Behavioral Sequence Used by LEWIS and CLARK for Collecting and Delivering Objects in the Find-Life-on-Mars Event.

The behaviors were tested in simulation, then on robots in the Mobile Robot Laboratory. At the AAAI competition, the Georgia Tech team planned to use the laboratory-developed behaviors as is, but lighting, floor coloring, and the paint used on the rock hazards caused unexpected perceptual difficulties. The floor of the arena included black splotches that were sometimes confused with rocks. The spectrum of light in the arena, in combination with the paint used on the rocks, caused the robots to occasionally mistake rocks for blue Martians. The perceptual difficulties were compounded by the fact that the hazards were too low to be detected by the robots' sonar-ranging sensors. These perceptual difficulties led CLARK to scrape a rock hazard, causing it to lose its arm in the challenge round. Between the challenge and the finals rounds, the hazard-detection problem was solved: The ultrasonic sensors were reaimed downward at a 45-degree angle. Hazards could then be detected reliably.

A change in the task for the final round presented a new challenge. The robots had to collect and deliver Martians painted six different colors instead of only two, but the robots' vision systems can only track three colors at a time. The use of multiple robots enabled a workaround: Each robot was programmed to specialize in the collection of three of the six types of Martian. The improved hazard sensing and refined vision strategy enabled the robots to collect 10 attractors and place 9 of them in the correct delivery bin.

Acknowledgments

Georgia Tech's team was led by Tucker Balch. Darrin Bentivegna and David Huggins programmed the vision hardware, and Harold Forbes built the grippers. Erwin Oei helped with the construction of the grippers and their repair at the contest. Tom Collins served as the team's faculty adviser. We thank Ron Arkin (director of Georgia Tech's Mobile Robot Laboratory) for use of the robot hardware and Nomadic Technologies for providing servocontrol software.

– Tucker Balch

(see sidebar), with Brandeis University placing second and University of Arkansas third. In the single-robot, manipulator category, Georgia Tech was a decisive first (see sidebar), with the University of Minnesota second. The results were the same in the multirobot, manipulator category; in addition, Brandeis participated, but did not place, in the multirobot category.

Both the changes the judges made to the environment and the changes the participants made to their algorithms combined to produce markedly better performances. Almost every entry performed better in the finals round. In particular, each first-place entry scored almost as many, or more, points in the one trial as they did in the two challenge-round trials combined. Particularly apparent was the robots' improved abilities to find and approach objects while they avoided obstacles, something that appeared to be more problematic in the earlier round.

One surprising result is that in all the trials, but especially in the finals round, the single-robot manipulator entries performed nearly as well as, and in some cases better than, the multirobot entries. Given this result, it might not be necessary to have separate single-robot and multirobot categories for future collection-type events (although multiple robots might have made a bigger difference if the arena were large enough such that no one robot could cover it in the allotted time).

One general observation that we can take from this event is that the state of visual object recognition, at least for bright, uniformly colored objects, is reasonably good. However, navigation (obstacle avoidance and position estimation), especially when combined with mobile manipulation, is less robust in "natural" environments. The problems that entries had in navigation and localization are in contrast to the reasonably good performances in previous AAAI competitions, where the collection-type events were held in office-type or obstacle-free environments. Although the state of the art in mobile robotics continues to improve, there is definitely significant work ahead of us if we want to be able to "find life" autonomously by the time the next Mars rover lifts off in 2001.

Acknowledgments

I would like to thank Tom Henderson and Doug MacKenzie, my fellow rules committee members, and Jim Hendler and Sridhar Mahadevan, my fellow judges, for all their help in running a successful competition event. Ron Arkin and Jim Firby provided much needed support (and rocks). However, most

especially I'd like to thank the seven participating teams that spent hundreds of hours getting the robots to work so well, both before and during the competition.

References

- Arkin, R. C. 1989. Motor Schema-Based Mobile Robot Navigation. *International Journal of Robotics Research* 8(4): 92-112.
- Balch, T. 1997. CLAY: Integrating Motor Schemas and Reinforcement Learning, Technical Report, GIT-CC-97-11, College of Computing, Georgia Institute of Technology.
- Hinkle, D.; Kortenkamp, D.; and Miller, D. 1996. The 1995 Robot Competition and Exhibition. *AI Magazine* 17(1): 31-45.
- Kortenkamp, D.; Nourbakhsh, I.; and Hinkle, D. 1997. The 1996 AAAI Mobile Robot Competition and Exhibition. *AI Magazine* 18(1): 25-32.
- Simmons, R. 1995. The 1994 AAAI Robot Competition and Exhibition. *AI Magazine* 16(2): 19-30.



Reid Simmons is a senior research scientist in the School of Computer Science at Carnegie Mellon University. He earned his Ph.D. from the Massachusetts Institute of Technology in 1988 in AI. Since coming to Carnegie Mellon in 1988, Simmons's research has focused on developing self-reliant

robots that can autonomously operate over extended periods of time in unknown, unstructured environments. He has helped develop nearly a dozen autonomous robot systems and has been involved in many of the American Association for Artificial Intelligence robot competitions as both participant and organizer. His e-mail address is reids@cs.cmu.edu.



Gregory Dudek is an assistant professor with the School of Computer Science and a member of the McGill Research Centre for Intelligent Machines and an associate member of the Department of Electrical Engineering at McGill University. He directs the McGill Mobile Robotics Laboratory. His e-mail address is dudek@cim.mcgill.edu.



Tucker Balch will join Carnegie Mellon's Computer Science Department as a postdoctoral researcher this fall. Balch's work focuses on multirobot cooperation, learning, and diversity. He is also interested in making robotics research more easily accessible through the development of a

freely available JAVA-based multirobot simulation system. His web page address is www.cc.gatech.edu/~tucker.