AI Magazine Volume 17 Number 1 (1996) (© AAAI)

The 1995 Robot Competition and Exhibition

David Hinkle, David Kortenkamp, and David Miller

■ The 1995 Robot Competition and Exhibition was held in Montreal, Canada, in conjunction with the 1995 International Joint Conference on Artificial Intelligence. The competition was designed to demonstrate state-of-the-art autonomous mobile robots, highlighting such tasks as goal-directed navigation, feature detection, object recognition, identification, and physical manipulation as well as effective human-robot communication. The competition consisted of two separate events: (1) Office Delivery and (2) Office Cleanup. The exhibition also consisted of two events: (1) demonstrations of robotics research that was not related to the contest and (2) robotics focused on aiding people who are mobility impaired. There was also a Robotics Forum for technical exchange of information between robotics researchers. Thus, this year's events covered the gamut of robotics research, from discussions of control strategies to demonstrations of useful prototype application systems.

This article describes the organization and results of the 1995 Robot Competition and Exhibition, which was held in Montreal, Canada, on 21 through 24 August 1995, in conjunction with the 1995 International Joint Conference on Artificial Intelligence (IJCAI-95).

The 1995 Robot Competition and Exhibition was the fourth annual competition in a string of competitions that began at the Tenth National Conference on Artificial Intelligence (AAAI-92) in San Jose, California. During this inaugural competition, 10 robots searched for and approached a set of tall poles in a large arena (Dean and Bonasso 1993). The next competition, held at AAAI-93 in Washington D.C., saw robots participating in events that involved maneuvering around an office-building layout and moving large boxes into patterns (Konolige 1994; Nourbakhsh et al. 1993). The third competition, held at AAAI-94 in Seattle, Washington, included such events as office-building navigation and trash pickup (Simmons 1995). The fourth competition (the first at an IJCAI conference) built on the successes of the previous competitions.

The goals for the competition and exhibition have remained the same over the years. The first goal is to allow robot researchers from around the world to gather (with their robots) under one roof, work on the same tasks, and exchange technical information. Second is to assess (and push) the state of the art in robotics. Third is to contrast and compare competing approaches as applied to the same task. Fourth is to provide a public forum in which robotics research can be seen by the AI community, the media, and the general public. Because of this broad range of goals, determining an appropriate format for the competition and exhibition was a challenge.

After much discussion with robotics researchers, we decided on a format similar to the previous year's. There were four basic parts: (1) a formal competition with fixed tasks and scoring, (2) a wheelchair exhibition (a new addition) in which the results of mobile robotics research can be shown in a practical application, (3) a robot exhibition in which researchers can display robotics research that is not applicable in the competition, and (4) a forum in which all participants and invited researchers can discuss the state of robotics research.

The formal competition was itself divided into two different tasks: The first task involved navigating within an office-building

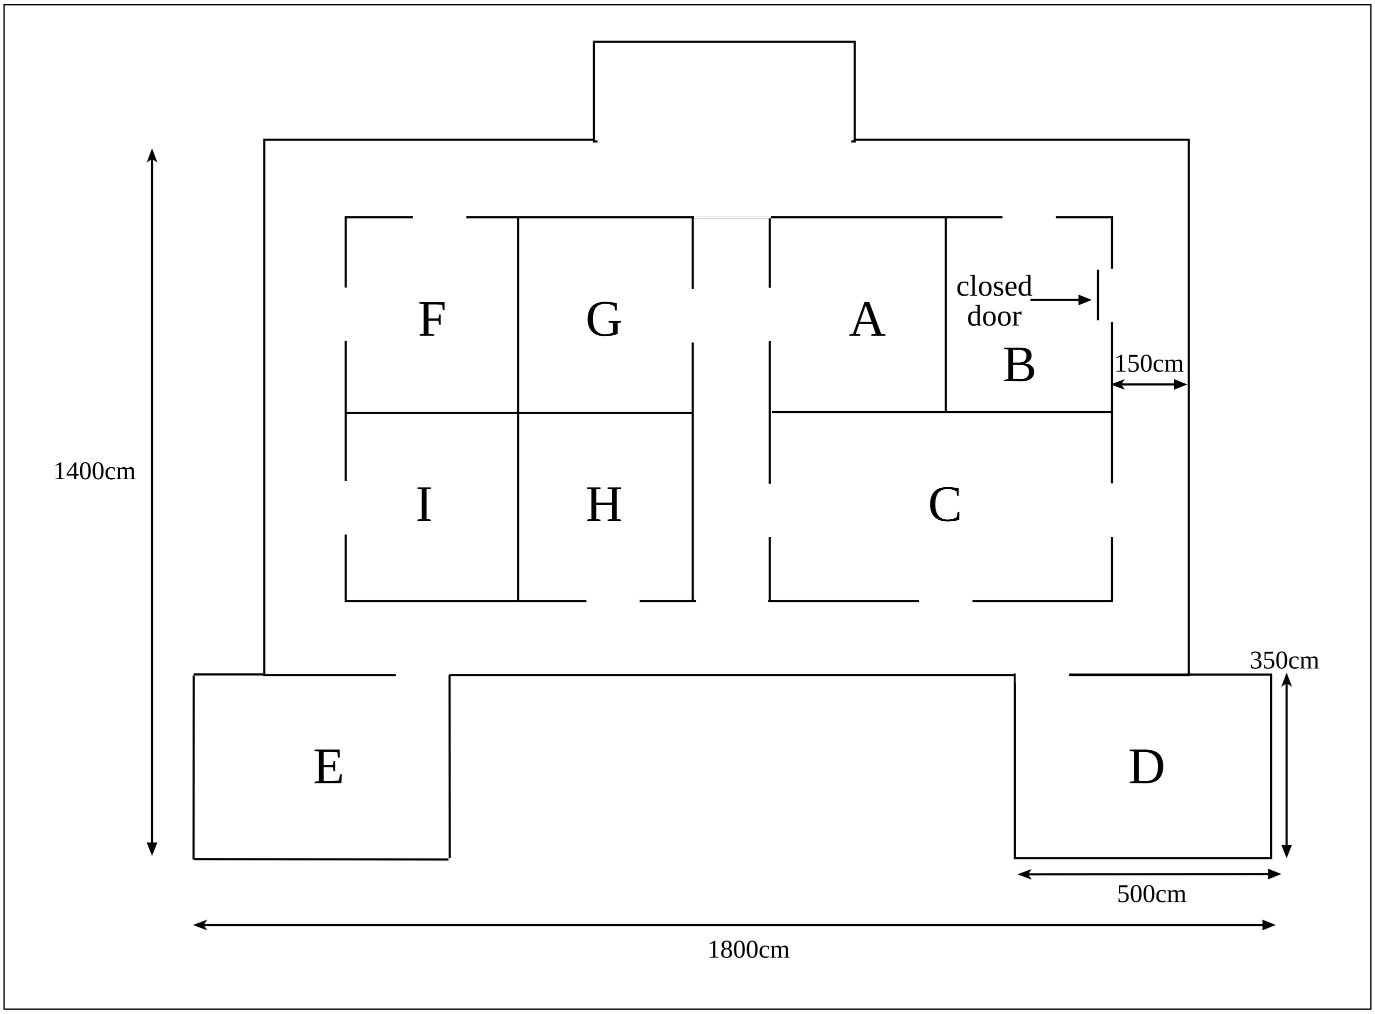

Figure 1. The Competition Arena.

A closed door (room B) is inset from the walls.

environment using directions entered at run time by a human. The second task involved distinguishing between trash and recyclables and depositing each in the correct receptacle. For each task, teams had two preliminary trials in which to demonstrate to the judges their ability to perform the task and a final trial in which teams competed against each other for points. Each task is described in detail in the following sections, and results are given. Then, the wheelchair exhibition, the robot exhibition, and the forum are discussed in turn.

Robot Competition Event 1: Office Delivery

In addition to the traditional goal-directed navigation, the first event of the robot competition was designed to promote the ability of robots to detect when there is a problem and ask for help through effective robot-human interaction. The event took place in a re-creation of a typical office environment with partitioned offices (figure 1). The robots were asked to follow a series of instructions that told them to which room they were supposed to go. The instructions consisted of statements such as “exit the room and turn left,” “go to the end of the hallway and turn right,” “enter the room on your third right,” and so on. (See figure 2 for the actual instructions for a trial in event 1). The robots would then proceed to search the start room for the door to the hallway, exit the room, and then follow the instructions to the goal room. However, sometimes the instructions that were given by the human contained an error. The given instructions would not lead to the goal room—or to any room in fact! The

Original Instructions

Start	Room	E
Exit	Room	Left
Turn	First	Right
Turn	Third	Right
Go	Past	Lobby
Turn	***Third***	***Right***
Turn	Third	Right
Enter	First	Left

Corrected Instructions

Turn	First	Right
Turn	Third	Right
Enter	First	Left

Figure 2. Instructions for Getting from Room E to Room D in Event 1.
The italicized instruction is incorrect.

robots were required to monitor their progress as they carried out the instructions, detect the error, and request further assistance (new instructions) when an error was detected. The corrected instructions would be given to the robot, which would then proceed to the goal room, enter, and announce that it had arrived.

The robots were awarded points based on completion of the task and the time it took to complete. Exiting the room was worth 20 points, detecting an error in the instructions was worth 40 points, and entering the goal room was worth 40 points. The total points for completing the task was 100. The time taken to complete the task (in minutes) was subtracted from the maximum time points of 25; so, if a robot took 5 minutes to complete the task, it would receive 20 points (25 points – 5 minutes). Extra points were awarded based on human-robot communication. Robots received 10 bonus points for effective communication and an additional 10 points for accepting the instructions verbally. Points were deducted for marking objects such as doors or hallways, at one point for each marker. Penalties were assessed for any mistakes the robots made. Points were deducted if the robot became confused and required assistance. The number of points varied between –5 and –40, depending on how many times assistance was required. However, only half as many points were deducted if the robot realized that it had become confused and requested the assistance. Also, in the event of a collision with a stationary obstacle or wall, 30 points were deducted.

This event was similar to last year's Office Delivery event but with greater emphasis placed on recovery from mistakes and human-robot interaction. At the third robot competition in 1994 (Simmons 1995), instead of being given just a set of directions to follow, the robots were given a topological map showing all the connections between hallways and rooms. Only one team (Nourbakhsh, Powers, and Birchfield 1995) was able to complete the event in 1994.

Results

This year, three teams were able to successfully complete the event: (1) Korea Advanced Institute of Science and Technology (KAIST), (2) North Carolina State University (NCSU), and (3) Kansas State University (KSU). A fourth entry, from the University of New Mexico (UNM), was damaged in transit and was unable to compete. The results of the final round of competition were

1. CAIR-2 (KAIST), 130.5 points
2. LOLA (NCSU), 114.0 points
3. WILLIE (KSU), 70.5 points
4. LOBOT (UNM)

Teams

Now we briefly describe each team's robot and its approach to the event.

CAIR-2—Korea Advanced Institute of Science and Technology CAIR-2 is a custom-built prototype robot, developed at KAIST. It uses sonar and infrared sensors and has a stereo vision system with pan and tilt. It has a custom-built voice-recognition system and speech synthesizer. Although its native language is Korean, it spoke perfect English throughout the competition. The control architecture was designed to combine both behavior-based and knowledge-based approaches. Major components of the navigation system are the collection of behaviors (such as find-a-door, go-forward, and avoid-obstacles), a high-level task executor, a fuzzy-state estimator, information extractors, and a coordinator for behaviors.

CAIR-2's basic strategy was to scan the room looking for a special marker on the doorway. The doorways for the start and goal rooms were marked with a special symbol to assist in detection. Once the robot detected the doorway, it proceeded to exit the room, keeping its two cameras focused on the doorway marker. The human-to-robot instructions were entered verbally. As each instruction was entered, the robot would repeat it back for confirmation. If the robot misunderstood an instruction, the speaker would supply a corrected instruction.

CAIR-2 consistently performed well throughout the trials and in the finals. In the finals, CAIR-2 exited the start room in about a minute and took a total of 5 minutes and 25 seconds to complete the event. Because the instructions were given to CAIR-2 verbally, supplying the corrected instructions to the robot took an extra minute or two, which was, of course, not counted as part of the running time. For more information on CAIR-2, see "CAIR-2: Intelligent Mobile Robot for Guidance and Delivery," also in this issue.

LOLA: North Carolina State University LOLA is a NOMAD 200 robot equipped with a pan-tilt–mounted color camera and sonar sensors. LOLA's on-board computation includes a dual C40 digital signal–processing (DSP) image processor and a 486 DX2-66 running LINUX. Communication with the robot for delivery of the directions and feedback from the robot was done by radio Ethernet. The control architecture comprises four modules: (1) *state-set progression* for establishing a probabilistic framework for feature-based navigation in a topological space; (2) *feature detection* for identifying doorways, hallways, and so on; (3) *low-level motion control* for determining the direction and velocity of the robot as well as performing obstacle avoidance; and (4) *registration* for determining the direction of the hallway from sonar data.

LOLA's basic strategy was to maneuver from the initial starting position toward the center of the start room. This strategy afforded the robot a better position from which to scan the room for the exit door. Each doorway was marked with a colored circle to assist in detection. Once the robot had detected the doorway of the start room, it proceeded to exit into the hallway and get its bearings by aligning itself with the hallway walls.

LOLA performed well throughout both the finals and the trials. The scores for both the trials and the finals were virtually identical. In the finals, LOLA exited the start room in about 1-1/2 minutes and completed the event in 6 minutes and 15 seconds. For more information on LOLA, see "LOLA: Probabilistic Navigation for Topological Maps" and "LOLA: Object Manipulation in an Unstructured Environment," also in this issue.

WILLIE: Kansas State University The KSU team used a NOMAD 200 robot (WILLIE) from Nomadic Technologies. The robot was equipped with two sonar rings. The robot relied on a radio Ethernet to communicate between the control program running on a workstation and the actual robot.

The basic strategy was navigation using sonar widths to position the robot and identify doors in the hallway. A subsumption architecture, with threads running the sonar-detection and sonar-avoidance routines, was used. For example, the exit-room strategy first involved finding a wall; then, the robot did wall following while a separate thread detected doorways. When a door was found, the robot aligned itself on the door and exited the room.

The performance during the first trial run was good. WILLIE exited the start room in 1-1/2 minutes, detected the error in the human-supplied instructions, accepted the new corrected instructions, and completed the task within 5 minutes. However, during the second trial run, because of radio interference, the performance of the radio Ethernet degraded severely.

Continuing in the contest required porting the 10,000 lines of the control program from the UNIX workstation to the 486 processor on board the robot. It also meant installing LINUX and a thread package on the 486 processor. The porting took 12 hours and involved retuning the control program to account for differences in timing related to running directly on the robot instead of on the workstation.

The final run was not as successful as the team would have liked. Some of the timing problems caused the robot to miscount a doorway and caused a collision with the wall. However, the team finished with a sense of accomplishment and a desire to prepare for the 1996 competition.

LOBOT: University of New Mexico The UNM LOBOT is a custom-built mobile robot designed by UNM engineering students. The LOBOT is driven by a two-wheel differential configuration with supporting casters. It is octagonal in shape, stands about 75 centimeters tall, and measures about 60 centimeters in width. Sensing is achieved using a layout of 16 ultrasonic transducers. The on-board distributed processing and control system consists of a 486 PC-based master and networked MC68HC11 slaves.

The LOBOT uses an object-oriented behavioral approach based on a task decomposition of event 1 (that is, exit-room, navigate-hallway, enter-room). It is written in C++ and uses commercial software for generating speech and recognizing high-level verbal instructions. The design philosophy views the LOBOT as a collection of objects that cooperate to achieve goals. In this hierarchical architecture, three main objects exist at the highest level: (1) navigation; (2) the helm; and (3) object detection, recognition, and ranging (ODRR). The *helm* and the *ODRR* encapsulate the usable resources of the LOBOT, such as motor control and sensor data. *Navigation* is a client of the ODRR and the helm, using them for perception, sensor-based decision making, and the initiating of appropriate actions. The motion behaviors are implemented as objects within navigation that operate using an artificial force-field approach. In particular, the exit-room behavior is based on an augmented force-field approach adapted from Kweon et al. (1992). That is, actual sonar-range data are augmented with additional virtual-range data that collectively act as a virtual wall behind the robot. Repulsive forces from the virtual wall combine with actual range forces to push the robot toward free space. Starting from any initial location in a single-exit room, these artificial forces eventually flush the robot out through the exit. For robustness, a doorway-traversal behavior is also activated in the vicinity of the exit.

Unfortunately, the LOBOT was damaged during shipping and was unable to compete. However, the team vowed to try again next year.

Robot Competition Event 2: Office Cleanup

The second event of the robot competition was designed to promote interaction between mobile robots and their environment. The event took place in room C of the competition arena (figure 1). The exits from this room were blocked and, on the floor, were empty soda cans and empty Styrofoam coffee cups. Also in the room, in the four corners, were trash and recycling bins (two of each). The task had the robots pick up the soda cans and deposit them in the recycling bin and then pick up the Styrofoam cups and deposit them in the trash bin. Scoring was based on the number of objects picked up and correctly deposited within 20 minutes. Penalties were assessed for modifying the environment or colliding with obstacles.

In designing this event, the competition organizers wanted to promote the research area of mobile manipulation. Although virtual manipulation was allowed (that is, the robot could approach the trash and announce that it was picking up the trash without actually doing so), the penalty was severe enough that only one team used this approach. All the other robots had some form of manipulation, which is a large improvement over last year's competition in which a similar event attracted only two teams that performed actual manipulation (see Simmons [1995]).

The competition organizers also wanted to promote the use of computer vision to distinguish between different objects and then have intelligent-control software make decisions based on these perceptions. Thus, the robots needed to recognize two classes of object (trash and recyclables) and two classes of container (trash bins and recycling bins). The classes of container were marked with the symbols *T* and the recycling closed circle for trash and recyclables, respectively (figure 3). Again, there was an advancement over last year's event, where there was only one type of trash and only one type of receptacle.

We saw a vast improvement over last year in the performance of robots in this event.

Results

We saw a vast improvement over last year in the performance of robots in this event. All the robots successfully used vision to distinguish among the objects. All the robots successfully distinguished between the bins and navigated to the correct bin for each object. The final results reflected the differences in manipulation and object recognition. The

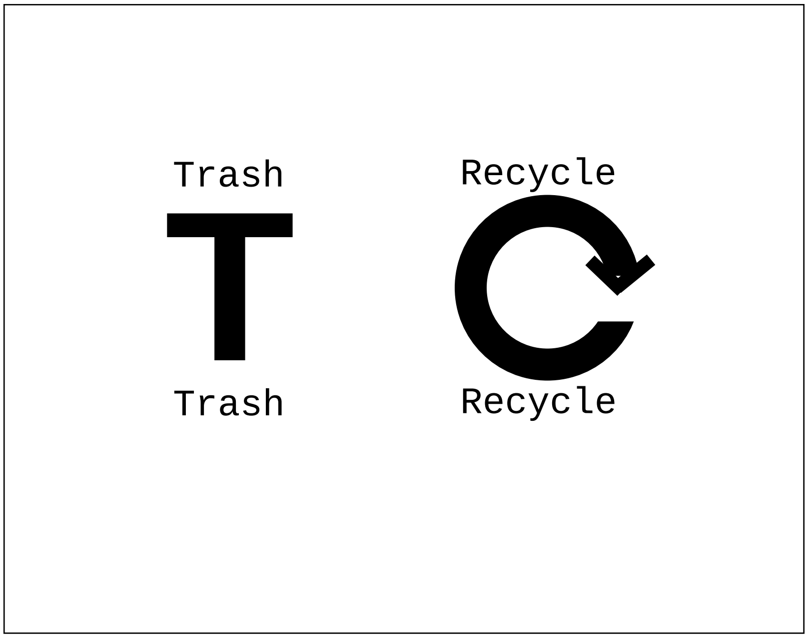

Figure 3. Symbols for the Trash Bin (left) and Recycle Bin (right).

top robots (see related articles in this issue) actually picked up the objects and deposited them in the correct bins. The next tier of robots did not pick up objects but pushed them next to the appropriate trash and recycling bins; they received fewer points than those robots that actually picked up objects. One team used virtual manipulation. Also, several teams modified the trash or the bins and received penalties for doing so. The final results were

1. LOLA (NCSU), 196 points

2. CHIP (University of Chicago), 165 points

3. WALLEYE (University of Minnesota), 135 points

4. NEWTON 1 & 2 (Massachusetts Institute of Technology [MIT] and Newton Labs), 105 points

5. CLEMENTINE (Colorado School of Mines), 79 points

Teams

Now we briefly describe each team's robot and its approach to the event.

LOLA: North Carolina State University LOLA is a NOMAD 200 robot equipped with a prototype Nomadics manipulator, pan-tilt–mounted color cameras, and sonar sensors. LOLA's on-board computation included a dual C40 DSP image processor and a 486 DX2-66 running LINUX. LOLA's basic methodology was as follows: LOLA first locates trash using predefined color-histogram models of the different types of trash and histogram back projection. Second, LOLA heads off in pursuit of the trash. Third, during pursuit, LOLA tracks the centroid of the trash as it moves down the image plane and utilizes a nonlinear least squares algorithm to calculate its location relative to the robot. Fourth, once within range, LOLA grasps trash using position estimation.

Once LOLA grasps a piece of trash, it looks for the appropriate receptacle and deposits the trash using the same method just described. The trash can and recycle bin are distinguished by a color marker at the base of the receptacle (pink for trash, yellow for recyclable).

LOLA performed well in the preliminary round, depositing 13 objects. In the final round, LOLA was again performing well until optical sensors on its prototype arm started to give false readings. It was later determined that the optical sensors were being triggered by the audience cameras! However, by this point, LOLA had already deposited enough trash to win the event. In the final round, LOLA correctly deposited seven objects. For more information on LOLA, see "LOLA: Object Manipulation in an Unstructured Environment" and "LOLA: Probabilistic Navigation for Topological Maps."

CHIP: University of Chicago CHIP is a small mobile robot built on a Real-World Interface B12 base. It has a single arm for manipulating objects in the world, sensors on the gripper for detecting touch and pressure, and eight sonar sensors to help with obstacle avoidance. CHIP's primary sensor is stereo, color vision. CHIP is controlled by the animate agent architecture. The low level of the architecture is a collection of soft real-time routines that can be rearranged into different control loops at different times. At a high level, the reactive action package (RAP) system manipulates the set of routines running at any given time to create a sequence of control states to accomplish a specific task.

CHIP systematically searches any area by recording where it looks in a global frame relative to its "wake-up" position. It always looks in nearby unsearched areas simply by panning when possible but moving around as needed. CHIP can recognize a broad class of small objects when seen against a relatively clean background. It segments an edge image into regions of possible objects and, for each segment, computes the size, aspect ratio, edge density, average color, fraction of white, and contour regularity. The resulting feature vector is classified against a set of fuzzy exemplars by choosing the nearest neighbor within a maximal distance.

CHIP steadily improved over the two preliminary rounds and into the finals. In the initial preliminary round, CHIP was only able to correctly deposit one object. In the second preliminary round, CHIP deposited three objects. Then, after many hours of late-night hacking, CHIP really shined in the finals, giving LOLA a run for her money by depositing four objects. For more information on CHIP, see "Programming CHIP for the IJCAI-95 Robot Competition," also in this issue.

WALLEYE: University of Minnesota The chassis of WALLEYE is built from an inexpensive radio-controlled car with the body shell of the car and the original electronics removed and replaced by specially designed boards. All boards are built around the 68HC11 microcontroller and have, at most, 16K of EPROM (electronically programmable read-only memory) and 32K of random-access memory. WALLEYE uses three microcontrollers, one for the main board, one to control the motor, and one for the vision system. The vision system uses a CCD (charge-coupled device) chip with digital output, a wide-angle lens, and a frame grabber board on which all the vision processing is done. The images are 160 x 160 pixels, with 256 gray levels. The camera can grab as many as 20 frames a second. WALLEYE has a gripper with a beam across the fingers to detect when something has been grasped. The gripper cannot lift objects, only hold them.

The basic strategy of WALLEYE was to look around for a cup or a can. When a blob that could correspond to such objects is found in the image, WALLEYE starts driving toward it, tracking the object while approaching. If the object seen was not really an object, the illusory object disappears while it is being tracked. Then, WALLEYE again starts its search for another object. Tracking an object is, in general, easier than finding it and much faster. When WALLEYE gets close to an object, the beam in the fingers is broken, signaling the presence of something between the fingers. To guarantee that the object is indeed a cup or a can, WALLEYE backs up and verifies that the fingers are still holding on to the object. In this way, WALLEYE will not confuse legs of chairs or other immovable objects with trash. Once an object has been grasped, WALLEYE looks for the appropriate trash bin and pushes the trash in front as it moves toward the receptacle. When WALLEYE gets within the trash zone, it lets go of the trash, depositing it. It then continues the search for more trash.

WALLEYE performed well in the first preliminary round, pushing 11 objects to the correct bin. Reflections in the floor because of an overnight cleaning were responsible for a subpar performance in the second round, where WALLEYE only deposited three pieces of trash. In the finals, WALLEYE's performance improved somewhat, and it pushed four objects. Because it could not place the objects in the bins, only near them, it did not receive as many points for each object as LOLA or CHIP, landing it in third place. However, WALLEYE showed that the task can be performed with extremely limited computing power under a variety of environmental conditions.

NEWTON 1 and 2: MIT and Newton Labs The MIT–Newton Labs entry in the trash-collection contest was two vision cars. The vision car uses an off-the-shelf remote-control car as its robot base. A small vision system (the COGNACHROME vision system made by Newton Research Labs) and a color camera are mounted on the car. The COGNACHROME vision system includes a 68332-based processor board and a custom video-processing board. The video-processing board takes NTSC (National Television Systems Committee) input from the color camera, digitizes the signal, and classifies the pixels on the basis of color. This board sends a 1-bit signal for each color channel to the 68332 board. (The system allows three color channels, although only one was used for the contest.) The 68332 board processes this signal to find all the objects of the specified color in the scene; it processes these data at 60 Hertz and uses the results to control the car. The camera is the only sensor on the car.

The COGNACHROME vision system includes software for tracking objects on the basis of color. Beyond this software, about 65 lines of C were written for the low-level interface to the car and about 300 lines of C to implement the control software for the contest. There were two cars in the contest. Each car focused on one color. Trash was colored blue, and recyclables were colored orange. The trash and recycling bins were goals in the corners, with blue or orange swatches above them. In the competition, the cars moved randomly until they saw a piece of trash and a goal of the same color in the same image. The car would then move toward the trash and push it toward the goal.

The MIT–Newton Lab robots were entered less than 48 hours before the competition, but they proved to be crowd pleasers because they moved quickly, and they moved constantly. The robots knocked cans and cups into the bin area with great force. They sometimes rammed into each other and even tipped over. At the end of their frenzied activity, the robots managed to push four objects near the correct bin in the first trial and five objects near the correct bin in the final round. Because the team modified both the bins and the objects and did not place the objects in the bins, they received fewer points for an object than other teams ahead of them.

CLEMENTINE: Colorado School of Mines

CLEMENTINE is a Denning-Branch MRV4 mobile robot with a ring of 24 ultrasonic sensors, a color camcorder, and a laser-navigation system. CLEMENTINE was the only entry without a manipulator. CLEMENTINE is controlled by an on-board 75MHz Pentium PC. The team consisted of four undergraduate computer science students who programmed the robot as part of their six-week senior practical design course. The team took a behavioral approach, focusing on the issues of recognition, search, and sensor fusion.

CLEMENTINE began the task by systematically looking for the red regions using the color camcorder. If a red region was close to the appropriate size of a can seen from that distance, CLEMENTINE would move to the can and ask a helper to pick up the can. CLEMENTINE would then continue to look for another can, to a maximum of three. If CLEMENTINE did not find another can, it would go to the nearest recycle bin, drop off the can (again, asking a helper to deposit the can), and then return to the center of the ring and scan for more trash. CLEMENTINE used its laser-navigation system, which triangulated its position from three barcode–like artificial landmarks. It also knew a priori where the trash bins were.

The trash-recognition process was successful, and in a preliminary round, CLEMENTINE detected all 10 cans, depositing 7 of them. In the second round, CLEMENTINE deposited nine cans. However, the algorithm was sensitive to lighting changes and, in the final round, deposited only seven cans, tying the number deposited by the first-place team. However, because CLEMENTINE was performing virtual manipulation, each object was worth fewer points.

Wheelchair Exhibition

A robotic wheelchair exhibition was added to this year's event to demonstrate how the robotics technology that has been developed over the last several years could be applied successfully. Many people are mobility impaired but are unable to safely operate a normal power wheelchair (Miller and Grant 1994). This year's exhibitors concentrated on supplementing the control system of a power wheelchair to endow it with some semiautonomous navigation and obstacle-avoidance capabilities. The chairs, of course, also had to be able to integrate continuous human commands as well as follow their programmed instructions.

Three chairs with automatic guidance systems were brought to IJCAI-95. The NAVCHAIR from the University of Michigan has been under development as a research project for several years. WHEELESLEY from Wellesley College and TAO-1 from Applied AI Systems were both built for this event. PENNWHEELS, from the University of Pennsylvania, was also exhibited. PENNWHEELS uses an innovative mobility system but does not have any guidance system (in fact, it is tethered to its power supply and computer).

Although there was no formal contest for the chairs, a wheelchair limbo contest was held. The contest consisted of the chairs automatically aligning and passing through a continually narrowing set of doorways. Although the NAVCHAIR and WHEELESLEY use totally different sensors, both were able to go through narrow doorways, and both got stuck at the same point (when there were less than two inches of clearance on a side). TAO-1 was demonstrated successfully but suffered an electronics failure during some maintenance right before the limbo contest.

NAVCHAIR

The NAVCHAIR assistive-navigation system is being developed to provide mobility to those individuals who would otherwise find it difficult or impossible to use a powered wheelchair because of cognitive, perceptual, or motor impairments.

By sharing vehicle-control decisions regarding obstacle avoidance, safe-object approach, maintenance of a straight path, and so on, it is hoped that the motor and cognitive effort of operating a wheelchair can be reduced.

The NAVCHAIR prototype is based on a standard Lancer powered wheelchair from Everest and Jennings. The Lancer's controller is divided into two components: (1) the *joystick module*, which receives input from the user through the joystick and converts it to a signal representing desired direction, and (2) the *power module*, which converts the output of the joystick module to a control signal for the left- and right-wheel motors. The components of the NAVCHAIR system are attached to the Lancer and receive power from the chair's batteries.

The NAVCHAIR system consists of three units: (1) an IBM-compatible 33MHz 80486-based computer, (2) an array of 12 Polaroid ultrasonic transducers mounted on the front of a standard wheelchair lap tray, and (3) an *interface module* that provides the necessary interface circuits for the system. During operation, the NAVCHAIR system interrupts the connection between the joystick module and the power module. The joystick position (representing the user's desired trajectory) and the readings from the sonar sensors (reflecting the wheelchair's immediate environment) are used to determine the control signals sent to the power module.

During the course of developing NAVCHAIR, advances have not only been made in the technology of smart wheelchairs but in other areas as well. Work on the NAVCHAIR has prompted the development of an obstacle-avoidance method, called the *minimum vector field histogram* (MVFH) *method* (developed by David Bell). MVFH is based on the *vector field histogram algorithm* by Borenstein and Koren (1991) that was originally designed for autonomous robots. MVFH allows NAVCHAIR to perform otherwise unmanageable tasks and forms the basis of an adaptive controller.

A method of modeling the wheelchair operator, stimulus-response modeling to make control-adaptation decisions, has also been developed and experimentally validated as part of the research on the NAVCHAIR. Current work focuses on using probabilistic reasoning techniques from AI research to extend this modeling capability (Simpson 1995).

WHEELESLEY

Robotics researchers do not often discuss user interfaces when explaining their systems. If they do, it is usually in terms of a programming interface. However, when we move from autonomous robots to wheelchair robots, we need to carefully consider the user interface. A robotic wheelchair must interact with the user and must do it well. The user should control the wheelchair system, not be controlled or constrained by it.

Unlike other wheelchair robots at the workshop that used a joystick as the sole interface, WHEELESLEY's user has the option of interacting with the robot through a joystick or the user interface. The joystick mode is similar to the other teams' joystick modes, so only the user interface is discussed here.

The user interface runs on a MACINTOSH POWERBOOK. Although the input to the interface is currently through the touch pad and button, a system could be built on top of this interface to customize the system for the user. Some wheelchair users have some upper-body control, but others need to use a sip-and-puff system. Some users can use voice; others cannot. The interface that was shown at IJCAI-95 is general and would have to be tailored to the needs of specific users.

WHEELESLEY's interface provides information while it allows the user to control the system. The user can track the speed of the wheelchair and can set the default speed of the wheelchair. (The *default speed* is the maximum traveling speed when no obstacles are present.) For users who are unable to turn their heads to see obstacles, a map of the wheelchair shows where obstacles are present. The interface allows the user to switch between *manual mode* (no computer control), *joystick mode* (navigation using the joystick with computer assistance), and *interface mode* (navigation using the interface with computer assistance).

The system was demonstrated at IJCAI-95. WHEELESLEY was the only system that could drive through doorways without being steered by a human in the chair.

The wheelchair was built by the KISS Institute for Practical Robotics. Software and user interface development was done by a team of five undergraduates at Wellesley College, supervised by Holly Yanco.

WHEELESLEY was the only system that could drive through doorways without being steered by a human in the chair.

TAO-1

The autonomous wheelchair development at Applied AI Systems, Inc. (AAI), is based on a behavior-based approach. Compared to more conventional AI approaches, this approach allows greatly increased performance in both efficiency and flexibility. In this approach, the concepts of "situatedness" and embodiment are central to the development of the autonomous control system. *Situatedness* emphasizes the importance of collecting information through sensors directly interfacing the real world, and *embodiment* stresses the significance of doing things in physical terms in the real operational environment. The robustness and graceful degradation characteristics of a system built using the behavior-based approach also make it attractive for this development.

The base wheelchair used for the current implementation of the autonomous wheelchair (TAO-1) is produced by FORTRESS of

Quebec. The electronics and control mechanics that came with the wheelchair were left intact. In fact, the chair can still be operated using the joystick; the user can override the autonomous control mode whenever he/she wishes.

The control system for the autonomous wheelchair developed at AAI is based on a Motorola 68332 32-bit microcontroller (a single-chip computer with on-chip memory and control electronics). It has AAI's own multitasking, real-time operating system that allows the controller to receive real-time signals from a large number of sensors, and it sends control output to two motors to drive the left and the right wheels. It looks after both forward-backward and left-right movements of the chair.

This year's robot exhibition was an extraordinary crowd pleaser because all the robots that were demonstrated were highly interactive with the audience.

Two color CCD cameras mounted on the chair detect free space and motion as much as 10 meters in front of the chair. Six active infrared sensors detect obstacles in close vicinity, to one meter from the chair. The signal from the cameras is processed by an intelligent vision-processing unit that is also built on behavior-based principles. The control program for all the vision processing occupies 9.5 kilobytes (KB), and the other behavior control occupies 2.75 KB. This program is significantly smaller than similar vision-based control programs operating in a real environment implemented using conventional AI methods.

Development is expected to continue in a staged approach. We are now in the first phase, the safety phase, where work concentrates on improved ability to avoid obstacles and dangerous situations in the environment. On completion of this phase, the chair will be mobile without hitting any objects or other moving things while it avoids common pitfalls that currently require human attention.

In the future, the vision system will be capable of detecting many other things, such as landmarks found in the path of the wheelchair, unusual appearances in the pavement, and traffic signals. The number of infrared sensors will be increased to allow it to move in more confined spaces. Later phases will develop sophisticated interactions between the human and the chair, improve mobility aspects of the chair, and introduce evolutionary computation methodologies to facilitate the chair adjusting to the needs and the situations of each individual user.

PENNWHEELS

PENNWHEELS is a prototype mobility system under development at the University of Pennsylvania. The robot uses two motorized wheels and two caster wheels to move over flat surfaces—just like a typical power wheelchair. However, PENNWHEELS also has two large two-degree-of-freedom arms that can lift the front or rear wheels off the ground. By using the arms and powered wheels in concert, PENNWHEELS is capable of negotiating single steps, moving on to podiums, and so on.

Although PENNWHEELS can go where few other wheelchairs dare tread, it is definitely still in the conceptual prototype stage. The robot is tethered to its power system and to a computer that calculates the arm and the wheel movements. The motors are not sufficiently powerful to lift the chair's weight, let alone that of a passenger. Even with these limitations, PENNWHEELS was able to give an impressive demonstration of the possibilities of using hybrid wheel-legged mobility.

Robot Exhibition

This year's robot exhibition was an extraordinary crowd pleaser because all the robots that were demonstrated were highly interactive with the audience. The KISS Institute for Practical Robotics demonstrated some of its educational robot systems, giving elementary school students from the audience a chance to operate and control the robots. Newton Labs demonstrated its height-speed color-tracking system by having its robots chase after objects tossed into the ring by audience members. Finally, the Stanford University CHESHM robot interacted directly with large crowds of people as they tried to fool the robot and trick it into taking a dive down the stairwell. Everyone came out of the exhibition better educated and well entertained.

ED.BOT and FIRE-FLY CATCHER

ED.BOT, built by the KISS Institute, is a small mobile Lego robot the size of a shoe box. Its on-board brain is an MIT 6.270 board, and standard equipment includes front bump sensors, phototransistors, and wheel encoders. Powered by a small internal rechargeable battery pack, ED.BOT's Lego motors enable it to move forward or in reverse at the lightning speed of almost two miles an hour.

ED.BOT's purpose is purely educational. It is designed for classroom use at all elementary-school–age levels. ED.BOT's Lego structure is both familiar and understandable to young students. Its on-board programs demonstrate each of the sensors and motors that are used both individually and in combination to achieve simple tasks such as hiding in dark places, moving through figure eights, and hunting down light bulbs. Grade-school students use ED.BOT to gain an understanding of robot fundamentals, including exploring the basic systems and learning about design, system integration, and navigation. The little robot is also used as a base on which to build more complicated mechanisms.

DD.BOT participated at IJCAI-95 as an exhibition and hands-on demonstration of an educational robot; therefore, it was accessible to the many children walking by. Children as young as five years old were interested in leading this colorful little robot around by shining a light at its phototransistors. Even the youngest were able to grasp that the robot would turn toward the phototransistor that received the most light. Older children and adults could understand that the phototransistors were wired crosswise to the opposing motor-wheel unit, making the unit turn faster and the robot turn toward the light.

Perhaps it was best demonstrated by seven-year-old Kate Murphy, who enjoyed leading the little robot around with a flashlight and reading the appropriate light values off the displays as she assisted during one of ED.BOT's official demos in the arena. Murphy especially liked to make ED.BOT hide in the dark using its version of *Hide,* a program that teaches the concept of calibration, among other things.

ED.BOT's cousin, FIREFLY CATCHER, was also a big hit with the younger "roboteers." FIREFLY CATCHER, which was built as a design exercise for a robot class for 10 year olds, uses a similar robot base equipped with a large green net in a raised position in front. The net snaps down whenever front bumpers register contact, and the three phototransistors show light values in the correct pattern. A light bulb with toy wings on a small pedestal served as our "firefly." Occasionally, the children would start the robot angled away from the goal so that it would have to turn several times, orienting itself toward the light, and bump a few times against the pedestal before centering itself and swinging down the net on its innocent prey. It never missed.

NEWTON and Many Colored Things

The NEWTON vision cars originally came to Montreal as part of the robot exhibition. It was not until after they arrived that their code was modified so that they could compete in the Office Cleanup event of the robot competition.

The vision cars use the same hardware and color-tracking algorithms described earlier in the section NEWTON 1 and 2: Massachusetts Institute of Technology and Newton Labs. The key difference in programming was that for the exhibition, the robots went at full speed and tried to keep the objects they were looking for centered in their visual field.

The effectiveness of the tracking algorithms could best be seen in the Man versus Machine Contest, where an audience member was given the joystick to a radio-controlled car. The car was colored orange, and the driver's goal was simply to keep the car away from the NEWTON vision car. This task proved difficult. The audience member's turn ended when the vision car had rammed the radio-controlled car off its wheels or into a dead-end corner.

Tasks that were beyond the reach of robots a few years ago are now being done routinely in the competition.

The vision cars also chased rubber balls, went after frisbees, and were even able to keep hoops rolling indefinitely—at least until the far wall came up to meet them (at about 20 miles an hour!).

CHESHM

The umbrella project at Stanford University under which CHESHM was developed is called the Bookstore Project. The immediate goal of the Bookstore Project is easy to state: Create a totally autonomous robot that goes from the Stanford Computer Science Department to the bookstore and returns with a book. The more general goal is to create an autonomous navigator that can travel the entire campus, coexisting with bicyclists, cars, tourists, and even students.

Three important pieces of the Bookstore Project puzzle have been addressed over the past few years: (1) the ability to interleave planning and execution intelligently, (2) the ability to navigate (that is, move without becoming lost), and (3) the ability to stay alive.

CHESHM is Stanford's best attempt at solving the problem of staying alive, that is, designing a general-purpose obstacle-avoidance system. The real challenge is perception: A safe robot must detect all sorts of static and moving obstacles, not to mention pot holes, ledges, and staircases. CHESHM uses a totally passive vision system to perceive obstacles robustly. The current system has three important features: First, the depth-recovery system makes no domain assumptions. As long as there is enough light to see by, and the obstacle has some contrast, it will be avoided. Second, the vision system is totally passive and, therefore, does not have the interference or washout problems that active sensor systems such as infrared sometimes have. Third, the vision system is entirely on board, a necessity for truly autonomous mobile robotics.

CHESHM comprises a NOMAD 150 base and a vision system. The Nomad 150 has no sonar, infrared, or tactile sensors. The vision system is an on-board Pentium PC with a frame grabber and three Sony CCD cameras. The three cameras are pointed in the same direction so that the images received from the cameras are almost identical. Our depth-recovery system is based on the idea of depth from focus; so, the focusing rings of the three cameras are at different but known positions.

One of CHESHM's greatest weaknesses proved to be its willingness to run over its victims' feet.

By examining which of the three cameras maximizes sharpness for each image region, CHESHM can form a scene depth map. Obstacle recognition is easy because the angle of the cameras to the ground is known. Therefore, the floor is expected to be a specific distance away in the image. If the depth-map distance is closer than the floor for a particular region, then there is an obstacle there. If the depth-map distance is farther than the floor should be, then there is a pot hole or a staircase. This simple method for detecting steps has proven to be surprisingly reliable and might be a somewhat novel achievement for mobile robots.

CHESHM's motion is programmed using a MACINTOSH POWERBOOK 170 that is fixed on top of the robot. The POWERBOOK receives depth-map information (by serial port B) from the vision system and communicates velocity commands to the Nomad 150 base (by serial port A). The program that is being used to test CHESHM is an almost purely functional wandering program that turns away from dangerously close obstacles or stairs. This program performs no filtering or sensor interpretation; therefore, it is a transparent tool for examining the reliability of the vision module through observation of the robot's wandering behavior.

IJCAI-95 is the last in a series of three major tests of CHESHM's wandering behavior exclusively using this passive vision system. The first experiment consisted of wandering the third floor of Stanford's Computer Science Department. The greatest danger in this environment, other than the open staircase, proved to be graduate computer science students, who are possibly the most evil robot-testing group in existence. The robot succeeded in avoiding static and moving obstacles in this environment and even outsmarted several graduate students, to their dismay.

The second experiment involved wandering Stanford's Memorial Court, which is a large concrete and tile outdoor area bounded by bushes, ledges, and steps. CHESHM successfully interacted with more than 40 invited humans who herded the robot and tested its obstacle- avoidance capabilities. During a 2-hour experiment, the robot was herded toward and successfully recognized the stairs more than 15 times, with 100-percent reliability, and avoided all sizes of humans, save 1 head-on collision with a black dress. The interaction of children with CHESHM was fascinating; at one point, the children played *Ring around the Rosie* with CHESHM, dancing round the robot while it spun about, trying to find an escape route.

IJCAI-95 was CHESHM's final test: The robot wandered upstairs during three separate coffee breaks over the course of the conference. Each run was more than one hour long and again involved herding toward a nearby staircase in an attempt to force CHESHM down the stairs. Over the course of three hours, CHESHM experienced standing-room-only crowds (at the beginning of the coffee breaks) as well as intense stress testing from individual conference participants. CHESHM again avoided the staircase with perfect reliability and avoided the attendees well.

One of CHESHM's greatest weaknesses proved to be its willingness to run over its victims' feet. The field of view of the camera system simply does not see low enough to allow CHESHM to recognize feet and dodge them. When feet are located directly underneath legs, as is customary, the feet rarely pose a problem. However, when individuals try to trip CHESHM by sticking their feet out, they are asking for a painful experience. Over the course of more than three hours of testing upstairs among conference attendees, CHESHM successfully avoided all body parts (save feet) and all static and moving obstacles save four direct collisions with humans. Given that the robot successfully avoided hundreds of

The 1996 Robot Teams.

humans over the course of this experiment, the researchers were extremely pleased with the results.

Stanford researchers are convinced that their obstacle-avoidance solution is a good approach for the Bookstore Project. Now, they are revisiting navigation, this time using purely passive vision as the only sensor. The Bookstore Project is exciting because real-time, passive perception is beginning to look tenable using off-the-shelf processing power.

Robot Forum

The Robot Forum was held after the competition to allow for in-depth dialogue. At the forum, each team gave a short presentation on its robot entry. A small group of noncompetition researchers, including Reid Simmons, Tom Dean, and Leslie Pack Kaelbling, also gave their impressions of the competition and its impact on robotics. Then, a free-wheeling discussion occurred. The primary focus of the discussion was on the direction of the competition over the next few years. There was a general consensus that the competition needs to move toward more natural environments with moving obstacles and that longer tasks requiring more robustness should be encouraged. Many participants in the discussion felt it was time to start moving the competition out of a constructed arena and into the actual hallways and rooms of the conference center or the conference hotel. There was also a call for more interaction between the robots and the conference attendees. The discussions at the forum will help next year's organizers shape the AAAI-96 robot competition.

Conclusion

Overall, we were pleased with how the robots performed. Historically, robots tend to get stage fright. When the crowds start to gather, the robots seem to become unpredictable. It is not uncommon to hear, "I have no idea why it is doing that. It has never done *that* before!" Typically, the explanation turns out to be that all the camera flashes, infrared focusing sensors, and cellular phones interfered with the robots' sensors and affected communications. Although there were a few problems this year, as in past years, the reliability of the robots has definitely improved. A major contributing factor to this improvement was that the majority of teams did all their computing on board. History has clearly shown that this approach is a much more reliable configuration.

One objective this year was to define the role of the robot contests in the greater scheme of robotics research. The wheelchair exhibition did just that. The NAVCHAIR used a sonar-processing algorithm first demonstrated by a 1992 contest winner. TAO-1 used a vision system demonstrated in the 1993 robot exhibition, and WHEELESLEY is the next-generation refinement of another 1993 contest entry. The wheelchair application is an important and practical use for intelligent robotics, and much of the research that went into the prototype systems exhibited this year can be linked directly to robot contests of a few years ago.

On a more detailed level, this year we wanted to develop a core set of rules that outline the tasks to be completed but also allow teams some flexibility in making what would otherwise be arbitrary choices. For example, one of the objectives of the second event was to demonstrate object recognition and manipulation. The rules stated that the trash to be manipulated was Coke cans and Styrofoam cups. However, we allowed teams (at no penalty) to substitute other types of cans (Pepsi perhaps) if they worked better for them, as long as the substituted trash was of the same approximate size and available in stores. One team (University of Chicago) chose to use shape instead of color to distinguish between objects. Therefore, it decided to use multiple types and colors of soda cans to demonstrate that extra generality. Also, to reduce needless anxiety among the teams, a rough outline of the arena was provided in advance, allowing teams to anticipate potential last-minute difficulties.

As the designers of previous years' competitions have attested to (Simmons 1995), designing a set of rules and a scoring mechanism that is fair to such a diverse group of robots and strategies is a difficult task. A lot of careful thought went into designing the scoring and penalties. The objective was to take the lessons learned from past years and construct an unambiguous 100-percent-objective scoring criteria and to not deviate from the announced scoring once the events began.

One of the key difficulties was in how to design a manipulation task that was fair to both physical- and virtual-manipulator robots. Although physical manipulation is obviously preferred over virtual (because of its inherently autonomous nature), past competitions have had few successful physical-manipulation robots. Because virtual manipulation can be so much faster than physical manipulation, we had to compensate somehow.

Based on past contests, we decided that physically placing trash inside the trash can would take approximately three times as long as virtual manipulation and that placing the trash near the trash can would take about twice as long. Thus, the final rules said that actually placing the trash in the trash can was worth 35 points each, pushing trash into the trash zone (near the trash can) was worth 20 points, and virtually placing the trash in the trash can was worth 10 points. In addition, the event would have two first-place winners, one overall winner based on the total score, and one winner in the physical-manipulator category. It turned out that all but one of the robots used physical manipulation and that the overall winner (NCSU) used physical manipulation and took both awards. We were also heartened by the fact that the final results, based on an objective scoring system, matched most observers' subjective impressions of each robot's abilities.

Overall, the last four years of robot competitions have been successful at pushing the state of the art in mobile robotics. Tasks that were beyond the reach of robots a few years ago are now being done routinely in the competition. This steady upward trend is primarily the result of advances in vision-processing techniques (especially color vision processing) and mobile manipulation. The competitions have allowed for sharing of technical information across the community of researchers, and a benchmark set of tasks has evolved that allows for comparison of competing technology approaches.

Acknowledgments

The competition organizers want to thank the many people and companies that made the competition successful. We especially want to thank the members of all the robot teams who went through extraordinary lengths, numerous border crossings, and severe sleep deprivation to bring their robots to IJCAI-95. Additional thanks go to Microsoft, the Advanced Research Projects Agency, IJCAI, and AAAI for their generous donations. Members of the AAAI staff who helped tremendously were Carol Hamilton, Mike Hamilton, Sara Hedberg, and Rick Skalsky. Awards for the winning teams were provided by Network Cybernetics Corporation (Irvington, Texas), distributors of the AI CD-ROM. Thanks to Pete Bonasso for videotaping the entire competition and Linda Williams for helping with timing and general organization. Finally, NASA Johnson Space Center's Robotics, Automation, and Simulation Division and the MIT AI Lab graciously allowed us to establish competition World Wide Web sites on their computers.

References

Borenstein, J., and Koren, Y. 1991. The Vector Field Histogram for Fast Obstacle Avoidance for Mobile Robots. *IEEE Journal of Robotics and Automation* 7(3): 535–539.

Dean, T., and Bonasso, R. P. 1993. 1992 AAAI Robot Exhibition and Competition. *AI Magazine* 14(1): 35–48.

Konolige, K. 1994. Designing the 1993 Robot Competition. *AI Magazine* 15(1): 57–62.

Kweon, I.; Kuno, Y.; Watanabe, M.; and Onoguchi, K. 1992. Behavior-Based Mobile Robot Using Active Sensor Fusion. In Proceedings of the 1992 IEEE International Conference on Robotics and Automation. Washington, D.C.: IEEE Computer Society.

Miller, D. P., and Grant, E. 1994. A Robot Wheelchair. In Proceedings of the AAIA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space, 407–411. NASA Conference Publication 3251. Houston, Tex.: National Aeronautics and Space Administration.

Nourbakhsh, I.; Powers, R.; and Birchfield, S. 1995. DERVISH: An Office-Navigating Robot. *AI Magazine* 16(2): 53–60.

Nourbakhsh, I.; Morse, S.; Becker, C.; Balabanovic, M.; Gat, E.; Simmons, R; Goodridge, S.; Potlapalli, H.; Hinkle, D.; Jung, K.; and Van Vactor, D. 1993. The Winning Robots from the 1993 Robot Competition. *AI Magazine* 14(4): 51–62.

Simmons, R. 1995. The 1994 AAAI Robot Competition and Exhibition. *AI Magazine* 16(2): 19–30.

Simpson, R. C.; Levine, S. P.; Bell, D. A.; Koren, Y.; Borenstein, J.; and Jaros L. A. 1995. The NAVCHAIR Assistive-Navigation System. In Proceedings of the IJCAI-95 Workshop on Developing AI Applications for the Disabled. Menlo Park, Calif.: American Association for Artificial Intelligence.

David Hinkle is a senior scientist at Lockheed Martin's Artificial Intelligence Research Center. He received his B.S. in computer science from Syracuse University in 1987 and his M.S. in computer science from Northeastern University in 1989. His research interests include machine learning, neural networks, and robotics. He is currently conducting research in the area of data mining. He is a member of the American Association for Artificial Intelligence. He is also an organizer for the 1996 AAAI Robot Competition and Exhibition in Portland, Oregon.

David Kortenkamp is a research scientist with Metrica Incorporated and is responsible for directing research projects at NASA Johnson Space Center's Robotics Architecture Lab. He received his B.S. in computer science from the University of Minnesota in 1988. He received his M.S. and Ph.D. in computer science and engineering from the University of Michigan in 1990 and 1993, respectively. His research interests include software architectures for autonomous agents, mobile robot mapping, navigation and perception, and human-robot interaction and cooperation. Kortenkamp has been involved in all the American Association for Artificial Intelligence (AAAI) robot competitions, including leading the group that finished first in the inaugural competition in 1992. He is also an organizer for the 1996 AAAI Robot Competition and Exhibition in Portland, Oregon.

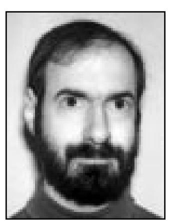

David Miller received his B.A. in astronomy from Wesleyan University and his Ph.D. in computer science from Yale University in 1985. He has been building and programming robot systems for more than a dozen years. He spent several years at the Jet Propulsion Laboratory, where he started the Planetary Micro-Rover Program. He won the Most Innovative Entry Award in the First AAAI Robot Competition and Exhibition for SCARECROW, the robot he built with his then–five-year-old son. He is currently the cochair of the Robotics, Resources, and Manufacturing Department of the International Space University Summer Program and is also the technical director of the KISS Institute for Practical Robotics.