The 1995 IJCAI Robot Competition and Exhibition

David Hinkle, David Kortenkamp and David Miller

December 5, 1995

Abstract

The fourth annual 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, Office Delivery and Office Clean-up. The Exhibition also consisted of two events: demonstrations of robotics research that was not related to the contest; and 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

1 Introduction

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

This is the fourth annual competition in a string of competitions that began at AAAI-92 in San Jose, CA. During this inaugural competition, ten robots searched for and approached a set of tall poles in a large arena [1]. The next competition, 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 [5, 2]. The third competition, held at AAAI-94 in Seattle, Washington, contained events that included office building navigation and trash pickup [7]. The fourth competition (the first at an IJCAI) built on the successes of the previous competitions.

The goals for the competition and exhibition have remained the same over the years. First, 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, to assess (and push) the state-of-the-art in robotics. Third, to contrast and compare competing approaches as applied to the same task. Finally, to provide a public forum in which robotics research can be seen by the artificial intelligence community, the media and the general public. Because of these broad range of goals, determining an appropriate format for the competition and exhibition was quite challenging.

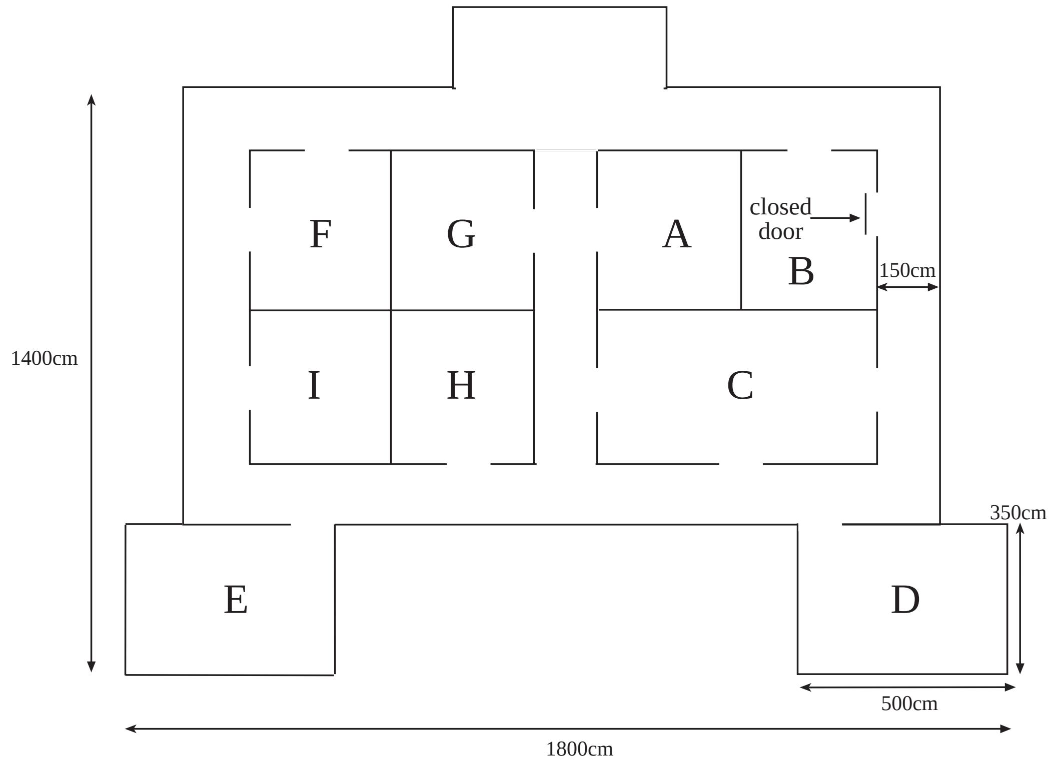

Figure 1: The competition arena. A closed door (room B) is inset from the walls.

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 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 will be described in detail in the following sections and results will be given. Then the wheelchair exhibition, the robot exhibition and the forum will be discussed in turn.

2 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 (see 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 like "Exit the Room and turn Left", "Go to the End of the Hallway and turn Right", "Enter the room on your Third Right", etc. (See

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.

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 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, who 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 how long 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 they would receive 20 points (25 - 5 min.). 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 1 point per 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 more of an emphasis placed on recovery from mistakes and human-robot interaction. At the third AAAI robot competition in 1994 [7], 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 [6] was able to complete the event in 1994.

2.1 Results

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

1. Korea Advanced Institute of Science and Technology (130.5 points)

2. North Carolina State University (114.0 points)

3. Kansas State University (70.5 points)

4. University of New Mexico
   (unable to compete due to robot damage during transportation to Montreal).

2.2 Teams

Now we will briefly describe each team's robot, and their approach to the event.

2.2.1 Korea Advanced Institute of Science and Technology

CAIR-2 is a custom built prototype robot, developed at the Korea Advanced Institute of Science and Technology (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. Though 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 had 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 very 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 5:25 to complete the event. Since 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 their running time. For more information on CAIR-2 please see the article in this issue.

2.2.2 North Carolina State University

Lola is a Nomad 200 robot equipped with a pan/tilt-mounted RGB camera and sonar sensors. Lola's on-board computation included a dual C40 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 via radio Ethernet. The control architecture was comprised of 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, etc.

3. Low level motion control: for determining the direction and velocity of the robot as well as to perform obstacle avoidance.

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 afforded the robot a better position from which to scan the room for the exit door. Each doorway was marked with a round 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 very well throughout both the finals and the trials. The scores for both trials and the finals were virtually identical. In the finals, Lola exited the start room in about one and a half minutes, and completed the event in 6:15. For a more information on Lola please see the article in this issue.

2.2.3 Kansas State University

The Kansas State University 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 to identify doors in the hallway. A subsumption architecture with threads running the sonar detection and avoidance routines was used. For example, the exit room strategy involved first 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 one and a half minutes, detected the error in the human supplied instructions, accepted the new corrected instructions, and completed the task within five minutes. However, during the second trial run, due to radio interference, the performance of the radio Ethernet degraded severely.

Continuation in the contest required porting the 10K lines of the control program from the UNIX workstation to the 486 processor on board the robot. This also required installing Linux and a threads package on the 486 processor. The porting took 12 hours and involved re-tuning 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.

2.2.4 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 cm tall and measures about 60 cm in width. Sensing is achieved using a layout of 16 ultrasonic transducers. The onboard distributed processing and control system consists of a 486 PC-based master and networked MC68HC11 slaves.

The LOBOt employs an object-oriented behavioral approach based on a task decomposition of Event 1 (i.e. exit-room, navigate-hallway, and 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: Navigation, the Helm, and ODRR (Object Detection, Recognition, and Ranging). The Helm and ODRR encapsulate the usable resources of the Lobot, such as motor control and

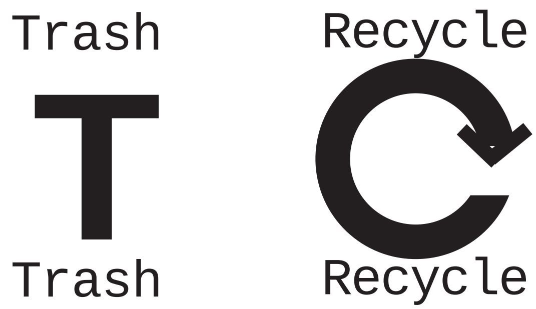

Figure 3: Symbols for the trash bin (left) and recycle bin (right).

sensor data. Navigation is a client of the ODDR and Helm, utilizing them for perception, sensor-based decision-making, and initiating appropriate actions. The motion behaviors are implemented as objects within Navigation which operate using an artificial force field approach. In particular, the exit-room behavior is based on an augmented force field approach adapted from [3]. That is, actual sonar range data is augmented with additional *virtual* range data which 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. But the team vowed to try again next year!

3 Robot Competition: Event 2 (Office Clean-up)

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 (see 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 was for the robots to pick up the soda cans and deposit them in the recycling bin and to 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. While "virtual" manipulation was allowed (i.e., 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 of the other robots had some form of manipulation. This is a large improvement over last year's competition in which a similar event attracted only two teams that performed actual manipulation (see [7]).

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 those perceptions. Thus, the robots needed to recognize two classes of objects (trash and recyclables) and also recognize two classes of containers (trash bins and recycling bins). The classes of containers were marked with the symbols 'T' and the recycling closed circle for trash and recyclables respectively (see Figure 3). This was an advancement over last year's event in which there was only one class of objects, trash, and only one type of receptacle.

3.1 Results

We saw a vast improvement in the performance of robots in this event over what was demonstrated last year. All of the robots successfully used vision to distinguish amongst the objects. All of 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 in object recognition. The 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 bin; 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**: North Carolina State University (196 points)
2. **CHIP**: University of Chicago (165 points)
3. **Walleye**: University of Minnesota (135 points)
4. **Newton 1 & 2**: MIT and Newton Labs (105 points)
5. **Clementine**: Colorado School of Mines (79 points)

3.2 Teams

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

3.2.1 Lola: North Carolina State University

Lola is a Nomad 200 robot equipped with a prototype Nomadics manipulator, pan/tilt-mounted RGB 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:

1. Lola locates trash using predefined color histogram models of the different types of trash and histogram backprojection.
2. Lola heads off in pursuit of the trash.
3. During pursuit, Lola tracks the centroid of the trash as it moves down the image plane and employs a nonlinear least-squares algorithm to calculate its location relative to the robot.
4. Once within range, Lola grasps trash using position estimation.

Once Lola has grasped 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 extremely well in the preliminary round depositing 13 objects. In the final round, Lola was performing extremely 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 audiences 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 the accompanying article in this issue.

3.2.2 CHIP: University of Chicago

CHIP is a small mobile robot built on an RWI 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 Packages (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 the accompanying article in this issue.

3.2.3 Walleye: University of Minnesota

The chassis of Walleye is built out of 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 and 32k of RAM. Walleye uses 3 microcontrollers, one for the main board, one to control the motor, and one for the vision system. The vision system uses a CCD 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 up to 20 frames/second. Walleye has a gripper with a beam across the fingers to detect when something has been grasped. The gripper can not 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 towards it, tracking the object while approaching. If the object seen was not really an object, the illusory object will disappear while being tracked. Walleye would then start again 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, so 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 not movable 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 he lets go of the trash, depositing it. He then continues the search for more trash.

Walleye performed extremely well in the first preliminary round, pushing eleven objects to the correct bin. Reflections in the floor due to an overnight cleaning were responsible for a sub-par performance in the second round, where Walleye only deposited three 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 per 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.

3.2.4 Newton 1 & 2: MIT/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 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 this data at 60 Hz 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, we wrote about 65 lines of C for the low-level interface to the car, and about 300 lines of C to implement the control software for the contest. We had two cars in the contest. Each car would focus 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 towards 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 very 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 they did not place the objects in the bins, they received fewer points per object than other teams ahead of them.

3.2.5 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 6-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, up 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 bar-code-like artificial landmarks. It also knew a priori where the trash bins were.

The trash recognition process was successful and in a preliminary round, detected all 10 cans depositing seven of them. In the second round, Clementine deposited 9 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.

4 Wheelchair Exhibition

A robotic wheelchair exhibition was added to this year's event in order to demonstrate how the robotics technology that has been developed over the last several years could be usefully applied. There are many people who are mobility impaired but are not able to operate a normal power wheelchair safely [4]. This year's exhibitors concentrated on supplementing the control system of a power wheelchair in order to endow it with some semi-autonomous navigation and obstacle avoidance capabilities. The chairs of course also had to be able to integrate continuous human commands as well as following their programmed instructions.

Three chairs with automatic guidance systems were brought to IJCAI. 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 is tethered to its power supply and computer).

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

4.1 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 due to cognitive, perceptual or motor impairments.

By sharing vehicle control decisions regarding obstacle avoidance, safe object approach, maintenance of a straight path, etc., 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 & Jennings. The Lancer's controller is divided into two components: the joystick module, which receives input from the user via the joystick and converts it to a signal representing desired direction, and 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 which 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 Bell). MVFH is based on the Vector Field Histogram (VFH) algorithm by Borenstein & Koren, which was originally designed for autonomous robots. MVFH allows the 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 on the NavChair focuses on using probabilistic reasoning techniques from artificial intelligence research to extend this modeling capability [8].

4.2 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 with the joystick or with the user interface. The joystick mode is similar to the other teams' joystick mode, so the user interface will be discussed here.

The user interface runs on a Macintosh Powerbook. Although the input to the interface currently is 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 while others need to use a sip and puff system. Some users could use voice; others can not. The interface that was shown at IJCAI is very general but would have to be tailored to the needs of the specific user.

Wheelesley's interface provides information while allowing 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, we have provided a map of the wheelchair that 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. Wheelesley was the only system that could drive through doorways without needing to be 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.

4.3 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, both in 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 micro controller (a single chip computer with on-chip memory and control electronics). It has AAI's own multi-tasking, real-time operating system which allows the controller to receive real time signals from a large number of sensors and it sends control outputs to two motors to drive the left and right wheels. It looks after both forward/backward and left/right movements of the chair.

Two color CCD cameras mounted on the chair detect free space and motion up to 10 meters in front of the chair. Six active infra-red (IR) sensors detect obstacles in close vicinity, up to 1 meter from the chair. The signal from the cameras is processed by an intelligent vision processing unit which is also built on behavior-based principles. The control program for all the vision processing occupies 9.5 KBytes and the other behavior control occupies 2.75 KBytes. This is significantly smaller than similar vision-based control programs operating in 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 is concentrated 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 avoiding common pitfalls which 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 appearance of the pavement, and traffic signals. The number of IR sensors will be increased to allow it to move in more confined spaces. Later phases will developed sophisticated interactions between the human and the chair, improve mobility aspects of the chair, and introduce evolutionary computation (EC) methodologies to facilitate the chair adjusting to the needs and situations of each individual user.

4.4 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 normal power wheelchair. But PennWheels also has two large 2-degree of freedom arms which 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, etc.

While 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 wheel movements. The motors are not sufficiently powerful to lift the chair's weight, let alone that of a passenger. But even with these limitations, PennWheels was able to give an impressive demonstration of the possibilities of using hybrid wheel/legged mobility.

5 Robot exhibition

This year's robot exhibition was an extraordinary crowd pleaser as all of the robots that were demonstrated were highly interactive with the audience. KISS Institute 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 their height speed color tracking system by having their robots chase after objects tossed into the ring by audience members. And 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 entertained.

5.1 Ed.Bot & Fire-Fly Catcher

Ed.Bot, built by the KISS Institute for Practical Robotics, is a small mobile LEGO robot the size of a shoe-box. His onboard brain is an MIT 6.270 board, and standard equipment includes front bump sensors, photo-transistors, and wheel encoders. Powered by a small internal rechargable battery pack, Ed.Bot's LEGO motors enable him to forward or reverse at the lightning speed of almost 2 mph.

Ed.Bot's purpose is purely educational. He is designed for classroom use at all elementary school age levels. Ed.Bot's LEGO structure is both familiar and understandable to young students. His 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 upon which to build more complicated mechanisms.

Ed.Bot participated at IJCAI as an exhibition and hands-on demonstration of an educational robot, therefore he was quite 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 photo-transistors. Even the youngest were able to grasp that the robot would turn towards the photo-transistor that received the most light. Older children and adults could understand that the photo-transistors were wired crosswise to the opposing motor/wheel unit, making that unit turn faster and the robot turn towards the light.

Perhaps he 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. Kate especially liked to make Ed.Bot hide in the dark using his 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, who 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 photo-transistors show light-values in the correct pattern. A light bulb with toy wings on a small pedestal served as our "firefly." Occasionally the children at IJCAI 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.

5.2 Newton & 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 contest.

The Vision Cars use the same hardware and color-tracking algorithms as described in section 3.2.4. 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 be best seen in the "Man vs. 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 simple to keep the car away from the Newton Vision Car. This task proved quite difficult. The audience members turn ended when the Vision Car had rammed the R/C car off its wheels or into a dead-end corner.

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

5.3 Cheshm

The umbrella project at Stanford University under which Cheshm 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.

There are three important pieces of the Bookstore Project puzzle that we've been addressing over the past few years:

1. The ability to interleave planning and execution intelligently
2. The ability to navigate (i.e., move without becoming lost)
3. The ability to stay alive

Cheshm is our 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 system we have implemented has three important features: 1) The depth recovery system makes no domain assumptions. Technically, it will fail (but will recognize that fact!) if the environment is too dark or if the obstacle has zero contrast. 2) The vision system is totally passive and therefore does not suffer from any interference or washout problems like infrared and laser-rangefinder. 3) The vision system is entirely on-board, a necessity for truly autonomous mobile robotics.

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

By examining which of the three cameras maximizes sharpness for each image region, Cheshm can form a scene depthmap. Obstacle recognition is easy because we know the angle of the cameras to the ground. Therefore, we expect the floor to be a specific distance away in the image. If the depthmap distance is closer than the floor for a particular region, then there is an obstacle there. If the depthmap distance is farther than the floor ought to be, then there is a pothole or staircase. This simple method for detecting steps has proven to be surprisingly reliable and may be a somewhat novel achievement for mobile robots.

We program Cheshm's motion using a Macintosh Powerbook 170 that is fixed on top of the vision system. The powerbook receives depthmap information (via serial b) from the vision system and communicates velocity commands to the Nomad 150 base (via serial a). The program that we have been using to test Cheshm is an almost purely functional wandering program that turns away from dangerously close obstacles and steps. This program performs no filtering nor 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 is the last in a series of three major tests of Cheshm's wandering behavior exclusively using this passive vision system. Our 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 C.S. 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.

Our 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 forty invited humans who herded the robot and tested its obstacle avoidance capabilities. During a two- hour experiment, the robot was herded toward and successfully recognized the stairs more than 15 times, a 100% reliability, and avoided all sizes of humans, save one 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 it while it spun about, trying to find an escape route.

IJCAI 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 very 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 humans over the course of this experiment, we were extremely pleased with the results.

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

6 Robot Forum

The robot forum was held after the competition to allow for an in-depth dialogue. At the forum each team gave a short presentation on their robot entry. A small group of non-competition researchers, including Reid Simmons, Tom Dean and Leslie Pack Kaelbling, also gave their impression of the competition and it's impact on robotics. After this, 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 towards 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 of 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.

7 Conclusion

Overall we were quite 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 affecting communications. Although there were a few problems this year, as in the past, the robots have definitely improved in reliability. A major contributing factor to this was the fact that a majority of teams did all of their computing on-board. History has clearly shown that this 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 the 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 these prototype systems shown this year can be directly linked to robot contests a few years back.

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 to 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 that worked better for them, as long as the substituted trash was of the same approximate size and available in stores. One team (Chicago) chose to use shape instead of color in order to distinguish between objects. Therefore they decided to use multiple types and colors of soda cans to show off that extra generality. Also, in order 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 [7], designing a set of rules and scoring mechanism that is fair to such a diverse group of robots and strategies is a difficult task. A lot of careful thought had to go into designing the scoring and penalties. The objective was to take the lessons learned from past years and construct an unambiguous 100% 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. While physical manipulation is obviously preferred over virtual (due to its inherently autonomous nature) past competitions have had few successful physical manipulation robots. Since 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. So 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 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 very 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 due to advances in vision processing techniques (especially color vision processing) and in 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.

8 Acknowledgments

The competition organizers wish to thank the many people and companies that made the competition successful. We especially wish to thank all of the members of all of the robot teams who went through extraordinary lengths, numerous border crossings and severe sleep deprivation to bring their robots to IJCAI. Additional thanks go to Microsoft, ARPA, 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, Linda Williams for helped 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 WWW sites on their computers.

References

[1] Thomas Dean and R. Peter Bonasso. 1992 AAAI robot exhibition and competition. *AI Magazine*, 14(1), 35-48, 1993.

[2] Kurt Konolige. Designing the 1993 robot competition. *AI Magazine*, 15(1), 57-62, 1994.

[3] I Kweon, Y. Kuno, M. Watanabe, and K. Onoguchi. Behavior-based mobile robot using active sensor fusion. In *Proceedings of 1992 IEEE Int'l. Conf. on Robotics and Automation*, 1992.

[4] D. P. Miller and E. Grant. A robot wheelchair. In *Proceedings of the AAIA/NASA Conference on Intelligent Robots in Field, Factory, Service and Space*, 407-411, 1994.

[5] Illah Nourbakhsh, Sarah Morse, Craig Becker, Marko Balabanovic, Erann Gat, Reid Simmons, Steven Goodridge, Harsh Potlapalli, David Hinkle, Ken Jung, and David Van Vactor. The winning robots from the 1993 robot competition. *AI Magazine*, 14(4), 51-62, 1993.

[6] Illah Nourbakhsh, Rob Powers, and Stan Birchfield. DERVISH: An office-navigating robot. *AI Magazine*, 16(2), 53-60, 1995.

[7] Reid Simmons. The 1994 AAAI robot competition and exhibition. *AI Magazine*, 16(2), 19-30, 1995.

[8] R.C. Simpson, S. P. Levine, D.A. Bell, Y. Koren, J. Borenstein, and L.A. Jaros. The navchair assistive navigation system. In *Proceedings of the IJCAI'95 Workshop on Developing AI Applications for the Disabled*, 1995.