The 1996 AAAI Mobile Robot Competition and Exhibition

David Kortenkamp, Illah Nourbakhsh and David Hinkle

Abstract

The Fifth Annual Mobile Robot Competition and Exhibition was held in Portland, Oregon in conjunction with AAAI-96. The competition consisted of two events: 1) Call a meeting and 2) Clean up the tennis court. The first event stressed navigation and planning. The second event stressed vision sensing and manipulation. In addition to the competition, there was a mobile robot exhibition in which teams demonstrated robot behaviors that did not fit into the competition tasks. The competition and exhibition were unqualified successes, with nearly 20 teams competing. The Robot Competition raised the standard for autonomous mobile robotics, demonstrating the intelligent integration of perception, deliberation and action.

1 Introduction

This article describes the Fifth Annual AAAI Mobile Robot Competition and Exhibition, which was held in conjunction with AAAI-96 in Portland, Oregon from August 3, 1996 to August 8, 1996. This competition built on the successes of four earlier competitions at AAAI-92 [1], AAAI-93 [3, 4], AAAI-94 [5] and IJCAI-95 [2]. The 1996 competition was the most widely attended, with nineteen teams from three countries and fourteen states. In addition the PBS series "Scientific American FRONTIERS" covered the competition and series host Alan Alda watched the finals and interacted with the robots.

The events of the 1996 competition built upon those of previous years, offering incrementally harder challenges to push state-of-the-art robotics. As in the past, there were two tasks, one stressing navigation and planning and the other stressing sensing and manipulation. There was also a free-form exhibition in which teams could demonstrate robots and techniques that were innovative but did not lend themselves to the competition tasks. The two events and the exhibition will be discussed in detail in the next three sections.

2 Event 1: Call a meeting

The first event required the robots to "call a meeting" between two professors and the Director. The event was held in an arena with an office building floorplan, shown in Figure 1. The robot's first task was to navigate from the Director's office to one of two conference rooms in order to detect whether the room was occupied or not. If it was occupied, the robot checked to see if the second conference room was available. If the second conference room was also occupied, the robot was to schedule the meeting in the Director's office. Once the location of the meeting was established, the robot had to give the two professors and the director two pieces of information: 1) The location of the meeting; and 2) The starting time of the meeting. The meeting should start one minute after all three participants have been informed. This required the robot to predict as accurately as possible how long it would take to find each person and announce the meeting. Shortly before the competition the robots were given a graphical representation of the office building, showing rooms and hallways and rough distances. They could use this representation for planning and for reasoning about time. In order to simulate a realistic environment, the rules noted that there could be people moving about in the office building, possibly blocking hallways and doorways.

This event was more difficult than similar events in previous years requiring several new skills. 1) The robots had to detect occupancy of rooms; 2) they had to predict their completion time; and 3) moving

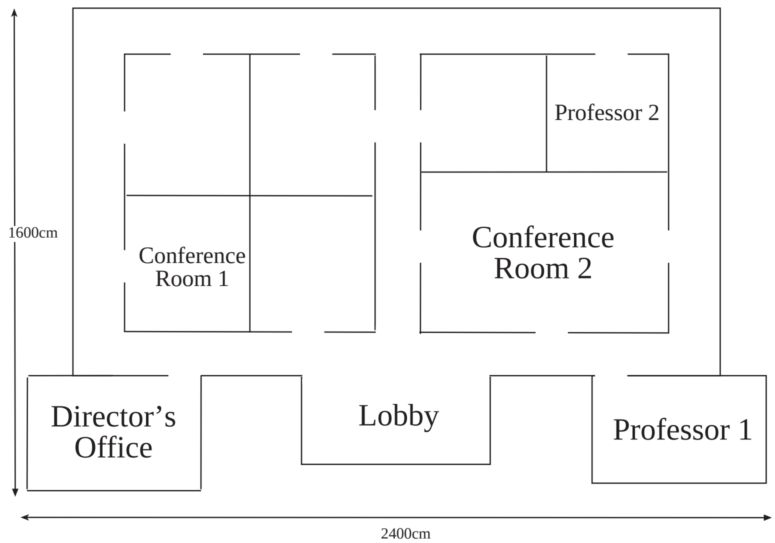

Figure 1: The arena used in Event 1.

obstacles (people) were in the halls and doorways. In addition the task overall was longer than in previous years, both in terms of the total distance the robots had to travel and the overall time the robots needed to operate without making a mistake. Scoring was based on the robots accomplishing specific portions of the task (e.g., entering the first conference room, correctly determining occupancy, etc.), on how quickly the robots were able to perform the task, and on the accuracy of their time prediction. Robots might have points deducted from their score for several reasons:

1. Modifications to the arena were penalized: One of the objectives is to encourage the development of algorithms that can handle uncustomized environments.

2. Unexplained or inefficient actions by the robot: The robots were required to solve the task as efficiently as possible. As a default measure of efficiency we used shortest path necessary to acomplish the task. However, this is not the only possible definition of efficiency. For example, the shortest route may not always be the fastest route. Therefore, we allowed the teams to develop their own measures of efficiency as long as the robot explained the rational for its actions at each decision point in the event.

3. Collisions: If the robot collided with any of the stationary objects in the arena or a person standing in the hallway, they were penalized. Slight contact between the robot and an object was not penalized. Some teams used a bumper sensor to detect any initial contact and then reacted appropriately before a stronger contact was made.

4. Assistance to the robot: The robots were required to operate fully autonomously throughout the event. If a robot became confused and needed to be restarted teams were allowed to intervene, at a penalty. If the robot was able to recognize that it needed assistance and requested it, the penalty was cut in half.

Robots were also scored on their occupancy detection method. Robots that could detect the motion of someone pacing about in the room (giving a lecture, for example) were given high scores, while those robots that needed to ask the occupants of a room to perform some action (such as press a key on the robot's keyboard) were given lower scores.

The event was an unqualified success, with many of the teams able to complete the entire task. There was controversy however, as the top finisher, SRI International, took a multi-agent approach to the task (see related article in this issue). While a multi-agent approach was not prohibited by the rules (in fact, the University of Minnesota also planned a multi-agent approach, but mechanical difficulties caused them to withdraw) some of the other top teams felt that dividing the task amongst robots that could communicate with each other over radio modems gave an unfair time advantage. Indeed, the SRI robots finished the entire task in just under five minutes, while the next closest robot (Kansas State Team 1) took just over nine minutes. The points SRI's multi-robot entry received for being that much faster than all the other single robot-entries enabled them to overcome their lower score for room occupancy detection and a collision penalty.

The competition organizers wanted to encourage innovative, multi-agent entries because cooperative mobile robots is a significant research area. In fact, most multi-agent entries in past competitions had been failures because of the added complexity of keeping several robots working and maintaining communication and control over the robots. However, as SRI proved, multi-agent solutions can offer significant time advantages over single agent solutions in some domains. In future competitions, separate first place awards may have to be given for multi-agent and single agent entries.

The Kansas State Team 1 and the USC/ISI (Yoda) teams tied for second place and both scored perfectly in every aspect of the task, received no penalties, and their time predictions were exactly correct. The Kansas State Team I was the fastest single-agent robot, finishing twice as fast as USC. Several other teams had very good showings in the final round. The University of Texas at El Paso, with a robot called Diablo, used the three-tiered architecture approach to complete the task. [Can we add anything else about the architecture? another phrase or so would be fine] Kansas State Team 2 also finished very well and demonstrated an occupancy detection system that was typical of most robot entries. They performed occupancy detection by taking successive frames from the robot's cameras and splitting these images into vertical slices. Differences between these vertical slices on successive frames were used to determine if something had moved. The Colorado School of Mines entered their robot Clementine for the fourth straight year. Each year a new team of undergraduate students, as part of a class project, enters the competition. This year the team successfully completed Event 1 and finished in fourth place.

Several teams used sophisticated Partially Observable Markov Decision Process (POMDP) for navigation. The POMDP framework takes into account the uncertainties associated with action and perception and builds a belief distribution on the topological map. The topological node with maximum belief is assumed to be the current location of the robot. Robots using POMDP models included North Carolina State and Carnegie Mellon University, with its robot Amelia.

Several teams competed in the preliminary rounds of Event 1, but due to mechanical break-downs or software bugs they did not advance to the final rounds. The McGill University team, composed of undergraduates, had several slow, but successful trial runs. They worked all night to speed up their robot, but introduced bugs that prevented them from advancing to the finals. The team from the University of Stuttgart had an impressive entry that smoothly and competently completed several trial runs. However, a failure of their sonar sub-system at the last minute prevented them from competing in the finals. The third Kansas State University team had software difficulties that prevented them from competing. All three Kansas State teams were groups of undergraduate students working independently of each other, but sharing a common hardware platform.

Finally, the University of Minnesota had an innovative multi-agent approach in which a "mother" robot would visit the two conference rooms and then, after determining which was free, would dispatch two "baby" robots that it was carrying to inform the professors. Mechanical problems with this complicated approach caused them to withdraw from the competition at the last minute.

3 Event 2: Clean up the tennis court

In the second event the robot was placed in an enclosed, rectangular room. In the room were ten ordinary tennis balls and two moving "squiggle" balls. Squiggle balls are motorized balls found in toy stores that move around fairly quickly and, upon encountering obstacles, bounce off of them and head in another direction. In one corner of the room the teams placed a pen of their own design. The task was to collect tennis balls and squiggle balls and place them in the pen. The addition of moving objects added a level of complexity to this task that was not present in previous competitions. The fast moving squiggle balls would challenge robots to act and react quickly, placing a premium on fast vision and manipulation capabilities, coupled with an effective search strategy. Fortunately, several teams were up to the challenge.

Scoring for this event was fairly simple. Teams scored twenty points for each ball (either squiggle or tennis) in the pen at the end of the round (15 minutes). Teams scored an additional fifty points for capturing a squiggle ball, even if they did not deposit it. In addition, teams received thirty points for demonstrating that they could track the squiggle ball and thirty additional points for intentionally touching the squiggle ball even if they did not capture it (these additional points could only be accrued once). The purpose of the tracking and contact points was two-fold: 1) To allow teams that could not do manipulation quickly enough to capture squiggle balls, but could track and hit them, to compete and gain points; and 2) To provide some way to reward teams that explicitly track and capture squiggle balls instead of simply scooping up everything in the pen. Finally, penalties were given for teams that destroyed squiggle balls in any way and teams that marked tennis balls. Teams could mark squiggle balls and the pen in any way they chose.

Two teams completed the entire event, capturing all ten tennis balls and both squiggle balls. These teams used very different approaches (see the related articles in this issue). The team from Newton Research Labs had a single, small robot with a gripper that could hold a single tennis or squiggle ball. Their robot had a fast color vision system that could sense the balls and provide immediate feedback to the robot controller. When the robot spotted a squiggle ball it would set off in pursuit, rapidly accelerating its speed to overtake the fleeing squiggle ball, dramatically scooping it up from behind. The robot was an instant crowd pleaser because of its fast, animate action. The team from the University of Bonn/RWI/CMU (combined team) took a different approach. They had a large robot with a sweeper in front. The robot could hold many tennis and squiggle balls. The robot then performed a systematic sweep of the arena. Not just a sweeper, this robot employed sophisticated vision and sonar sensing to detect errant tennis balls and go after them. This robot could also track the squiggle balls, although in the final round it retrieved the two squiggle balls during its sweep without doing any tracking.

The University of Utah, with its robot IGOR, came very close to completing the entire task. Utah also used the sweeper approach, with a large mechanical device that was attached to the front of the robot. IGOR demonstrated tracking of squiggle balls and captured both Squiggles successfully. It also captured all of the tennis balls with the exception of those laying against the outer wall. Two other teams successfully competed in Event 2. The University of Minnesota had a small robot called Walleye with a flipper-type gripper that could hold up to three squiggle or tennis balls. Both squiggle balls and tennis balls were painted black and Walleye could detect them using a black and white camera. Walleye demonstrated tracking and capture of squiggle balls. Mechanical glitches in the final round prevented Walleye from scoring as highly as it had in the preliminary rounds, when it was able to pick up most of the balls. North Carolina State University had the only entry that used a conventional four-degree-of-freedom robot manipulator and gripper. While it was an impressive design by an undergraduate team, the arm and gripper were much too slow to have any hope of catching a squiggle ball. However, the robot could track squiggle balls using a color histogram matching technique and it did an exceptional job at picking up stationary tennis balls. The North Carolina State manipulator demonstrated the impressive ability to choose between multiple gripping strategies, depending on whether the target ball was flush with a side wall or free on all sides. The University of Stuttgart achieved impressive scores in several preliminary rounds using a sweeper robot that was strikingly similar to the Bonn/RWI/CMU robot. However, this robot was withdrawn from the finals because of a mechanical breakdown.

4 Exhibition

This year's robot exhibition offered an extremely diverse set of technology demonstrations. An important common theme throughout the exhibition was how robotics and AI technologies could provide value for solving real-world problems.

University of Michigan's Rich Simpson demonstrated the NavChair Assistive Navigation System, a "smart wheelchair," which uses a suite of sonar sensors to provide navigation assistance to the wheelchair operator. The NavChair has particular value for individuals with only gross motor control: the user indicates the general direction of travel using a joystick, and the wheelchair takes care of fine corrections to avoid obstacles along the way. Rich also demonstrated a voice recognition component integrated into the robotic wheelchair, which allowed the computer to accept directions verbally rather than through the traditional joystick. One of the reasons that voice navigation of a wheelchair has not been practical in the past is that it is difficult to convey fine navigation corrections through voice commands. The navigation assistance provided by the NavChair allows for the successful integration of voice commands by effectively handling the necessary fine control automatically.

Iowa State University (Chad Bouton, Richard Cockrum, Deven Hubbard, Brian Miller, Kelly Rowles, Sophia Thrall) demonstrated Cybot, a six foot tall, 200 pound robot endowed with a six degree of freedom manipulator. The entire robot, manipulator included, was designed and built by Iowa State students. Cybot is capable of complex manipulation tasks such as pouring a drink from a can to a glass. Cybot performed this very task very successfully during the robot exhibitions. Like Michigan's wheelchair robot, Cybot contains on-board voice recognition allowing it to interact with members of the audience, asking them if they would like a drink, then responding to their answer appropriately.

Stanford University (Thomas Willeke and Clayton Kunz) demonstrated yet another practical skill: the ability to automatically map office buildings quickly and efficiently. During each robot exhibition session, InductoBeast was let loose in the robot competition maze with no a priori knowledge concerning the floorplan of the simulated office building. By the end of the Stanford presentation, InductoBeast successfully completed a map of the arena, displaying the map on its monitor and proceeding to stress-test its map by traveling to randomly chosen doorways. An unusual characteristic of InductoBeast is that it uses a form of induction during the map-building process, proposing the existence of hypothetical hallways during mapping based on knowledge about the symmetries that commonly occur in office buildings.

The University of New Mexico (Chaouki Abdallah, John Garcia, Dave Hickerson, Ales Hvezda, Dave Mattes, Eddie Tunstel) demonstrated another homebuilt robot, LOBOtomous. LOBOtomous was designed and built by UNM engineers for a senior-level design class with hardware loaned from Sandia National Laboratories. The five foot tall robot uses a ring of sonar sensors to avoid immediate obstacles with purely on-board computing power. During the exhibitions, LOBOtomous demonstrated the necessary robot skill of mingling safely with a crowd, smoothly avoiding the crowd while maintaining a reasonable forward speed. LOBOtomous demonstrated more than navigation, however: it wandered from person to person, prompting each individual to play hangman with the robot using a laptop computer on LOBOtomous' "head."

Newton Research Labs (Bill Bailey, Jeremy Brown, Randy Sargent, Carl Witty, Anne Wright) demonstrated their always-popular small robot cars, which can visually track colored objects using Newton Labs' own vision system. Because this team participated in the competition as well, their exhibition consisted of a more complete example of their robot's Squiggle ball chasing prowess than was possible in the competition environment. Several Squiggle balls were let loose in a smaller pen, and the Newton Labs' robot showed off its fast and unerring ability to track, chase, and grab Squiggle balls using a small gripper.

Dartmouth College (Simon Court, Ed Fein, Marjorie Lathrop, Artyom Lifshits, David Lillarama, David Zipkin) presented two inexpensive robots, Serial Killer and eSPAM. These robots, based on A.K. Peters' Rug Warrior kit, are excellent embodiments of cost- effective navigation. At a cost of less than one twentieth of the cost of the production robots at the competition, Serial Killer and eSPAM were able to navigate a portion of the competition maze effectively.

Finally, the University of Chicago brought its robot CHIP, hoping to demonstrate the ability to recognize and act on natural human gestures. Unfortunately, CHIP was lost in shipping for several days. When it finally did arrive at the exhibition, team members did not have enough time to put together a demonstration for the audience.

5 Conclusions

The AAAI mobile robot competitions provide a good yardstick with which to measure progress in the field (although they are certainly not the only measure). The first competition five years ago involved finding tall poles sticking up amongst small static obstacles. Teams could mark the poles in any way they wanted. The object was simply to visit the poles. In the most recent competition, the first task was to use a sparse map to visit two conference rooms in an office building and determine if they were occupied or not. Along the way, people could be walking and hallways and doorways could be blocked. The robots also had to estimate how long it would take them to finish the task. The second task involved picking up tennis balls and also picking up a moving "squiggle" ball and placing them in a pen. This is a significant amount of progress in only five years.

Some of this progress can be attributed to the competition itself. A core group of organizers has steadily "raised the bar" in the competition; each year adding some additional level of complexity. In addition, simply gathering some of the top robot researchers in the same room, with their robots, all tackling the same tasks creates an environment in which researchers share their ideas and their experiences. We were very excited at the number of teams that competed this year as it included new participants as well as teams that have competed for many years. We expect this to continue next year at AAAI-97, when Ron Arkin and Jim Firby organize the competition (contact Ron at arkin@cc.gatech.edu for additional information).

Every year the competition organizers are asked by spectators "Where's the AI in these robots?" Indeed, a task like catching squiggle balls at first might not seem to require much AI, as it is understood at the conference. We hope that the articles by several of the most successful teams in this issue will help people see inside the heads of the best robots and see the AI.

The essence of autonomous mobile robot research is that it forces people to connect perception to action in an intelligent fashion in order to accomplish complex tasks. Connecting perception to action in an intelligent way is at the heart of AI and the mobile robot competitions allow the community to see just how much progress has been made.

6 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 the robot teams who went through extraordinary lengths and severe sleep deprivation to bring their robots to AAAI. The Hart Corporation, based in Vancouver, Washington, graciously donated twenty four Squiggle Balls to the competition. Members of the AAAI staff who helped tremendously were Carol Hamilton, Mike Hamilton, Sara Hedberg, Rick Skalsky and Daphne Daily. Thanks to Pete Bonasso for serving as a liaison between the competition and Scientific American FRONTIERS. Robin Murphy helped tremendously in fund raising activities. Microsoft Corporation donated a computer for Internet access. The NASA Johnson Space Center Robotics, Automation and Simulation Division hosted the competition WWW page and mailing list.

References

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

[2] David Hinkle, David Kortenkamp, and David Miller. 1995 robot competition and exhibition. *AI Magazine*, 17(1), 1996.

[3] Kurt Konolige. Designing the 1993 robot competition. *AI Magazine*, 15(1), 1994.

[4] 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), 1993.

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