

The 1997 AAAI Mobile Robot Exhibition

Holly A. Yanco

■ A wide variety of robotics research was demonstrated at the 1997 American Association for Artificial Intelligence Mobile Robot Exhibition. Twenty-one robotic teams participated, making it the largest exhibition ever. This article describes the robotics research presented by the participating teams.

The 1997 American Association for Artificial Intelligence (AAAI) Mobile Robot Exhibition was held in conjunction with the 1997 AAAI Mobile Robot Competition at the Fourteenth National Conference on Artificial Intelligence (AAAI-97). Twenty-one robotic teams participated, making this the largest robot exhibition ever. See figure 1 for a photo of the exhibition participants. Since the first Mobile Robot Competition and Exhibition at AAAI-92, the exhibition has served to demonstrate robotics research that is beyond the scope of the competition tasks. Conference attendees are able to see current research of many robotics researchers. The 1997 exhibition consisted of demonstrations, a video loop, and posters.

At AAAI-97, robot demonstrations were given by Applied AI (various robots), Brown University (RAMONA), Colorado School of Mines (SILVER BULLET and BUJOLD), Iowa State University (CYBOT) (figure 2), KISS Institute for Practical Robotics (Kids' Demo and robots of the Robot Building Lab), Massachusetts Institute of Technology AI Lab (PEBBLES), MIT AI Lab and Boston College (WHEELSLEY and EAGLEEYES), MIT Leg Lab (monoped hopper), Michigan State University (SHOSLIF), Navy Center for Applied Research in Artificial Intelligence (COYOTE), Northwestern University (KLUDGE), Rob Turner (ARB1), University of Minnesota (TBMN), University of Virginia (BRUCE), and University of Waterloo (HEXOTICA). Four exhibitors participated in the video loop only: Brandeis University, McGill University, MIT AI Lab (COG) and Metrica, Inc.

A wide variety of robotics research was demonstrated at the exhibition. Some highlights are a robot that played fetch in the hall-

ways of the convention center, an urban search and rescue team, a robotic wheelchair controlled by an eye tracker, a self-stabilizing monoped hopper, and a robot controlled by voice commands and gestures. Most robots that were exhibited are described here, in alphabetical order by robot name.

BRUCE, University of Virginia

BRUCE plays hide and seek with a human-controlled opponent. The robot uses a layered architecture for integrating planning and action. It differs from the usual approach of interfacing a planner to a reactive system in a layered architecture because the reactive system is replaced with a different kind of action system. This action system uses small task-dependent representations called *markers*. This system can no longer truly be called reactive (because it has state), and it is termed a perception-action system (Brill et al. 1998).

A perception-action system selects its actions from the current values of its sensors and the state of its representation. The exhibit discussed how an agent's perception-action system can effectively use and maintain its representation. The addition of representation to the action layer of an architecture (which controls the sensor-effector subsystems) facilitates the communication of goal information from higher layers of the architecture. The hide-and-seek application showed how the agent used representation to facilitate behaviors that are difficult for reactive systems and how the highly structured representation of the perception-action system provides hooks for communication with the higher layers of the architecture.

COYOTE, Navy Center for Applied Research in Artificial Intelligence

The Navy Center for Applied Research in Artificial Intelligence demonstrated two mobile



Figure 1. Participants in the 1997 Robot Exhibition.

robot systems at the 1997 AAAI Mobile Robot Exhibition: (1) ARIEL (autonomous robot for integrated exploration and localization) and (2) INTERBOT (a multimodal interaction system). Both these systems are implemented on a NOMAD 200 mobile robot equipped with sonar, infrared, and laser range sensors.

A central dilemma in robot exploration is that for a robot to add perceptions to a map, it needs to know its own location—but for a robot to determine its location, it often needs a map. This problem is addressed with ARIEL (Yamauchi, Schultz, and Adams 1998). ARIEL combines frontier-based exploration strategy (Yamauchi 1997) with continuous localization methods (Schultz and Adams 1998).

At the robot exhibition, ARIEL explored the environment setup to resemble the interior of a house for the Home-Vacuum event. Starting with an empty map, ARIEL was able to explore this simulated home yet maintain an accurate estimate of its position, building a map that accurately represented the layout of the rooms and obstacles.

The Navy Center also demonstrated INTERBOT, a multimodal interaction system (Perzanowski, Schultz, and Adams 1998). INTERBOT allows humans to interact with robots using natural language and gestures. At the exhibition, the human operator told the robot to “go over there” and gestured with his hand; the robot moved in the corresponding direction. Then the human told the robot to “back up this far,” showing how far with his hands, and the robot moved back the appropriate distance. These interactions demonstrated how the combination of speech and gestures allows for natural interaction between humans and robots.

HEXOTICA, University of Waterloo

HEXOTICA’s design employs a combination of control philosophies, using the inverse kinematic control method from industrial robotic arms with high-level subsumptionlike architecture popular with walking robots to guide the leg movements. Thus, the robot has a degree of control and range of movement unmatched with other small walking robots but maintains the same robust and adaptable behavior. Low-level control algorithms for the legs based on industrial robotic control allow the foot to move in a straight-line path between any two points in the leg’s work envelope (an ability most walking robots of this size do not have). High-level behavioral-based programs define these vectors, which are adapted according to the desired direction of travel and obstacles encountered along the way.

Kids’ Demo, KISS Institute for Practical Robotics

The KISS Institute for Practical Robotics offered an interactive robot demo for kids who had accompanied their parents to the conference. Children attending the demonstration were able to work with many robots. All the kids (and many adults) enjoyed interacting with CAPTAIN KISS, the KISS Institute’s unofficial entry in the Hors d’Oeuvres, Anyone? event that delivered Hershey’s Kisses. Each child also had the chance to work with FIREFLY CATCHER, which is a KISS Institute educational mobile robot that is designed to show off the major subsystems of a robot and demonstrate how they

interact. Other KISS Institute educational robots, such as ROBOSKULL and EDBOT, were also available for experimentation.

Kids were asked to design and draw a robot that they thought would be useful. Designs ranged from Kate Murphy's Helper that would help save lives of accident victims to Max's robot that would do his chores to Natasha's Munchkin that would "be able to pick up rocks and crunch them and spit them out at the other side." The designs of these future AI researchers were displayed during the exhibition.

KLUDGE, Northwestern University

Northwestern University's KLUDGE (figure 3) is a demonstration of a new, integrated architecture for reasoning and sensory-motor control, called *role passing* (Horswill 1998). At AAAI-97, KLUDGE played fetch games with passers-by for several hours a day in a natural, unmodified environment near the registration area.

KLUDGE showcased four unusual technologies: (1) vision-based navigation and collision avoidance in unmodified environments, (2) adaptive color-based tracking of as many as three concurrent objects, (3) a novel real-time inference engine that supports limited quantification and modal reasoning, and (4) a simple finite-state natural language parser and instruction-following system. All systems ran concurrently in real time (10-Hertz update rate) on a low-cost 25 MIPS (million instructions per second) embedded processor with a power budget of under 10 Watts.

On each iteration of its control loop, the robot parses the next word in its input buffer (if any); reestimates the positions of any tracked objects; redetermines the truths, states of knowledge, and goal statuses of each predicate represented by the system; and produces a set of control output for the sensory-motor system designed to advance the robot toward its top-level goal. Because these output are completely recomputed every 100 milliseconds (ms), the system is responsive to changes in the environment. For example, suppose the robot is delivering a ball to a recipient, and the ball is stolen. Within 100 ms, KLUDGE automatically aborts the drive to the recipient subgoal and reinstates the acquire ball subgoal.

For its demonstration task, the robot performed the command "continually get red," causing it to drive to, and grab, its red toy ball. Conference attendees, for their part, performed the command "continually taunt robot by stealing red ball."

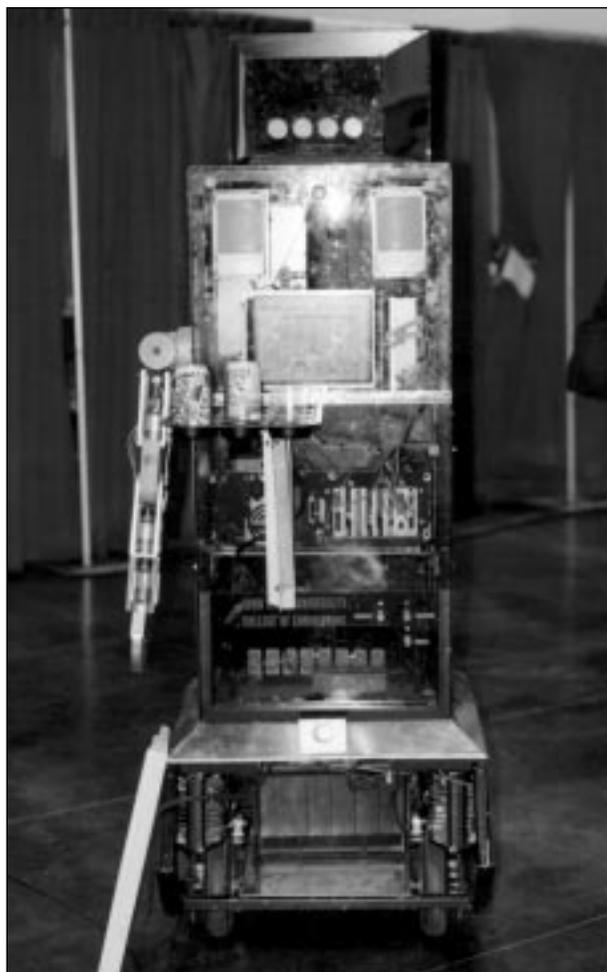


Figure 2. CYBOT from Iowa State University.

LEWIS AND CLARK, Georgia Institute of Technology

The Mobile Robot Laboratory at the Georgia Institute of Technology, directed by Ronald Arkin, demonstrated LEWIS and CLARK, a team of foraging robots. The two robots are able to track colored objects, pick them up, and deposit them in color-coded bins. They are based on the Nomadic Technologies' NOMAD 150 platform but modified by Georgia Tech to have grippers and real-time vision.

LEWIS and CLARK are two of a five-robot system used in research at Georgia Tech's Mobile Robot Laboratory. Team leader Tucker Balch is using the multirobot system to investigate how diversity and performance are linked in learning multirobot systems. The robots' control system software runs in JAVA on UNIX laptop computers mounted on the top of the robots. The robots are programmed using JAVABOTS, a new simulation and robot run-time system that supports behavior-based control system

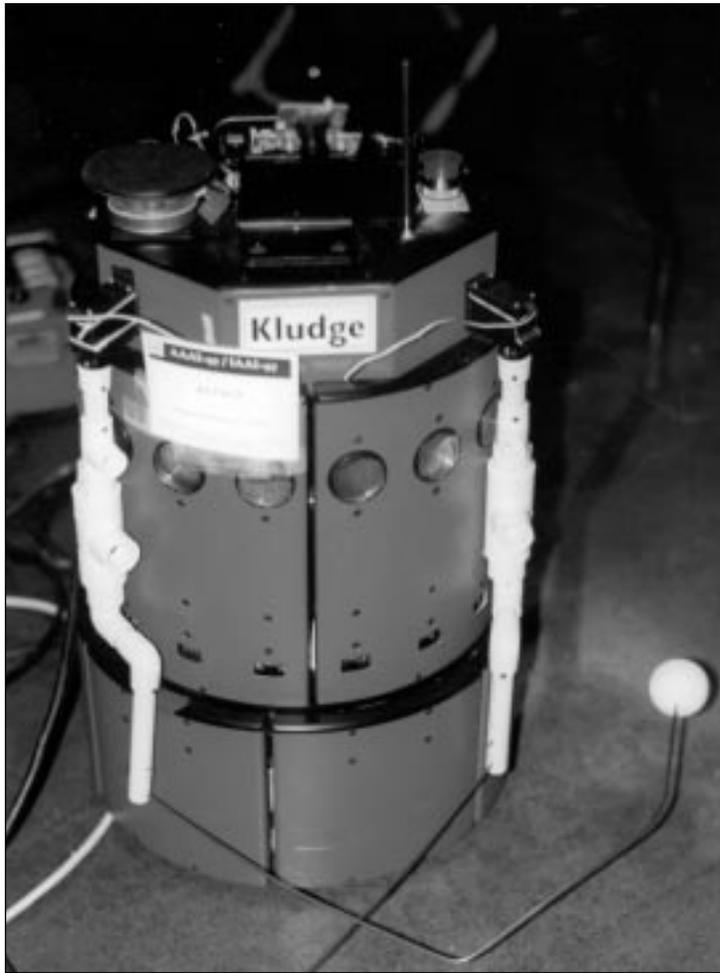


Figure 3. "Where Art Thou, Red Ball?"

KLUDGE from Northwestern University played fetch with the crowd in the hallways of the convention center.

development.

Monoped Hopper, MIT Leg Lab

The Leg Lab's monoped hopper (figure 4) is a one-legged hopping robot with no on-board computer or feedback, which illustrates the principle of self-stabilizing running. The robot was built by Robert Ringrose and David Robinson. Their exhibit combined videotape of multilegged self-stabilizing simulations with demonstrations of the physical monoped.

Legged robots can sustain stable dynamic locomotion without sensors or feedback. It is possible to construct a running robot that is inherently stable and needs no sensing to reject minor perturbations; it is therefore self-stabilizing. In contrast, most previous attempts to make robots run have used active, high-bandwidth feedback-control systems. Physically realistic simulations can run in a self-stabi-

lizing manner using one, two, and four legs. Additionally, the simplicity of a self-stabilizing mechanism makes it harder to break and easier to repair.

The robot consists of a battery pack and motor in the main body, a slider-crank mechanism to convert the motor's rotary motion into linear actuation (approximately sinusoidal), a curved foot attached to the linear actuator with a spring, and an on-off switch. There are no sensors and no computers. It hops in place consistently and reliably and recovers from external disturbances.

PEBBLES, MIT AI Lab

The PEBBLES robot being developed at the MIT AI Lab is a prototype microrover for the 2003 mission to Mars. PEBBLES's exploration tasks, which include navigation, visual exploration, and sample rock collection, are selected and performed in the complete absence of human teleoperation. The PEBBLES Mars rover is mounted on tread wheels suitable for rough-terrain navigation. The rover has a five-degree-of-freedom manipulator with a parallel jaw gripper mounted on the front end of the chassis. A single-color camera is attached to the end of this manipulator in an eye-on-hand configuration (figure 5).

PEBBLES's exploration behavior is composed of lower-level sensory motor reflexes, layered using the subsumption architecture (Brooks 1986). The lowest layer is obstacle-avoidance navigation (Lorigo, Brooks, and Grimson 1997). Layered on top of this navigation reflex is a second module for approaching a goal location marked by a visually salient target. To approach the target, the goalless navigation behavior is periodically interrupted to perform a rotation. When the goal is detected in the visual image, the robot moves toward it until an immediate obstruction forces the obstacle-avoidance module to take over again. This recursive orientation and approach algorithm successfully finds a path between the current and goal locations without explicit representation of the locations.

Higher behavior layers control manipulator motion. The joints of the manipulator use series elastic actuators (Pratt and Williamson 1995). This system enables contact with hard unstructured surfaces without damage to the motor and gearbox. In the digging algorithm, the desired end-effector position is specified to be at a certain depth below the ground. During a repeated scooping motion, if the gripper makes contact with the ground higher than the desired depth, soil will be removed.



Figure 4. Robert Ringrose with a Self-Stabilizing Monoped Hopper Developed at the MIT Leg Lab.

RAMONA, Brown University

RAMONA (figure 6) is a Real-World Interface B24 base with a single camera with a field of view. It uses the TELEOS AVP-100 vision system to extract optical flow at about 10 Hz. The system demonstrated how optical flow could be used for obstacle avoidance and a simple game of tag (Duchon, Warren, and Kaelbling 1998). For obstacle avoidance, the basic control law takes the average amount of optical flow magnitudes on each side of the image and turns the robot to the side with less flow. Given a few emergency mechanisms, this balance strategy allows the robot to wander around in complex, cluttered environments at speeds as high as 50 centimeters/second.



Figure 5. PEBBLES, a Mars Rover Prototype for the 2003 Mission to Mars, Has an Arm for Digging and Collecting Samples.

Pictured with Milyn Moy and Chandana Paul of the Massachusetts Institute of Technology AI Lab.

In the last five years, a number of groups have used similar strategies. With RAMONA, however, a number of factors need to be taken into consideration, most importantly the height of the camera (about 4 feet). Therefore, control laws for the tilt of the camera and robot speed have also been devised. These mechanisms have the robot look down and move slower in cluttered environments and look up and move faster in more open spaces.

For tag, the robot fixates moving targets, chasing them if they move away and running away if they move toward it. The robot can also be made to follow someone around, again using only optical flow. More recent work with a color camera allows the motion signals to be filtered to allow the robot to chase objects of a particular color.

More goal-directed actions are being integrated with these basic strategies. Work with simulations (Duchon 1996) has pointed to a number of methods for navigating in mazelike environments using optical flow.



Figure 6. Brown University's RAMONA Uses Optical Flow to Wander in Complex, Cluttered Environments at Speeds as High as 50 Centimeters/Second.

Pictured with Andrew Duchon of Brown University.

Robots of the 1997 Robot Building Lab

A rematch of the robots built during the 1997 Robot Building Lab (RBL-97) was held at the exhibition. Thirty people spent the Sunday and Monday prior to AAAI-97 building robots. Attendees were broken into 10 groups, and each group was given a robot kit. The kit included a four-wheel-drive mobile base, approximately 1000 pieces of LEGO Technic, a dozen sensors, several motors, and a Handy Board with software. The 36 hours were spent designing, building, programming, and testing the robots. Monday afternoon, all the robots competed in a Botball tournament, where the robot manipulated two dozen Ping-Pong balls on a 4- x 8-foot game area in a head-to-head competition. The robots had a variety of designs and strategies. One tried to lure the other robot to a false goal by using a shade and a reflector to create the illusion of the lighted target off to the side. Another simply tried to get to the target first and then sat there with a few balls over the target, blocking all latecomers. RBL-97 was run by the KISS Institute for Practical Robotics, as RBL-98 will be.

SHOSLIF, Michigan State University

Michigan State University brought its SHOSLIF-controlled robot to the exhibition. The group, headed by John Weng, has done research in computer vision and has a strong interest in sensor-based robotics. SHOSLIF has the capability to navigate in real time in an unstructured indoor environment using only one video camera. All the computation is performed in real time on board by a SUN SPARC-1 workstation equipped with an image frame grabber. Indoor navigation is different from outdoor because indoor visual features, such as floor edges, are often more unreliable than outdoor road edges. For example, they can often be occluded by passers-by or trash cans. For this reason, SHOSLIF does not use any predefined features. Instead, it uses the principal component analysis and the linear discriminant analysis to automatically derive useful features directly from the raw pixels of image frames. This approach of learning directly from raw pixels is called the appearance-based approach and has become a well-tested method for a variety of vision problems. Jim Firby, the robot competition cochair, commented during his speech at the introductory session in the conference hall that the grayish indoor scenes in which SHOSLIF navigates are quite difficult because of the lack of color and contrast, as shown by a SHOSLIF demonstration videotape that played during his speech.

In addition to autonomous navigation (Weng and Chen 1996), the SHOSLIF framework has been tested for several other tasks (Weng 1996), including vision-guide object manipulation, such as pouring a cup of milk into another cup (Hwang and Weng 1997); recognition of human faces and objects; and recognition of moving hand signs from American Sign Language. During the exhibition, the group displayed an enhanced version of the physical robot, which has a stereo setup with pan-tilt controls for each eye and a pan control for the neck as well as a robot arm.

SILVER BULLET and BUJOLD, Colorado School of Mines

Robin Murphy and two undergraduate representatives (Damian Diaz and Travis Flowers) presented marsupiallike robots for urban search and rescue (USAR). These robots are being constructed and programmed as part of a two-year National Science Research Experience for Undergraduates site grant at the Colorado School of Mines. The students built a

fully autonomous mobile robot from a 3- x 3.5-foot battery-powered children's jeep. The robot, *SILVER BULLET*, is able to navigate using sonars while it scans for signs of humans trapped in rubble by fusing visual cues (color, motion) with thermal data. If *SILVER BULLET* finds signs of a survivor, it attempts to position itself to provide the best viewing position possible for a human teleoperator, then alerts the teleoperator using radio ethernet. The teleoperator might then want to investigate further. Because many USAR sites have survivors trapped in small spaces or covered in rubble, *SILVER BULLET* might not be able to get close enough to adequately explore. Therefore, *SILVER BULLET* carries a small, rugged inspection (1- x 0.5-ft) robot, *BUJOLD*, which can be deployed and teleoperated (figure 7). *BUJOLD* actually sits inside a space in the rear of *SILVER BULLET*. When needed, *SILVER BULLET* lowers a ramp and guides *BUJOLD* out. *SILVER BULLET* can control *BUJOLD* directly, or a teleoperator can supervise. The students are exploring research issues in distributed control of heterogeneous mobile robots, human supervisory control, search patterns, and real-time sensor fusion for identifying possible survivors.

TBMN, University of Minnesota

The TBMN (trailer-backing minirobot) Project uses a neural network to learn to back a car and trailer rig to a target location by steering the front wheels of the car. A Kohonen self-organizing feature map is combined with eligibility traces to learn the appropriate output behavior for given the input values. The input into the network consist of the hitch angle and the angle of the rear trailer to the target. The output is the direction in which to turn the steering mechanism. If the rear of the trailer reaches the goal, success is signaled. If the angle to the target exceeds or the angle of the hitch exceeds, failure is signaled. The system is clocked to operate in discrete time units. No other sensory data or feedback are available. The system is designed for use on a robot that has limited computational power and has to acquire its proficiency in a small number of learning trials.

The system is implemented on a fully autonomous minirobot, TBMN, built of Legos and controlled by a HandyBoard. TBMN consists of two parts: (1) a cab unit and (2) a trailer unit. The cab is an ackerman-steering four-wheeled vehicle with rear-wheel differential drive. The front of the cab has a binary bumper switch that lets the robot know when it has contacted something from the front. The trail-



Courtesy, Holly Yanco.

Figure 7. *SILVER BULLET* and *BUJOLD* from the Colorado School of Mines Performing an Urban Search and Rescue Mission.

BUJOLD has been deployed from *SILVER BULLET* and is investigating a possible survivor.

er is attached to the cab by a hitch constructed from a single-turn potentiometer. A servo with a set of photoresistors (CdS cells) is mounted on the back of the trailer so that the robot can keep track of the target (a light bulb) that it backs up toward. The trailer also has a binary bumper switch on its rear to tell the robot when it is in contact with the target.

TBMN usually reaches good performance after 15 to 20 trials. Extensive simulation studies have determined that an average success rate greater than 85 is reached after 100 trials and greater than 95 after roughly 300 trials (Hougen, Rybski, and Gini 1998; Hougen, Gini, and Slagle 1997; Hougen et al. 1996).

Various Robots, Applied AI Systems, Inc.

Applied AI Systems, Inc. (AAI), exhibited a variety of autonomous robots at the exhibition. AAI's OCT-1B, an autonomous, multi-legged robot with intelligent behaviors that emulate those of an actual living creature, is based on the spiny lobster found off the coast of Japan. AAI uses evolutionary robotic techniques to produce the gait of OCT-1B. OCT-1B navigated its way around the demonstration area, using infrared sensors to avoid onlookers and obstacles.

AAI demonstrated the use of genetic algorithms to develop a wall-following behavior in KHEPERA robots. The KHEPERA allows interaction between the real world and simulated algorithms for trajectory execution, obstacle avoid-



Figure 8. WHEELSLEY, a Robotic Wheelchair from the MIT AI Lab, Receiving Commands from the User through EAGLEEYES, an Eye-Tracking System from Boston College.

ance, and preprocessing of sensory information. To demonstrate evolutionary robotics, AAI ran several KHEPERA in a maze. One KHEPERA was running a genetic algorithm to develop a wall-following behavior, and the other two KHEPERA robots roamed freely in the maze on an obstacle-avoidance program. By the end of the day, the KHEPERA running the genetic algorithm had produced an effective wall-following behavior that was distinctly superior to the wall-following behavior it exhibited during the early part of the day.

AAI used two LABO robots to demonstrate autonomous navigation and object detection by placing a bag of ice on a standard LABO and setting it loose to run a basic avoidance program. The other LABO, equipped with pyrosensors, was sent out to detect the first robot. In a short time, the pyrosensor-equipped LABO detected the ice on the other robot and began to follow it.

WHEELSLEY and EAGLEEYES, MIT AI Lab and Boston College

The WHEELSLEY robotic wheelchair system from the MIT AI Lab was demonstrated in conjunction with the EAGLEEYES system built at Boston College (figure 8). To drive the robotic wheelchair, the user looks at an arrow on the screen that corresponds to the desired direction (forward, right, left, or back). The robotic wheelchair executes the command using common-sense constraints, such as obstacle avoidance, until the user looks at the stop button (Yanco and Gips 1997).

The goal of the WHEELSLEY Project is the creation of a complete robotic wheelchair system to be used by people unable to drive standard powered wheelchairs (Yanco 1998). A complete robotic wheelchair system must be able to navigate indoor and outdoor environments and should switch automatically between navigation modes. For the system to be useful, it must easily be customized for the access methods required for each user.

EAGLEEYES (Gips et al. 1996) is a technology that allows a person to control the computer through five electrodes placed on the head. The electrodes measure the electrooculographic potential, which corresponds to the angle of the eyes in the head. Custom hardware and software translate the signals into coordinates on a computer screen.

Information on the Web

For more information on the robots described in this paper, see www.ai.mit.edu/people/holly/AAAI-97/ for links to each group's web site.

Acknowledgments

The following people provided information about their team's exhibition entry: Tucker Balch, Georgia Tech; Andrew P. Duchon, Brown University; Gabriel Ferrer, University of Virginia; Maria Gini, University of Minnesota; Ann Griffith, Applied AI Systems, Inc.; Ian Horswill, Northwestern University; David Miller, KISS Institute for Practical Robotics; Robin Murphy, Colorado School of Mines; Chandana Paul, MIT AI Lab; Robert Ringrose, MIT Leg Lab; Cathryne Stein, KISS Institute for Practical Robotics; John Weng, Michigan State University; and Brian Yamauchi, Navy Center for Applied Research in Artificial Intelligence.

Thanks to all the exhibition participants, Ron Arkin and Jim Firby for coorganizing the robot competition, and Rebecca Kuipers for videotaping the event. A special thanks to the AAI staff, especially Daphne Black, Annette

Eldredge, Mary Beth Jensen, Carol Hamilton, Mike Hamilton, and Rick Skalsky. None of the robot competitions and exhibitions would happen without them.

References

- Brill, F. Z.; Wasson, G. S.; Ferrer, G. J.; and Martin, W. N. 1998. The Effective Field of View Paradigm: Adding Representation to a Reactive System. *Engineering Applications of Artificial Intelligence* (Special Issue on Machine Vision for Intelligent Vehicles) 11(2): 189–201.
- Brooks, R. A. 1986. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation* 2(1): 14–23.
- Duchon, A. P. 1996. Maze Navigation Using Optical Flow. In *From Animals to Animats 4: Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, eds. P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, 224–232. Cambridge, Mass.: MIT Press.
- Duchon, A. P.; Warren, W. H.; and Kaelbling, L. P. 1998. Ecological Robotics. *Adaptive Behavior* 6:3–40.
- Gips, J.; DiMattia, P.; Curran, F. X.; and Olivieri, P. 1996. Using EAGLEEYES—An Electrode-Based Device for Controlling the Computer with Your Eyes—to Help People with Special Needs. In *Interdisciplinary Aspects on Computers Helping People with Special Needs*, eds. J. Klaus, E. Auff, W. Kremser, and W. Zagler, 77–83. Munich: R. Oldenbourg.
- Horswill, I. 1998. Grounding Mundane Inference in Perception. In *Autonomous Robots*, Volume 5. New York: Kluwer Academic. Forthcoming.
- Hougen, D. F.; Rybski, P.; and Gini, M. 1998. Repeatability of Real-World Training Experiments: A Case Study. Paper presented at the AAAI Spring Symposium on Integrating Robotics Research: Taking the Next Leap, 23–25 March, Stanford, California.
- Hougen, D. F.; Gini, M.; and Slagle, J. 1997. Partitioning Input Space for Control Learning. Paper presented at the International Conference on Neural Networks, 9–12 June, Houston, Texas.
- Hougen, D. F.; Fischer, J.; Gini, M.; and Slagle, J. 1996. Fast Connectionist Learning for Trailer Backing Using a Real Robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1917–1922. Washington, D.C.: IEEE Computer Society.
- Hwang, W., and Weng, J. 1997. Vision-Built Robot Manipulator Control as Learning and Recall Using SHOSLIF. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2862–2867. Washington, D.C.: IEEE Computer Society.
- Lorigo, L. M.; Brooks, R. A.; and Grimson, W. E. L. 1997. Visually Guided Obstacle Avoidance in Unstructured Environments. In *Proceedings of the IEEE Conference on Intelligent Robots and Systems*, 373–379. Washington, D.C.: IEEE Computer Society.
- Perzanowski, D.; Schultz, A. C.; and Adams, W. 1998. Integrating Natural Language and Gesture in a Robotics Domain. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*. Washington, D.C.: IEEE Computer Society. Forthcoming.
- Pratt, G. A., and Williamson, M. W. 1995. Series Elastic Actuators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-95)*, 399–406. Washington, D.C.: IEEE Computer Society.
- Schultz, A. C., and Adams, W. 1998. Continuous Localization Using Evidence Grids. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, 2833–2839. Washington, D.C.: IEEE Computer Society.
- Weng, J. 1996. CRESCETRON and SHOSLIF: Toward Comprehensive Visual Learning. In *Early Visual Learning*, eds. S. K. Nayar and T. Poggio, 183–214. New York: Oxford University Press.
- Weng, J., and Chen, S. 1996. Incremental Learning for Vision-Based Navigation. Paper presented at the International Conference on Pattern Recognition, 25–30 August, Vienna, Austria.
- Yamauchi, B. 1997. A Frontier-Based Approach for Autonomous Exploration. In *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 146–151. Washington, D.C.: IEEE Computer Society.
- Yamauchi, B.; Schultz, A.; and Adams, W. 1998. Mobile Robot Exploration and Map Building with Continuous Localization. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, 3715–3720. Washington, D.C.: IEEE Computer Society.
- Yanco, H. A. 1998. WHEELSLEY, A Robotic Wheelchair System: Indoor Navigation and User Interface. In *Lecture Notes in Artificial Intelligence: Assistive Technology and Artificial Intelligence*, eds. V. Mittal, H. A. Yanco, and J. Aronis, 261–273. New York: Springer-Verlag.
- Yanco, H. A., and Gips, J. 1997. Preliminary Investigation of a Semiautonomous Robotic Wheelchair Directed through Electrodes. In *Proceedings of the Rehabilitation Engineering Society of North America 1997 Annual Conference*, ed. S. Sprigle, 414–416. Arlington, Va.: RESNA.



Holly Yanco is a doctoral candidate at the Massachusetts Institute of Technology (MIT) in the Artificial Intelligence Laboratory. Her thesis research involves creating a robotic wheelchair system that will work in both indoor and outdoor environments. She received an M.S. in computer science from MIT in 1994. From September 1994 to January 1996, she was an instructor in the Computer Science Department at Wellesley College. Yanco is coeditor of a collection of papers on assistive technology and artificial intelligence. She organized the AAAI-97 Robot Exhibition and is serving on the rules committee for the Hors d'Oeuvres event at the AAAI-98 Robot Competition. Her e-mail address is holly@ai.mit.edu.



Android Epistemology

Edited by **Kenneth M. Ford, Clark Glymour, & Patrick J. Hayes**

Epistemology has traditionally been the study of human knowledge and rational change of human belief. *Android epistemology* is the exploration of the space of possible machines and their capacities for knowledge, beliefs, attitudes, desires, and action in accord with their mental states. From the perspective of android epistemology, artificial intelligence and computational cognitive psychology form a unified endeavor: artificial intelligence explores any possible way of engineering machines with intelligent features, while cognitive psychology focuses on reverse engineering the most intelligent system we know, us. The editors argue that contemporary android epistemology is the fruition of a long tradition in philosophical theories of knowledge and mind.

The sixteen essays by computer scientists and philosophers collected in this volume include substantial contributions to android epistemology as well as examinations, defenses, elaborations, and challenges to the very idea.

Contributors include Kalyan Basu, Margaret Boden, Selmer Bringsjord, Ronald Chrisley, Paul Churchland, Cary deBessonet, Ken Ford, James Gips, Clark Glymour, Antoni Gomila, Pat Hayes, Umar Khan, Henry Kyburg, Marvin Minsky, Anatol Rapoport, Herbert Simon, Christian Stary, and Lynn Stein.

ISBN 0-262-06184-8 336 pp., index. \$25.00 hardcover

The AAAI Press • Distributed by The MIT Press

Massachusetts Institute of Technology, 5 Cambridge Center, Cambridge, Massachusetts 02142

To order, call toll free: (800) 356-0343 or (617) 625-8569. MasterCard and VISA accepted.