

# The Mobile Robot RHINO

*Joachim Buhmann, Wolfram Burgard,  
Armin B. Cremers, Dieter Fox,  
Thomas Hofmann, Frank E. Schneider,  
Jiannis Strikos, and Sebastian Thrun*

■ RHINO was the University of Bonn's entry in the 1994 AAAI Robot Competition and Exhibition. RHINO is a mobile robot designed for indoor navigation and manipulation tasks. The general scientific goal of the RHINO project is the development and the analysis of autonomous and complex learning systems. This article briefly describes the major components of the RHINO control software as they were exhibited at the competition. It also sketches the basic philosophy of the RHINO architecture and discusses some of the lessons that we learned during the competition.

RHINO, shown in figure 1, is a b21 mobile robot platform manufactured by Real-World Interface. It is equipped with 24 sonar proximity sensors, a dual-color camera system mounted on a pan-tilt unit, and 2 onboard i486 computers. Sonar information is obtained at a rate of 1.3 hertz (Hz), and camera images are processed at a rate of 0.7 Hz. RHINO communicates with external computers (two Sun SPARCSTATIONS) by a tetherless Ethernet link.

The RHINO project is generally concerned with the design of autonomous and complex learning systems (Thrun 1994). The 1994 AAAI Robot Competition and Exhibition, sponsored by the American Association for Artificial Intelligence (AAAI), ended an initial six-month period of software design. Key features of RHINO's control software, as exhibited at the competition, are as follows:

**Autonomy:** RHINO operates completely autonomously. It has been operated repeatedly for durations as long as one hour in populated office environments without human intervention.

**Learning:** To increase the flexibility of the software, learning mechanisms support the adaptation of the robot to its sensors and the environment. For example, neural network

learning is employed to interpret sonar measurements.

**Real-time operation:** To act continuously in real time, any-time solutions (Dean and Boddy 1988) are employed wherever possible. Any-time algorithms are able to make decisions regardless of the time spent for computation. The more time that is available, however, the better the results are.

**Reactive control and deliberation:** RHINO's navigation system integrates a fast, reactive on-board obstacle-avoidance routine with knowledge- and computation-intense map building and planning algorithms.

RHINO's software consists of a dozen different modules. The interface modules (a base-sonar sensor interface, a camera interface, and a speech interface) control the basic communication to and from the hardware components of the robot. On top of these, a fast obstacle-avoidance routine analyzes sonar measurements to avoid collisions with obstacles and walls at a speed as high as 90 centimeters a second. Global metric and topological maps are constructed on the fly using a neural network-based approach combined with a database of maps showing typical rooms, doors, and hallways. RHINO employs a dynamic programming planner to explore unknown terrain and navigate to arbitrary target locations. It locates itself by continuously analyzing sonar information. In addition, a fast vision module segments images from two color cameras to find target objects and obstacles that block the path of the robot. RHINO's control flow is monitored by an integrated task planner and a central user interface.

The integration of a dozen different software modules, which all exhibit different timing and response characteristics, requires a flexible scheme for the flow and synchronization of the control software.



*Figure 1. The RHINO Robot from the University of Bonn.*

tion of information. The key principles for the design of RHINO's software are as follows:

Distributed control and communication: Each module communicates with several other modules through Ethernet (Fedor 1993). There is no single control unit, and communication is not centralized.

Asynchronous communication: RHINO's software lacks a central clock. Each of the modules runs independently of the other modules. To resolve conflicts, certain modules (such as the on-board obstacle-avoidance module) can take priority over other modules (such as the planner) in determining the robot's motion direction.

Software fault tolerance: RHINO's software is designed to accommodate sudden failures of most of its software components. Almost all modules can be stopped and restarted at any time. Effective mechanisms ensure that restarted modules will immediately obtain the currently available global information.

The following sections present some of the key components of the RHINO approach in more detail: (1) the obstacle-avoidance mod-

ule, (2) the modules concerned with sensor interpretation and map building, (3) the planner and explorer, and (4) the visual routines. The article concludes with a discussion that highlights some of the lessons that were learned during the AAAI competition.

## Fast Obstacle Avoidance

The obstacle avoidance runs on board, independent of other software components such as the planner. Every 0.25 seconds, a new velocity and motion direction are chosen according to the most recent sonar measurements. To rapidly adapt to new situations, only the last three sonar sweeps are considered. RHINO can react immediately to changes in the environment and hard-to-see and moving obstacles such as humans.

The obstacle-avoidance module controls both the velocity and the motion direction of the robot. At every instant in time, the velocity is determined such that no collision will occur within the next two seconds (two-second rule). The motion direction is determined based on target points, which are generated by the planner (see discussion later). To reach a given target, the robot can choose among different trajectories on which it will travel with different velocities. RHINO selects its motion direction by maximizing its translational velocity, denoted by  $v$ , while minimizing the angle to the target point, denoted by  $\theta$ .

To determine  $v$ , a simplified model of the robot's environment is constructed. Proximity information, obtained from RHINO's sonar sensors, is used to construct a two-dimensional obstacle line field. Every sonar reading is converted to a line in this field, as depicted in figure 2. To avoid collisions with obstacles, the obstacle-avoidance routine considers a variety of circular trajectories, one of which is shown in figure 2. For each trajectory, the distance between the robot and the closest obstacle line along the projected trajectory is computed. This distance determines the translational velocity  $v$ , according to the two-second rule. The projected angle to the target point,  $\theta$ , is calculated for the estimated robot position and orientation after 0.25 seconds. For both values  $v$  and  $\theta$ , a smoothed histogram is constructed. Because of the dynamic constraints, only a small number of trajectories are reachable within the next 0.25 seconds and are consequently considered in the histogram. Finally, the trajectory that maximizes a weighted difference of  $v$  and  $\theta$  is chosen. To increase the safety of the robot, a

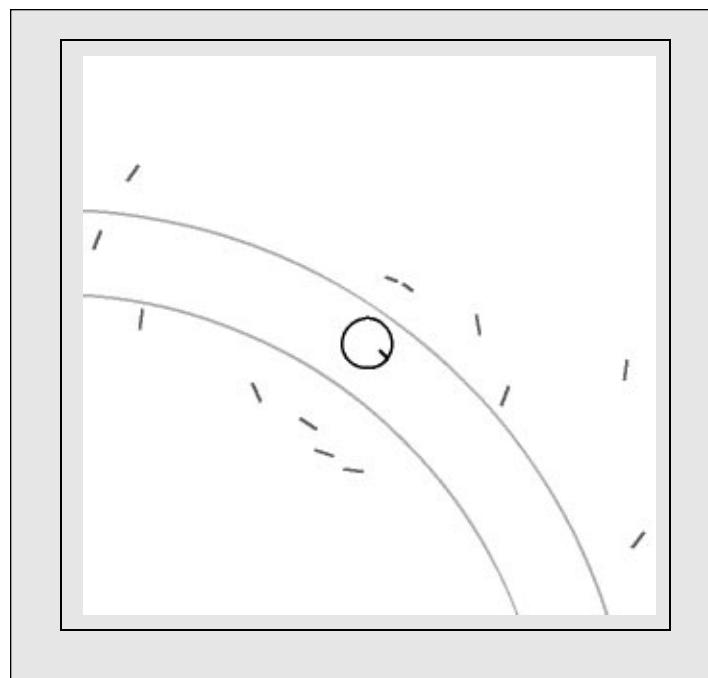
security distance of 10 centimeters is kept to surrounding objects. This security distance is increased to as much as 30 centimeters as the robot's velocity increases.

RHINO's obstacle-avoidance approach is easily extendible to other sensors. For example, prior to the competition, we successfully employed camera information to identify small obstacles on the floor that block the path of the robot, as described later. Each visually detected obstacle is mapped into a few lines in the obstacle field, much like the sonar information described earlier. However, visual information was not used by the obstacle-avoidance routine during the AAAI competition, basically because sonar information was fast and accurate enough in the competition ring.

## Map Building and Position Control

RHINO's global-navigation system builds and uses occupancy maps of the robot's environment. More specifically, when traveling through possibly unknown terrain, RHINO interprets its sonar readings to generate a two-dimensional, discrete probabilistic occupancy map. Sonar sensors are interpreted using an artificial neural network, which estimates the likelihood of occupancy of any point in a three-meter circle around the robot (Thrun 1993). Multiple measurements are integrated using Bayesian inference (Moravec 1993). Figure 3a shows a map that was constructed while we manually steered the robot through the competition arena. This map describes an area of approximately 30 x 20 meters. The hallways, rooms, large obstacles, and doors can clearly be recognized.

To navigate based on global metric information, it is imperative that the robot be able to locate itself accurately in its map. RHINO is equipped with fairly accurate wheel encoders. However, even small angular errors in dead reckoning can have devastating effects on the internal position estimation. To compensate for such error, the robot continuously matches its current sonar readings with its global occupancy map. If a mismatch is found between the occupancy map and the obstacles predicted based on the most recent sonar sweep, the internal position is corrected accordingly. In addition, RHINO registers the angular orientation of walls with respect to its current location to correct more accurately for rotational errors. This mechanism, which rests on the assumption that walls are typically perpendicular or parallel to each other, has

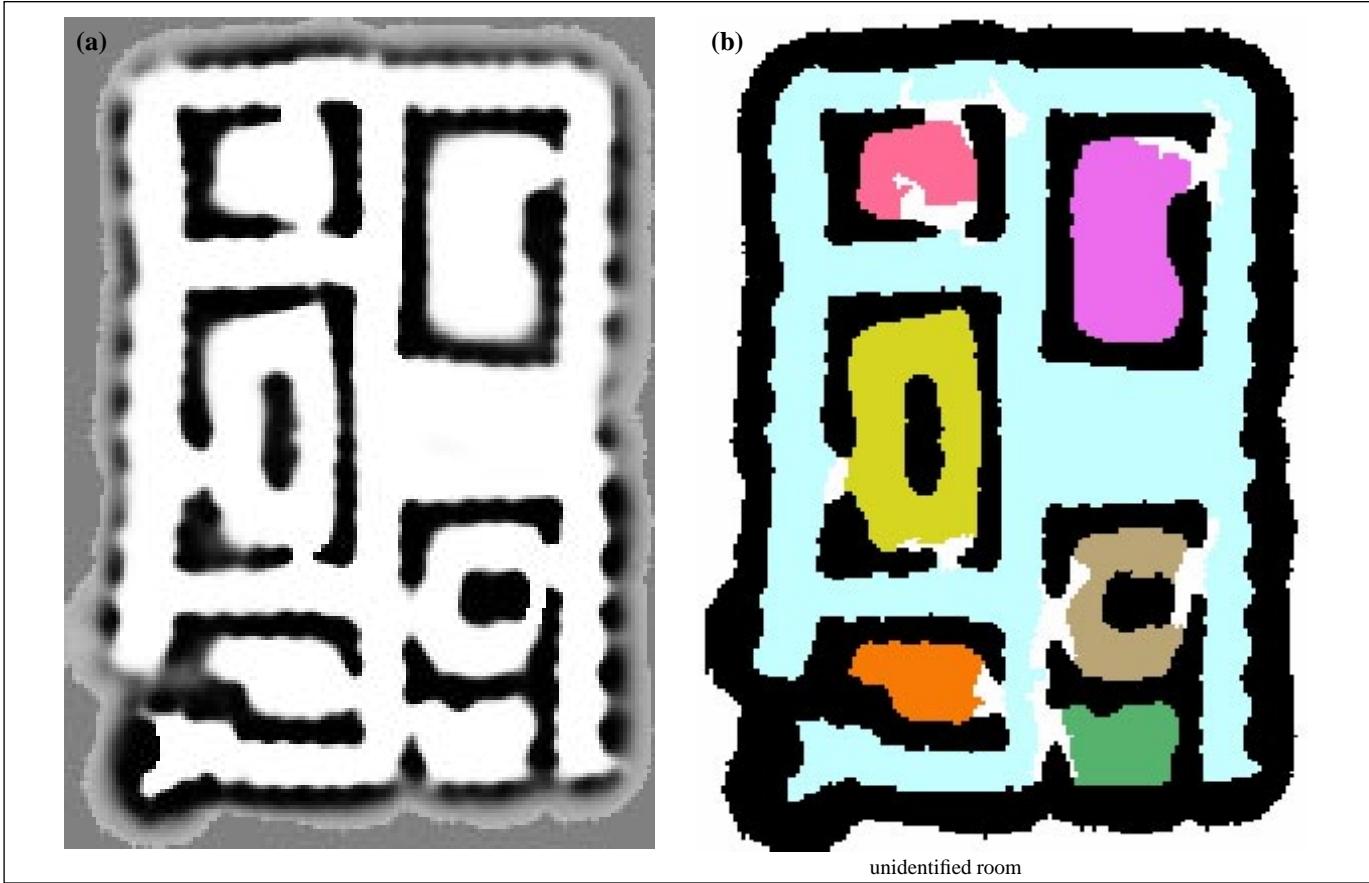


*Figure 2. Obstacle Line Field.*

Each sonar reading is indicated by a line centered around the robot. The trajectory, which is finally chosen by RHINO, is also shown.

been found to be effective for the detection of rotational errors at the competition as well as in various office environments. If RHINO operates some 30 minutes with velocities as high as 90 centimeters a second in unknown terrain, the total error is usually smaller than 30 centimeters. Without correcting the dead reckoning, this error often accumulates as much as 30 meters.

To obtain topological information concerning the location of rooms, doors, and hallways, RHINO analyzes its metric occupancy map continuously. Walls are identified by thresholding. A large database of examples of door regions, hallways, and rooms (and parts thereof) is continuously matched to assign topological labels to the unoccupied areas in the occupancy map. By analyzing the connectivity of the labeled map, RHINO is able to identify doors, hallways, and rooms. An example of a topologically labeled map is shown in figure 3b. This map, which is based on the metric map shown in figure 3a, subdivides the terrain into seven rooms or hallways (gray) and nine door regions (white). As is easy to see, most of the rooms and hallways have been identified correctly. In the bottom left corner of the figure, however, a small room has not been identified: Because of sensor noise, the occupancy map failed to cap-



*Figure 3. Two Views of the Competition Ring.*

A. This map of the competition ring was constructed from sonar measurements. Bright regions indicate free space, and dark regions indicate walls and obstacles. Walls and obstacles are enlarged by a robot diameter. B. Shown here is a topological analysis of the map. Obstacles are shown in black. As indicated by the different shading, the free space is divided into seven room-hallways (gray) and 9 door regions (white). The arrow points to an unidentified room in the competition ring.

ture a small wall—a problem that might particularly occur with thin walls, such as those found at the competition.

The topological map analyzer works continuously. At any point in time, it can be queried to output a topological map. However, the quality of the topological maps increases with time. The generation of the labels shown in figure 3b requires approximately 15 minutes of processing time on a Sun SPARCSTATION 10. Note that the underlying database of topological examples consists of preselected prototypes based on occupancy maps that were constructed at the University of Bonn prior to the competition.

### Planning and Exploration

In this section, we describe how occupancy maps are used when controlling the robot. RHINO's planner generates minimum-cost

paths to arbitrary goal locations or, as described later, to unexplored regions. These paths are constantly refined and communicated to the obstacle-avoidance routine, which then determines the final motion direction and velocity of the robot.

RHINO's main planning engine consists of a *dynamic programming routine*, which computes trajectories with minimum cost to a goal location (Howard 1960). The occupancy map is translated into a cost function, such that occupied territory results in a high traversal cost and free territory in a low traversal cost. Dynamic programming propagates path information from the goal(s) to arbitrary locations on the map. Consequently, steepest descent results in a minimum-cost path to the “cost-nearest” goal. Control can be generated at any time without any significant computation. However, deliberation time is traded for the quality of the resulting path.

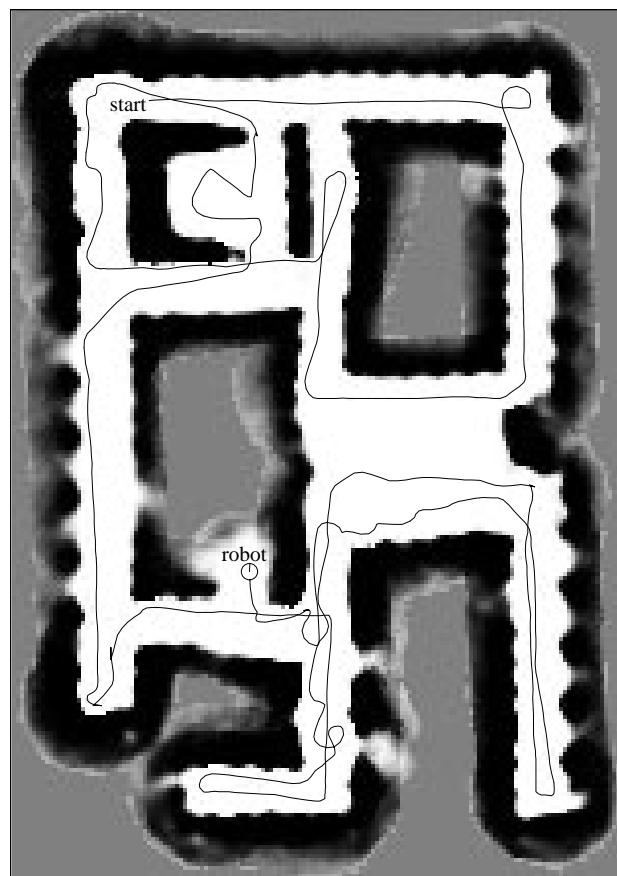
Because occupancy maps are often too inaccurate to generate collision-free motion control, in dynamic environments, RHINO's planner commands only the rough motion direction, which is then finalized by the collision-avoidance routine. Consequently, if unmodeled obstacles block the robot's path, the planner is faced with unexpected robot actions. Dynamic programming preplans for arbitrary robot locations. Goal information is propagated for every location in the map, not just the current location of the robot. Consequently, RHINO can quickly react if it finds itself to be in an unexpected location and generate appropriate motion directions without any additional computational effort. This rapid exception-handling ability provides the necessary freedom for the collision-avoidance routine to modify actions commanded by the planner at its own will.

In both stages of the competition, RHINO explored and mapped unknown terrain. RHINO's planning mechanism can easily be applied to generate explorative paths, lacking a specific goal point. If the set of goal positions is defined as the set of positions for which no map information is available, RHINO moves straight to the unexplored. Figure 4 illustrates the path of some 15 minutes of autonomous robot exploration in the competition ring. In this prototypical example, the main hallways have already been traversed, and RHINO continues to explore the unexplored rooms. RHINO's speed on the straight-line segments of the exploration path was generally between 50 and 90 centimeters a second. Further details on planning and exploration can be found in Thrun (1993).

## Vision and Object Recognition

The images from the color camera system are the input to a four-stage vision system, which solves two different tasks: First, it has to recognize important objects typically found in the environment (for example, objects in an office environment). Second, it supplies valuable information for the robot-navigation task by providing local occupancy maps to the map builder and obstacle locations relative to the robot's position to the collision-avoidance module. This second task, however, was not performed during the final runs at the competition. Here, map building and obstacle avoidance relied solely on sonar information, which turned out to be sufficiently reliable in the competition ring.

In the first stage of low-level processing, images are low-pass filtered and subsampled

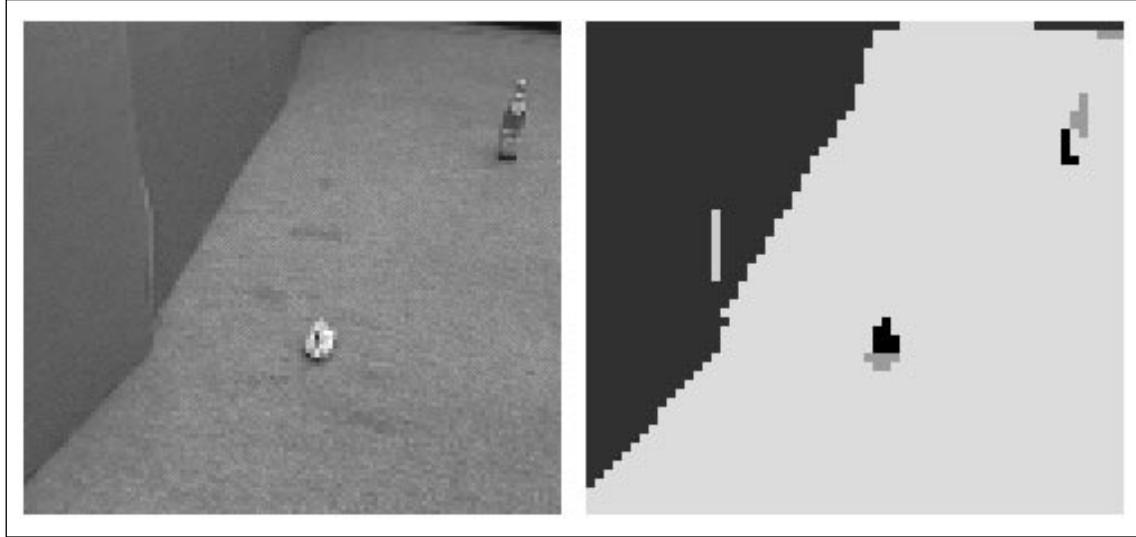


*Figure 4. Occupancy Map Constructed from Scratch during 15 Minutes of Autonomous Robot Exploration.*

The robot's path, which starts at the upper left corner, is also shown.

to reduce the data transfer by the radio link and preprocess the image for the next stage. This process is performed on one of the on-board i486 computers. Sampling in space (image size) and in time (frame rate) is done dynamically, dependent on the actual velocity of the robot. Thereby, the transmission-channel capacity is allocated in a task-driven way.

The second processing stage is done by an image-segmentation algorithm that partitions the transmitted image into homogeneous, connected regions (figure 5). Homogeneity is measured by a dissimilarity measure between neighboring image sites (pixels or blocks of pixels). For reasons of efficiency, the dissimilarity measure is restricted to a weighted squared sum of color and luminance differences between sites (with an additional threshold). Formally, the segmentation task can be described as a minimization problem of a cost function, which sums up the local



*Figure 5. In the Second Processing Stage, the Transmitted Image Is Partitioned.*

A. The raw image. B. The segmented image in coarse resolution (9 segments).

inhomogeneities of all regions for a given partition. To achieve real-time performance without the need for special hardware, the segmentation is implemented by a fast region-merging scheme. The decision about whether two neighboring regions should be merged depends on a comparison between the current costs and the costs after merging.

The third stage takes the segmented image as an input and seeks to identify and label certain elements of a typical indoor scene, for example, the floor, walls, and doors. Important additional information for both navigation and object recognition can be derived: The distances and sizes of all objects or regions located on the floor are calculated based on knowledge of the position, viewing angle, and so on, of the cameras. This distance and size information for walls and objects is incorporated into the occupancy map.

At the top of RHINO's vision-processing architecture, a feature-based object recognizer detects objects of interest in the environment. The recognition module is able to learn from labeled examples of feature vectors extracted from example images. For every type object, a Gaussian model (mean and covariance matrix) is estimated according to the maximum likelihood principle. Typical features are the normalized mean and the variance of object luminance; the mean and the variance of object color (hue and saturation); and geometric features such as the absolute size, width, and height of the object,

estimated based on the object location calculated in the third stage of the vision system. At the competition, we used between 30 and 50 training examples to model each of the 7 object classes. These example images showed different objects varying in object distance, lighting conditions, and the choice of the class representative. The assignment to a class is done by minimizing the Mahalanobis distance to the class mean (Duda and Hart 1973). If the likelihood is below a certain threshold, the candidate object is not accepted as a member of a known class and is considered unclassified.

Once a target object is found, its location is communicated to the planner and other modules concerned with task and motion control.

## Results and Discussion

This article surveys the software architecture of the fully autonomous RHINO robot, as it was exhibited at the 1994 AAAI Robot Competition and Exhibition. RHINO is controlled by a dozen software modules that work and communicate asynchronously. Special emphasis is put on real-time operation, learning, and the integration of reactivity and global map knowledge. RHINO does not require prior knowledge on the locations of walls or obstacles or the topology of its environment.

In the first stage of the competition (Office Delivery), RHINO had to move to a designated target location. This stage consisted of three

trials, two of which counted for the final score. Because we are specifically interested in navigation without prior information, we attempted to use the first trial for exploring and mapping the competition ring and the remaining two trials for the delivery task. However, although RHINO traveled fast, we learned that the arena could not be explored completely without prior information in the allotted time. Consequently, we had to "buy" a metric map for subsequent trials.

One of the major problems we encountered at the competition was RHINO's unreliable radio link. Unpredictable radio communication, possibly based on interference with other radio links, caused RHINO's on-board operating system (LINUX, a PC-version of UNIX) to suspend obstacle avoidance for periods of 10 seconds or more. In the second and third trial of stage one, RHINO suffered from severe communication failures and collided repeatedly with walls. Thus, the first stage of the competition could not be completed, and RHINO was excluded from the finals of this stage.

The communication problem was fixed in the second stage of the competition by moving the radio link closer to the competition ring. In this stage (Office Cleanup), RHINO was required to find and fetch objects such as soda cans and paper wads, pick them up, and drop them in groups of three into a nearby trash bin. Because RHINO is currently not supplied with a manipulator, it indicated its intention to pick up and drop objects by voice. RHINO used essentially the same exploration routines as in the first stage but at a reduced speed. In addition, the visual routines described earlier were employed for the identification of obstacles. At the competition, RHINO found most of the objects in the starting room and then continued to clean up the hallway. Here, RHINO scored second, defeated only by a collaborating team of three robots (see article by Tucker Balch and his colleagues, also in this issue).

The AAAI competition ended an initial six-month period of software engineering. RHINO's software is generally applicable to autonomous navigation in indoor environments. In the future, RHINO will operate 24 hours a day, interrupted only by battery charging. Our main scientific interest is the study and the design of autonomous, complex learning systems, which, in the domain of robotics, includes adaptive approaches to sensory processing and lifelong robot learning (Thrun 1994). We are currently implementing various learning techniques that allow RHINO to adapt to new situations and

acquire new skills necessary for achieving a broad variety of tasks.

### Acknowledgments

Some of the low-level software (TCX [Fedor 1993], device drivers for the speech board and the cameras) was provided by Carnegie Mellon University, which is gratefully acknowledged. We also acknowledge the steady and helpful support of Real-World Interface. Travel to the competition would not have been possible without a generous travel grant from AAAI and the invaluable assistance of Peter Lachart, Wolli Steiner, and Peter Wallossek. The work of Hofmann was supported in part by the Ministry of Science and Research of the state of North Rhine-Westphalia.

### References

- Dean, T. L., and Boddy, M. 1988. An Analysis of Time-Dependent Planning. In Proceedings of the Seventh National Conference on Artificial Intelligence, 49–54. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Duda, R. O., and Hart, P. E. 1973. *Pattern Classification and Scene Analysis*. New York: Wiley.
- Fedor, C. 1993. TCX: An Interprocess Communication System for Building Robotic Architectures. Programmer's Guide to Version 10.xx. Pittsburgh, Penn.: Carnegie Mellon University.
- Howard, R. A. 1960. *Dynamic Programming and Markov Processes*. Cambridge, Mass.: MIT Press.
- Moravec, H. P. 1988. Sensor Fusion in Certainty Grids for Mobile Robots. *AI Magazine* 9(2): 61–74.
- Thrun, S. B. 1994. A Lifelong Learning Perspective for Mobile Robot Control. In Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems, 23–30. Washington, D.C.: IEEE Computer Society.
- Thrun, S. B. 1993. Exploration and Model Building in Mobile Robot Domains. In Proceedings of the ICNN-93, 175–180. Washington, D.C.: IEEE Computer Society.



Joachim Buhmann is an associate professor of computer science at the University of Bonn. He received a Ph.D. in theoretical physics from the Technical University of Munich in 1988. He has held postdoctoral positions at the University of Southern California and the Lawrence Livermore National Laboratory. His current research interest focuses on the theory of neural networks and their applications to image understanding and signal processing.

Wolfram Burgard is a research associate at the University of Bonn. He received his Ph.D. in computer science from the University of Bonn in 1991

and a Master's degree from the University of Dortmund. Burgard's research focuses on expert systems, fuzzy logic, logic programming, and robotics.



Armin B. Cremers is a full professor of computer science and chair of the Computer Science Institute at the University of Bonn. Prior to this position, he was a full professor at the University of Dortmund and an assistant professor at the University of Southern California.

His scientific research is in the areas of database and information systems and AI.



Dieter Fox received a Bachelor's degree in 1990 and a Master's degree in 1993, both from the University of Bonn. Currently, he is enrolled in the Ph.D. program in computer science at the University of Bonn. His main scientific interests are in robotics, map building, and collision avoidance.



Thomas Hofmann received a Bachelor's degree in computer science in 1989 and a Master's degree from the University of Bonn in 1993. Currently, he is a Ph.D. student in the Computer Science Department at the University of Bonn. His main scientific interests are in computer vision, learning theory, and statistics.



Frank E. Schneider received a Master's degree from the University of Bonn in 1994. Since 1989, he has worked for the Research Institute of Applied Science in Wachtberg, Germany. His scientific interests are robotics; virtual reality; and vision and simulation, in particular, the visualization of sensory data using virtual reality techniques and video compression.



Jiannis Strikos received a Bachelor's degree in computer science from the University of Bonn in 1993. Since this time, he has been an active member of the Computer Vision Group. His current research is in the area of active vision, visual search, selective attention, and real-time face recognition.

Sebastian Thrun received a Bachelor's degree from the University of Hildesheim in 1988 and a Master's degree from the University of Bonn in 1992. Currently, he is enrolled in the Ph.D. program in computer science at the University of Bonn and is also affiliated with Carnegie Mellon University.



Thrun's main scientific interests are AI, machine learning, and computers on wheels. Thrun led the RHINO team at the 1994 AAAI Robot Competition and Exhibition.

## Florida AI Graduate (MS) Student Fellowship

The Institute for the Interdisciplinary Study of Human & Machine Cognition (IHMC) at the University of West Florida invites your application for the Fall of 1995. This is a special research fellowship restricted to Florida residents.

The IHMC was founded with legislative support in 1990 as an interdisciplinary research unit. Additionally, IHMC has succeeded in securing substantial extramural support and has established an enviable research and publication record. Interdisciplinary research is underway in the computational and philosophical foundations of AI, computers in education, knowledge-based systems, neural networks, internet-based autonomous agents, cognitive psychology, and reasoning under uncertainty.

The University is situated in a 1000-acre protected nature preserve bordering the Escambia River, and is approximately 14 miles north of the country's finest white-sand beaches. New Orleans is only 3 hours away by car.

The successful applicant(s) will have a BS in Computer Science with outstanding academic performance, a GRE score of at least 1400, and a strong interest in one or more of the following areas: AI, cognitive science, human/machine interaction, or any of several other areas.

The Florida AI Graduate Student Fellowship includes a tuition-waiver and a half-time graduate research assistantship at the IHMC while pursuing a MS in Computer Science.

Applicants should email <kford@ai.uwf.edu> their vita and a letter of application to Dr. Kenneth Ford, Director, Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32514.