Experiences with two Deployed Interactive Tour-Guide Robots

S. Thrun[1]**, M. Bennewitz**[2]**, W. Burgard**[2]**, A.B. Cremers**[2]**, F. Dellaert**[1]**, D. Fox**[1]**, D. Hähnel**[2]
G. Lakemeyer[3]**, C. Rosenberg**[1]**, N. Roy**[1]**, J. Schulte**[1]**, D. Schulz**[2]**, W. Steiner**[2]

[1]School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

[2]Computer Science Department III
University of Bonn
Bonn, Germany

[3]Computer Science Department
Aachen University of Technology
Aachen, Germany

Abstract

This paper describes and compares two pioneering mobile robot systems, which were recently deployed as interactive tour-guides in two museums. Both robots demonstrated safe and reliable navigation in proximity of people. They also interacted with museum visitor through various means, including the Web. Probabilistic algorithms and learning are pervasive in their software architectures. This article sketches the basic software, summarizes results, compares the robots, and discusses open problems.

1 Introduction

In the near future, an increasing number of service robots will have to work in close proximity to people, interact with them, and perform tasks in populated environments, highly dynamic and unpredictable. Problems of this sort arise in a range of application domains, including janitorial services, personal service robots, information kiosks, and robots in the health care sector (e.g., nursing robots).

This paper describes the collective experience with two deployed museum tour-guide robots. The first robot, Rhino (shown in Fig. 1), operated for six days in May 1997 in the Deutsches Museum Bonn. The second robot, Minerva (Fig. 2), was installed for a total of 14 days in the Smithsonian's National Museum of American History (NMAH), during August/September of 1998. Both robots are between 1.2 and 1.5 meters in height, and they are equipped with laser range finders, cameras, sonar sensors, and tactile sensors. Their tasks involve approaching people, interacting with them by replaying pre-recorded messages and displaying texts and images on on-board displays as well as safe and reliable navigation in un-modified and populated environments. The setting is uniquely suited to study some of the generic problem in service robotics.

2 Project Goals

As noted above, the goal of this research is to contribute to the development of a new generation of low-cost service robots that directly interact with people. In particular, our work pursues four central goals

1. **Save and rapid navigation** in dynamic and unpredictable environments shared with people. Both robots described here navigate safely in densely crowded public places.
2. **Effective human-robot interaction** with individuals, crowds of people and people who had no prior exposure to robotics. Minerva, in particular, possesses a collection of mechanisms targeted to appeal at people's intuition.
3. **Autonomous operation in unmodified environments.** Both robots differ from most closely related work in that they did not require modifications of their environments.
4. **Robotic tele-presence** for people at remote locations using the Internet. Both robots possess Web interfaces that enabled people all around the world to command them and perceive sensor data (e.g., camera images).

While the emphasis in the Rhino control system is on navigation, the Minerva system is designed to improve the capabilities of human-robot interaction and tele-presence. Minerva's navigation system is an improved version of Rhino's navigation components. Additionally, Minerva possesses a face to express moods, is able to learn how to best attract people, and possesses a much improved Web interface. As a result, Minerva was also much more effective in interacting with people. Furthermore, Minerva employs a collection of learning algorithms (e.g., for learning maps, or model for composing tours) that facilitate the installation and enable the robot to adapt continuously. These learning algorithms significantly accelerate the installation of mobile robots. Whereas Rhino's installation in the Deutsches Museum Bonn took several weeks, Minerva's installation took a few days.

3 Rhino: Focus on Reliable Navigation in Dynamic Environments

The control system of the first museum tour-guide robot Rhino mainly focuses on safe and reliable navigation in public and populated environments. Public environments, such as museums, differ from more confined environments (e.g., research labs) in that they are highly dynamic, unpredictable, and often even hostile: People often seek to confuse robots. Additionally, there generally are several unmarked hazards such as staircases, or objects that cannot be

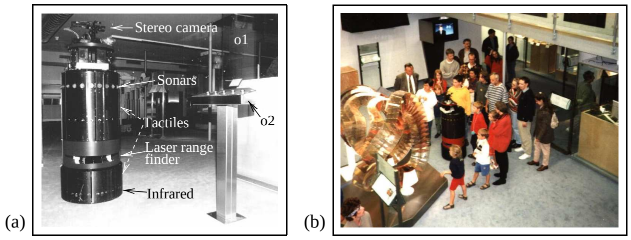

Fig. 1: (a) Rhino, (b) Tour in the Deutsches Museum Bonn (Germany).

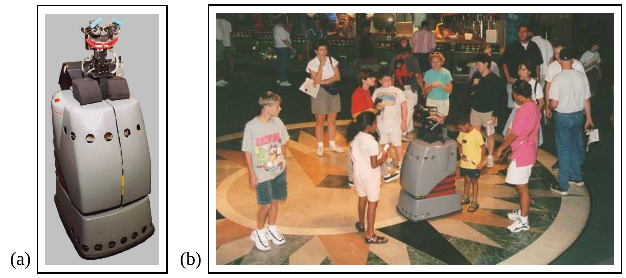

Fig. 2: (a) Minerva, (b) Tour in the Smithsonian's National Museum of American History (Washington, DC).

sensed, making safe navigation a challenging problem of utmost importance.

Rhino's navigation software is described in depth in [2]; therefore, we only summarize its main components here, referring the interested reader to the article above.

3.1 Localization

Localization refers to the determination of the robot's pose (x-y location and bearing θ) within its environment. Accurate localization was a prerequisite for a collection of functions: navigating to exhibits and taking images thereof for the Internet, avoiding collisions with hazards such as staircases that were otherwise undetectable with the robots' sensors, and finding people (as described below).

Rhino employs a grid-based version of *Markov localization* [16], which typically localizes the robot with 10cm accuracy. The (hand-drawn) metric map used by Rhino for localization in the Deutsches Museum Bonn is shown in Fig. 3 (black polygons only). In essence, Markov localization maintains a probability density over robot poses, which is updated whenever the robot moves (as measured by odometry) or when it senses (e.g., using a laser range finder). Fig. 3 illustrates Markov localization during *global localization* (i.e., localization from scratch): After a single sensor scan, the robot's belief is distributed according to the gray region in Fig. 3a (the darker a location, the more likely it is). Fig. 3b depicts the robot's belief after a second sensor scan. Here the probability mass is centered at the correct location, illustrating that two sensor scans are sometimes sufficient to globally localize the robot. A derivation and detailed description of the algorithm can be found in [4].

3.2 People Detection

In Rhino and Minerva alike, people detection serves a dual goal. On the one hand, it is necessary for the robots' interactive components. Static objects that block a robot's path should be treated differently from people. On the other hand, as shown in [5] it aids localization in crowded environments: By filtering out sensor data corrupted by people, the remaining sensor measurements are much better suited for localization, and people cannot confuse the robot by systematically blocking its sensors.

Rhino and Minerva detect people by applying dedicated filter on their proximity measurements [5]. Both robots use the current pose distribution (localization result) to check the plausibility of sensor scans. Sensor scans that "surprise" are attributed to people. Fig. 4a shows those measurements—taken form a single 360° scan, which are believed to correspond to obstacles. Those believed to correspond to people (in the same sensor scan) are shown in Fig. 4b. The filter worked extremely reliable, as detailed experimental results in [5] document.

3.3 Collision Avoidance

The collision avoidance module controls the momentary motion direction and velocity of the robot so as to avoid collisions with obstacles—people and exhibits alike. To determine the location of nearby objects, the robot screens its sensors in regular time intervals (typically 3-4 times a second). Additionally, the map is consulted to generate "fake" (or: virtual) sensor readings that correspond to undetectable hazards or obstacles (e.g., staircases).

At velocities of 70 cm/sec (maximum velocity during the day) to 163 cm/sec (Minerva's high speed under exclusive Web control), inertia and torque limits impose severe constraints on robot motion which may *not* be neglected. Thus, the collision avoidance module takes the dynamics into account. Under the typical dynamic constraints (limited torque), it generates collision-free motion that maximizes two criteria: progress towards the goal, and the robot's velocity. As a result, the robot can ship around smaller obstacles (e.g., a person) without decelerating. The velocity is updated four times a second [2].

3.4 Mapping

The map of the environment used by Rhino in the museum is shown in Fig. 3 (black polygons only). During operation, Rhino updates this map on-the-fly to accommodate changes (e.g., stools that were dropped in narrowly confined areas of the museum). Here Rhino employs occupancy grids [3]. Fig. 5b shows an example, where a large crowd of people blocks a path around an obstacle. The figure shows the modified map, along with the new path generated by the path planner described below.

Minerva goes a step further: The map is learned from data (images, range data, odometry) collected while manually steering the robot through the museum. The problem of building maps from data with imprecise odometry is known as the *concurrent mapping and localization problem* [11],

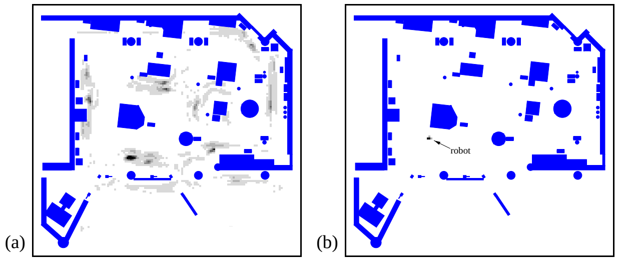

Fig. 3: Global localization in the Deutsches Museum Bonn: (a) Belief state after incorporating one laser scan. (b) Belief after incorporating a second scan, at which point the robot uniquely knows its position.

which highlights its chicken-and-egg nature. Fig. 6 shows two maps, one of the NMAH's floor plan, and one of the museum's ceiling. The first is constructed using laser range data, and the second using a B/W camera pointed upward. Space constraints prohibit us from describing in detail our approach to concurrent mapping and localization (see [19]).

3.5 Planning

Finally, each robot employs two planners, a motion planner for moving from one exhibit to another, and a mission planner for scheduling tours and battery changes (Minerva only).

The motion planner is a modified version of dynamic programming [2], an algorithm similar to the popular A*. It recursively computes the distance to a goal location from all other locations, as depicted by the grayly-shaded area in Fig. 5 (white symbolizes the goal location). This distance function has the advantage that motion commands can be computed for *arbitrary* locations, not just the current one; hence, if people step in the robot's way, the robot does not have to do further computing to recover from a detour. Additionally, our approach employs an efficient re-planning algorithm to adapt plans to the ever-changing map [2]. Minerva's motion planner additionally includes a strategy for staying close to obstacles, in order to minimize the danger of getting lost (a so-called *coastal planner* [12]).

The Rhino system uses GOLOG/GOLEX [6] for mission planning and execution, which essentially executes pre-programmed plans (with some modifications). Minerva's planner is based on RPL [1]. Using a *learned* a model of expected travel times it composes tours dynamically, to meet a target duration of 6 minutes per tour.

4 Minerva: Improved Human Robot Interaction

An analysis of the Rhino project clearly yielded the importance of interaction with people. In contrast to many other forms of human-robot interaction (e.g., gestures [8, 9]), the type interaction in public places like museums is *short-term* and *spontaneous*. Characteristic for the museum domain, as for many other applications of robots in public places, is the fact that people cannot be expected to "know" how to operate a robot. They typically spend very little time with the

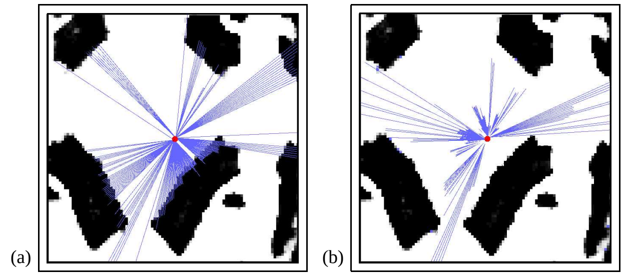

Fig. 4: The novelty filter sorts laser measurements into two categories: (a) obstacles and (b) people.

machine, they often approach robots in crowds, and there are no limitations on people's age and/or technical expertise. To accommodate this diverse mode of interaction, both robots are equipped with an interface that somewhat resembles people. While both robots are equipped with two cameras mounted on a 2 DOF pan/tilt unit and with a device for replaying pre-recorded messages, Minerva possesses a motorized face shown in Fig. 7a, which served as a *focal point* for interaction with people. Through four motors that controlled the mouth and the eyebrows, Minerva is able to exhibit a range of facial expressions, from happy to angry [13].

4.1 "Emotional" States

Minerva uses a finite state automaton to communicate intentions to people, where state emulated "emotions." The FSA is shown in Fig. 7b. It is extremely simple. The default mode is "happy," which is expressed through the appropriate face configuration. When people block the robot's path while giving a tour, it goes through progressively less happy moods, by changing its facial expression and replaying the appropriate message. In the worst case, Minerva reaches the state "angry" and yells at the people phrases like "get out of my way.". The emotional state is reset to "happy" when the robot manages to move 10 centimeters or more. During the deployment of Minerva we found the people clearly recognized Minerva's "emotional" state, acknowledged and respected it in most cases, and gave way to the robot.

Rhino, in contrast, does not possess a face to express "emotional" states. Instead, it uses its horn indicate its desire for space. Unfortunately, most people regard the horn rather entertaining, and intentionally step into the robots way. As a result, Rhino was significantly less successful in negotiating crowds.

4.2 Learning To Attract People

Minerva uses learning to attract people. The data were acquired in regularly scheduled phases in between tours lasting one minute each, during which the robot actively tried to attract people. Instead of telling it how to do this, feedback-driven learning was employed for determining the best strategy. Minerva uses a memory-based reinforcement learning approach [17] (*no* delayed award). Reinforcement is received in proportion of the proximity of people; coming too close, however, leads to a penalty (violating Min-

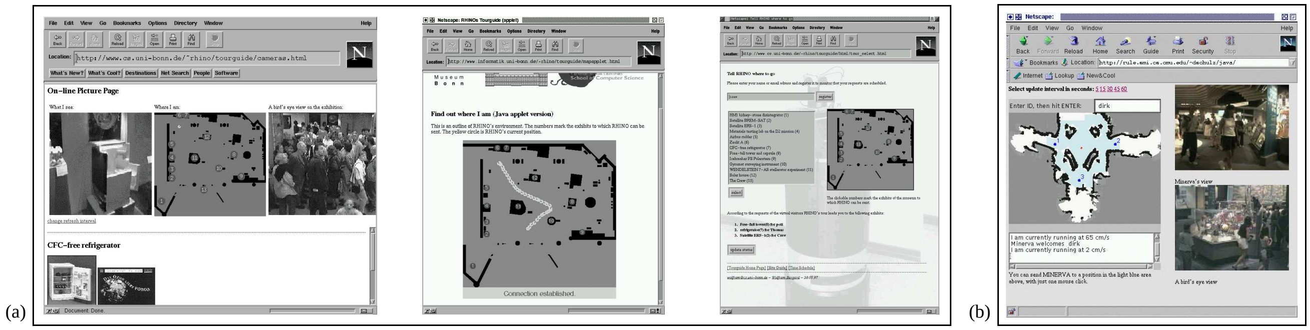

Fig. 8: Web interfaces: (a) Rhino, (b) Minerva.

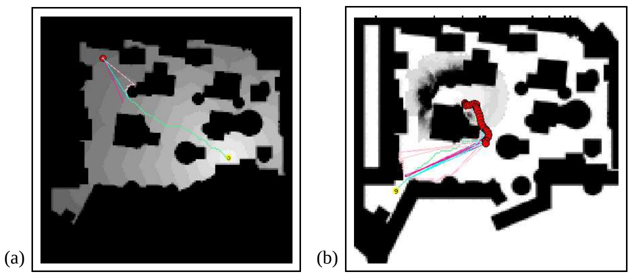

Fig. 5: (a) Path planning with dynamic programming (b) On-line mapping.

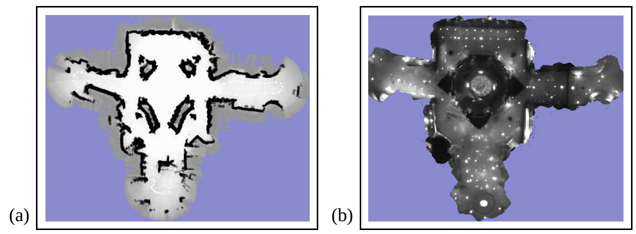

Fig. 6: Maps acquired in the NMAH: (a) Floor plan (b) Ceiling.

erva's space). Minerva's behavior is conditioned on the current density of people. Possible actions included different strategies for head motion (e.g., looking at the nearest person), different facial expressions (e.g., happy, sad, angry), and different speech acts (e.g., "Come over," "do you like robots?").

During the two weeks, Minerva performed 201 attraction interaction experiments, each of which lasted one minute. Over time, Minerva developed a "positive" attitude (saying friendly things, looking at people, smiling). As shown in Fig. 9, acts best associated with a "positive" attitude attracted people the most. For example, when grouping speech acts and facial expressions into two categories, friendly and unfriendly, we found that the former type of interaction performed significantly better than the first (with 95% confidence). However, people's response was highly stochastic and the amount of data we were able to collect during the exhibition is insufficient to yield statistical significance in most cases; hence, we are unable to comment on the effectiveness of individual actions. More findings are discussed in [13].

4.3 Web Interfaces

Finally, both robots possess Web interfaces discussed in detail in [14]. Here, too, Minerva's interface is a progression resulting from insights made with Rhino's interface. Rhino's control interface is spread over three different pages shown in Fig. 8a, one for watching the museum, one for monitoring the robot only, and one for control. While the former pages display the robot's location in the map, along with camera images recorded by the robot and a stationary camera, the latter enabled people to send the robot to 13 pre-defined locations. More than 2,000 people around the world controlled the robot successfully. At times, more than a hundred requests were pending in Rhino's queue.

Rhino's interface design suffers from three limitations: The involvement of more than page made it difficult to simultaneously control the robot and watch its operation; The definition of pre-selected target locations limits the users' choice; The interface does not allow for shared control between museum visitors and the Web.

Minerva's interface is designed to overcome these problems. Minerva possesses a day-time interface (not shown here) and one for exclusive control, shown in Fig. 8b. In both interfaces users control the robot through a single page. In the exclusive control interface users can select arbitrary target locations, which are assigned on a first-come-first-serve basis. Web users have to register with a name, and the robot announces every request which is fulfilled. Anecdotal evidence suggests [14] that this interface is significantly more successful towards our goal of providing people on the Web with a "robotic tele-presence;" however, bandwidth limitations, latency, and the lack of video/audio transmission from the Web user to the robot still make it hard to turn this concept into reality. Consequently, the issue of using robots and the Internet to realize tele-presence remains largely an issue for future research.

5 Experimental Results and Comparison

Both robots successfully led thousands of people through the respective museums, explaining exhibits and interacting

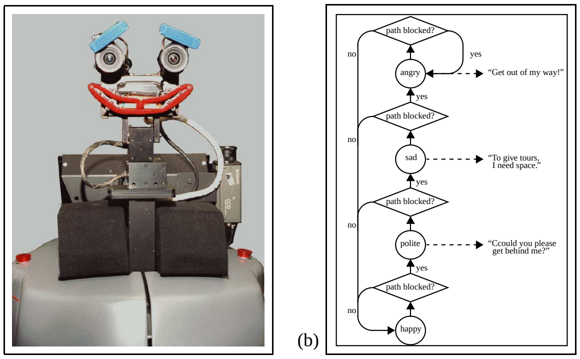

(a) (b)

Fig. 7: (a) Minerva's motorized face and (b) "emotional" states.

with them along the way. As Tab. 1 suggests, both robots on average traversed more than 3km per day, at a comparable average speed. During opening hours, the maximum speed was limited to 70 cm/sec, which is similar to people's walking speed in a museum. At several Internet-only nights, Rhino's maximum speed was 80cm/sec, whereas Minerva moved at a speed of up to 163 cm/sec. The difference is accounted by different hardware limitations. Rhino's presence caused an estimated 50% increase in number of visitors in this—relatively small—museum. No similar estimates exist for Minerva, as the Smithsonian is a huge museum with large, natural fluctuations.

Minerva's environment was an order of magnitude larger and more crowded than Rhino's. The center area of the NMAH was a large open area, which made localization more difficult than in Rhino's environment. This was accommodated by improved navigation algorithms (e.g., by using a camera pointed towards the ceiling and by coastal navigation). Both robots navigated safely and reliably, often at their physical speed limit.

Minerva's interactive strategy is more elaborate than Rhino's. Minerva's looks (in particular its face) and behavior appealed stronger to people's intuition than Rhino. When comparing both robots, we found that people understood Minerva's actions and intentions much better than that of Rhino, and they were typically more satisfied. For example, Minerva was highly successful in asking people for space when giving tours—when its face turned to a sad or angry expression, people typically cleared the way. We found that both robots maintained about the same average speed (Minerva: 38.8 cm/sec, Rhino: 36.6 cm/sec), despite the fact that Minerva's environment was an order of magnitude more crowded. These numbers illustrate the effectiveness of Minerva's interactive approach to making progress.

Minerva also possessed an improved Web interface, which enabled Web users to specify arbitrary target locations instead of choosing locations from a small pool of pre-specified locations. Rhino's Web interface prescribed a small set of 13 possible target locations, which corresponded to designated target exhibits. When under exclusive Web control, Minerva was more than twice as fast as Rhino (see Table 1). In everyday operation, however, the maximum speed of both robots was limited to walking speed (70 cm/sec).

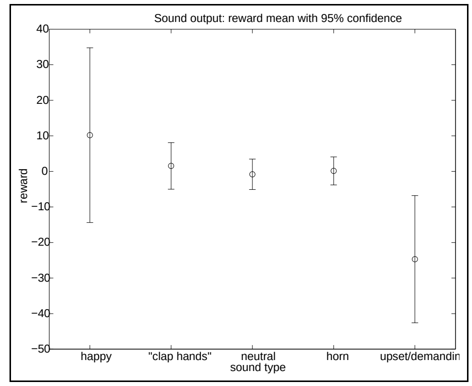

Fig. 9: Statistics of people's response to different styles of interaction (from friendly on the left to upset/demanding on the right). The data were input to a reinforcement learning algorithm, which learned on-line optimal interaction patterns.

6 Related Work

A rich body of related work is reviewed in [2] (over 160 references) and [18]. Probably the first mobile robotic tour-guide was developed by Horswill [7], whose robot Polly gave tours to visitors of MIT's AI Lab. This robot was mostly reactive, strictly relying on visual cues to find people, locations, and obstacles. Tape on the floor was necessary to limit the robot's operational range. More recently (a year after Rhino was deployed), Nourbakhsh and colleagues developed a similar tour-guide named Sage or Chips [10]. Sage/Chips was directly inspired by Rhino. It differs from the robots described here in several aspects: It does not plan; Therefore, it follows a pre-planned path through the museum; It uses large multi-colored markers in the museum for localization; Its interaction is limited to acknowledging when people block its way and replaying pre-recorded multi-media explanations that explain the various exhibits. On the positive side, Sage/Chips can plug itself into a wall-mounted charger and hence operation entirely without human intervention. Neither of these robots (Polly and Sage/Chips) has a Web interface. However, prior to Rhino, the robot Xavier [15] could be controlled through the Web.

7 Discussion and Open Problems

This article surveyed the major components of two robotic museum tour-guides, pointing out the progression from a first generation to a second. Probabilistic navigation algorithm, paired with algorithms for user interaction and learning, led to a highly robust system that has been proven to withstand the challenges that arise in densely populated

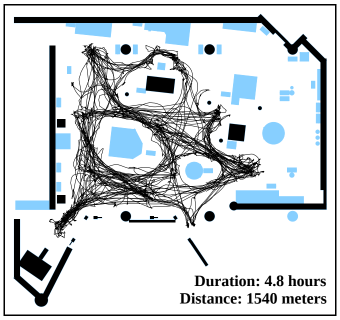

Fig. 10: Example of Rhino's trajectory for a 4.8 hour sequence, during which people corrupt more than 50% of the robot's sensor readings.

	Rhino	Minerva
deployment date	May 1997	August 1998
deployment duration	6 days	13 days
hours of operation	47	94.5
number of visitors	>2,000	> 10,000
number of Web visitors	2,060	n/a
total distance	18.6 km	44.0 km
maximum speed	80 cm/sec	163 cm/sec
average speed during motion	36.6 cm/sec	38.8 cm/sec
exhibits explained	2,400	2,668
additional visitors	>50%	n/a

Table 1: Summary of the robot's six-day deployment period.

public areas.

We believe that these robots are initial examples of a much richer, highly profitable application domain for service robots: entertainment and education. In contrast to a large number of existing service robot applications (e.g., janitorial robots), such robots do not stand in direct competition to human labor; instead, they add to the educational experience of people, and entertain them. As our initial experience suggests, the concept is highly viable for museums, which face an increasing struggle to attract people and deliver their educational messages. We also believe that similar robots can be deployed profitably in shopping malls, entertainment parks, hotels, trade shows, and so on.

The work reported here suggests a rich agenda for future research. While the issue of navigation is relatively well-understood, the issue of spontaneous short-term interaction remains largely open—despite the fact that interaction is a key ingredient of any successful service application. The integration of speech recognition and a natural language interface seems highly attractive for further developing this concept. Additionally, using the Internet to establish a truly bi-directional interaction (e.g., by transmitting acoustic and visual data in *both* directions) is an issue that warrants future research.

Acknowledgments

This research is sponsored in part by DARPA via AFMSC (contract number F04701-97-C-0022), TACOM (contract number DAAE07-98-C-L032), and Rome Labs (contract number F30602-98-2-0137). Additional financial support was received from Daimler Benz Research and Andy Rubin, all of which is gratefully acknowledged.

References

[1] M. Beetz and M. Bennewitz. Planning, scheduling, and plan execution for autonomous robot office couriers. In *Integrating Planning, Scheduling and Execution in Dynamic and Uncertain Environments*, AAAI Workshop Notes 98-02, 1998.

[2] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*. To appear.

[3] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):249–265, 1987.

[4] D. Fox, W. Burgard, and S. Thrun. Grid-based Markov localization for mobile robots. *Journal of Artificial Intelligence*. To appear.

[5] D. Fox, W. Burgard, S. Thrun, and A.B. Cremers. Position estimation for mobile robots in dynamic environments. In *Proceedings of AAAI-98*.

[6] D. Haehnel, W. Burgard, and G. Lakemeyer. GOLEX: Bridging the gap between logic (GOLOG) and a real robot. In *Proceedings of KI-98*.

[7] I. Horswill. Specialization of perceptual processes. Tech.Rep. AI TR-1511, MIT, AI Lab, 1994.

[8] R.E. Kahn, M.J. Swain, P.N. Prokopowicz, and R.J. Firby. Gesture recognition using the PERSEUS architecture. In *Proceedings of CVPR-96*.

[9] D. Kortenkamp, E. Huber, and P. Bonasso. Recognizing and interpreting gestures on a mobile robot. In *Proceedings of AAAI-96*.

[10] I.R. Nourbakhsh. The failures of a self-reliant tour robot with no planner. At `http://www.cs.cmu.edu/~illah/SAGE/index.html`, 1998.

[11] W.D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proceedings of IROS-93*.

[12] N. Roy, W. Burgard, D. Fox, and S. Thrun. Coastal navigation: Robot navigation under uncertainty in dynamic environments. In *Proceedings of ICRA-99*.

[13] J. Schulte, C. Rosenberg, and S. Thrun. Spontaneous short-term interaction with mobile robots in public places. In *Proceedings of ICRA-99*.

[14] D. Schulz, W. Burgard, A.B. Cremers, D. Fox, and S. Thrun. Web interfaces for mobile robots in public places. Submitted for publication.

[15] R. Simmons, R. Goodwin, K. Haigh, S. Koenig, and J. O'Sullivan. A layered architecture for office delivery robots. In *Proceedings of ICAA-97*.

[16] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of IJCAI-95*.

[17] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

[18] S. Thrun, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. MINERVA: A second generation mobile tour-guide robot. In *Proceedings of ICRA-99*.

[19] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31:29–53, 1998.