A Robotic Wayfinding System for the Visually Impaired

Vladimir Kulyukin* **Chaitanya Gharpure** **Pradnya Sute**

Nathan De Graw **John Nicholson**

Department of Computer Science
Utah State University
Logan, UT 84322
vladimir.kulyukin@usu.edu

Abstract

We present an emerging indoor assisted navigation system for the visually impaired. The core of the system is a mobile robotic base with a sensor suite mounted on it. The sensor suite consists of an RFID reader and a laser range finder. Small passive RFID sensors are manually inserted in the environment. We describe how the system was deployed in two indoor environments and evaluated by visually impaired participants in a series of pilot experiments.

Introduction

The inability to navigate unfamiliar environments remains a key barrier to equal access for the 11.4 million visually impaired people in the United States (LapLante & Carlson 2000). This inability denies the visually impaired adequate access to many buildings, impedes their use of public transit, and makes their integration into local communities difficult. Thus, there is a significant need for assisted navigation systems that help the visually impaired overcome the navigation barrier, especially in unfamiliar environments, where conventional aids, such as white canes and guide dogs, are of limited use.

Over the past three decades, considerable R&D effort has been dedicated to navigation devices for the visually impaired. Benjamin, Ali, and Schepis (Benjamin, Ali, & Schepis 1973) built the C-5 Laser Cane. The cane uses optical triangulation with three laser diodes and three photo-diodes as receivers. Bissit and Heyes (Bissit & Heyes 1980) developed the Nottingham Obstacle Detector (NOD), a hand-held sonar device that gives the user auditory feedback with eight discrete levels. Shoval et al. developed the NavBelt, an obstacle avoidance wearable device equipped with ultrasonic sensors and a wearable computer (Shoval, Borenstein, & Koren 1994). The NavBelt produces a 120-degree wide view ahead of the user. The view is translated into stereophonic audio directions. Borenstein and Ulrich (Borenstein & Ulrich 1994) built GuideCane, a mobile obstacle avoidance device for the visually impaired. GuideCane consists of a long handle and a ring of ultrasonic sensors mounted on a steerable two-wheel axle.

More recently, a radio frequency identification (RFID) navigation system for indoor environments was developed at the Atlanta VA Rehabilitation Research and Development Center (Ross 2001; Ross & Blasch 2002). In this system, the blind users' canes are equipped with RFID receivers, while RFID transmitters are placed at hallway intersections. As the users pass through transmitters, they hear over their headsets commands like "turn left," "turn right," and "go straight." The Haptica Corporation has developed Guido©, a robotic walking frame for people with impaired vision and reduced mobility (www.haptica.com). Guido© uses the onboard sonars to scan the immediate environment for obstacles and communicates detected obstacles to the user via speech synthesis.

While the existing approaches to assisted navigation have shown promise, they have had limited success for the following reasons. First, many existing systems increase the user's navigation-related physical load because they require that the user wear additional and, oftentimes substantial, body gear (Shoval, Borenstein, & Koren 1994), which contributes to physical fatigue. The solutions that attempt to minimize body gear, e.g., the C-5 Laser Cane (Benjamin, Ali, & Schepis 1973) and the GuideCane (Borenstein & Ulrich 1994), require that the user abandon her conventional navigation aid, e.g., a white cane or a guide dog, which is not acceptable to many visually impaired individuals. Second, the user's navigation-related cognitive load remains high, because the user makes all final wayfinding decisions. While device-assisted navigation enables visually impaired individuals to avoid immediate obstacles and gives them simple directional hints, it provides little improvement in wayfinding over white canes and guide dogs. Limited communication capabilities also contribute to the high cognitive load.

*This research has been supported, in part, by the National Science Foundation Grant IIS-0346880, a Community University Research Initiative Grant (CURI) from the State of Utah, and by a New Faculty Research grant from Utah State University.

Finally, few assisted navigation technologies are deployed and evaluated in their target environments over extended time periods. It is this lack of deployment and evaluation that makes it difficult for assistive technology (AT) practitioners to compare different solutions and choose the one that best fits the needs of a specific individual.

Can Robot-Assisted Navigation Help?

Yes, it can. First, the amount of body gear carried by the user is significantly minimized, because most of it can be mounted on the robot and powered from on-board batteries. Thus, robotic bases offer solutions to two of the hardest problems in assisted navigation: hardware miniaturization and portable power supply. Consequently, the navigation-related physical load is reduced. Second, the delegation of such key wayfinding capabilities as localization and path planning to the robot reduces the user's cognitive load. Third, the robot can interact with other people in the environment, e.g., ask them to yield or receive instructions. Fourth, robotic guides can carry useful payloads, e.g., suitcases and grocery bags. Finally, the user can use robotic guides in conjunction with her conventional navigation aids, e.g., white canes and guide dogs.

Are all environments suitable for robotic guides? No. There is little need for such guides in familiar environments where conventional navigation aids are adequate. However, unfamiliar indoor environments that are dynamic and complex, e.g., airports and conference centers, are a perfect niche for robotic guides. Guide dogs, white canes, and other navigation devices are of limited use in such environments, because they cannot help their users localize and find paths to useful destinations.

The idea of robotic guides is not new. Horswill (Horswill 1993) used the situated activity theory to build Polly, a mobile robot guide for the MIT AI Lab. Polly used lightweight vision routines that depended on textures specific to the lab. Thrun et al. (Thrun et al. 1999) built Minerva, a completely autonomous tour-guide robot that was deployed in the National Museum of American History in Washington, D.C. Burgard et al. (Burgard et al. 1999) developed RHINO, a close sibling of Minerva's, which was deployed as an interactive tour guide in the Deutsches Museum in Bonn, Germany. Unfortunately, these robots do not address the needs of the visually impaired. The robots depend on the users' ability to maintain visual contact with them, which cannot be assumed for the visually impaired. Polly has very limited interaction capabilities: the only way users can interact with the system is by tapping their feet. To request a museum tour from RHINO (Burgard et al. l1999), the user must identify and press a button of a specific color on the robot's panel. The approach on which Polly is based requires that a robot be evolved by its designer to fit its environment not only in terms of software but also in terms of hardware. This makes it difficult to produce replicable solutions that work out of the box in a variety of environments. Autonomous solutions like RHINO and Minerva require substantial investments in customized engineering to become and remain operational. For example, RHINO must run 20 parallel processes on 3 on-board PCs and 2 off-board SUN workstations connected via a customized Ethernet-based point-to-point socket communication protocol. Even with these high software and hardware commitments, RHINO reportedly experienced six collisions over a period of forty-seven hours, although each tour was less than ten minutes long (Burgard et al. 1999).

A Robotic Guide

We have built and deployed a prototype of a robotic guide for the visually impaired. Its name is RG, which stands for "robotic guide." Our basic research objective is to alleviate localization and navigation problems of purely autonomous approaches by incrementing environments with inexpensive and reliable sensors that can be placed in and out of environments without disrupting any indigenous activities. Effectively, the environment becomes a distributed tracking and guidance system (Kulyukin & Blair2003; Kulyukin, Gharpure, & De Graw 2004) that consists of stationary nodes, e.g., sensors and computers, and mobile nodes, e.g., robotic guides.

Additional requirements are: 1) that the instrumentation be fast, e.g., two to three hours, and require only commercial off-the-shelf (COTS) hardware components; 2) that sensors be inexpensive, reliable, easy to maintain (no external power supply), and provide accurate localization; 3) that all computation run onboard the robot; and 4) that human-robot interaction be both reliable and intuitive from the perspective of the visually impaired users. The first two requirements make the systems that satisfy them replicable, maintainable, and robust. The third requirement eliminates the necessity of running substantial off-board computation to keep the robot operational. In emergency situations, e.g., computer security breaches, power failures, and fires, off-board computers are likely to become dysfunctional and paralyze the robot if it depends on them. The fourth requirement explicitly considers the needs of the target population.

Scope limitations

Several important issues are currently beyond the project's scope. First, we do not address the issue of navigating large open spaces, e.g., large foyers in hotels. While a few recent references in the literature suggest that ultrasonic sensors could be used to address this issue (Addlesee et al. 2001), the proposed solutions are sketchy and are evaluated in carefully controlled, small lab environments. Thus, we assume that all environments in which RG operates have walls, hallways, aisles, T- and Xintersections, and solid and static objects, e.g., vending machines, that the onboard sensors can detect. Second, robotic guides prototyped by RG are not meant for individual ownership. Rather, we expect institutions, e.g., airports and large chain stores, to purchase such guides and operate them on the premises in the future. Third, the wayfinding technology proposed in this project could potentially assist sighted people with cognitive and mobile disabilities in navigating unfamiliar environments. While we intend to pursue these possibilities in the future, the target population for this project is visually impaired individuals (no more than light perception), at least 16 years of age, ambulatory, with no serious speech impediments, hearing problems, or cognitive disabilities. Finally, robotic guides prototyped by RG are not meant for outdoor navigation.

Hardware and software

RG is built on top of the Pioneer 2DX robot platform (www.activmedia.com). The platform has three wheels, 16 ultrasonic sonars, 8 in front and 8 in the back, and is equipped with three rechargeable Power Sonic PS-1270 onboard batteries that can operate for up to five hours at a time.

What turns the platform into a robotic guide is a Wayfinding Toolkit (WT) mounted on top of the platform and powered from the on-board batteries. The WT currently resides in a PCV pipe structure attached to the top of the platform. The WT includes a laptop connected to the platform's microcontroller. The communication between the laptop and the microcontroller is done through a usb to serial cable.

The laptop interfaces to a radio-frequency identification (RFID) reader through another usb to serial cable. The TI Series 2000 RFID reader is connected to a square 200mm by 200mm RFID RI-ANT-GO2E antenna that detects RFID sensors (tags) placed in the environment. We currently use TI RFID Slim Disk tags. These tags can be attached to any objects in the environment or worn on clothing. They do not require any external power source or direct line of sight to be detected by the RFID reader. They are activated by the spherical electromagnetic field generated by the RFID antenna with a radius of approximately 1.5 meters. Each tag is programmatically assigned a unique ID. Finally, the laptop is connected to a LMS 200 laser range finder from the SICK corporation (www.sick.com) mounted on the front of the robot. A dog leash is attached to the battery bay handle on the back of the platform. The upper end of the leash is hung on a PCV pole next to the RFID antenna's pole. Visually impaired individuals follow RG by holding onto that leash. Eventually, when we find a satisfactory solution to the hardware miniaturization problem, we may be able to mount the WT directly on the guide dog and take advantage of a natural vision system.

As a software system, RG is based on Kupiers' Spatial Semantic Hierarchy (SSH) (Kupiers 2000). The SSH is a framework for representing spatial knowledge. It divides spatial knowledge of autonomous agents, e.g., people and robots, into four levels: the control level, causal level, topological level, and metric level. The control level consists of low level mobility laws, e.g., trajectory following and aligning with a surface. The causal level represents the world in terms of views and actions. A view is a collection of data items that an agent gathers from its sensors. Actions move agents from views to views. For example, a robot can go from one end of a hallway (start view) to the other end of the hallway (end view). The topological level represents the world's connectivity, i.e., how different locations are connected. The metric level adds distances between locations.

The control level is implemented with the following low-level routines all of which run on the laptop: follow-wall, turn-left, turn-right, avoid-obstacles, go-thru-doorway, pass-doorway, and make-u-turn. In selecting the routines, we tried to find a minimal action set that can be used in many standard indoor environments. These routines are written in the behavior programming language of the ActivMedia Robotics Interface for Applications (ARIA) system from ActivMedia Robotics, Inc. These behaviors draw on well understood indoor navigation techniques of the potential fields approach (Murphy 2000). Thus, in RG, sensor readings are converted into attractive and repulsive vectors that are summed to decide where the robot should go next.

RFID-based localization allows the robot to overcome the problem of local minima that many potential fields approaches experience in the absence of reliable localization, because the robot can generate a necessary direction vector on the basis of its current location. The laptop also runs two other software components: 1) a speech recognition and synthesis engine that enables RG to receive and synthesize speech and 2) a path planner. The advantages and disadvantages of speech-based interaction are discussed in the next section.

The Path Planner realizes the causal and topological levels of the SSH. The Planner's knowledge base represents an aerial view of the environment in which RG operates. Currently, the knowledge base consists of tag connectivity graphs, tag to destination mappings, and low-level action scripts associated with specific tags.

Specifically, the environment is represented as a graph where nodes represent the RFID tags and the edges represent the actions required to travel from one tag to another. Views consist of sonar and laser range finder readings and the IDs of the RFID tags currently detectable. Thus, detecting a tag can trigger a specific behavior, e.g., move through a doorway and align with a right wall. The Planner uses the standard breadth first search algorithm to find a path from one location to the other. A path plan is a sequence of tag numbers and action sequences at each tag.

Human-Robot Interaction

Humans can interact with RG through speech, wearable keyboard, and GUIs. Speech-based and keyboard-based interactions are intended for visually impaired individuals. GUIs are intended for system administrators. Speech is received by RG through a wireless microphone placed on the user's clothing. Speech is recognized and synthesized with Microsoft Speech API (SAPI) 5.1, which is freely available from www.microsoft.com/speech. SAPI includes the Microsoft English SR Engine Version 5, a state-of-the-art Hidden Markov Model speech recognition engine. The engine includes 60,000 English words, which we found adequate for our purposes. SAPI couples the Hidden Markov Model speech recognition with a system for constraining speech inputs with context-free command and control grammars. The grammars constrain speech recognition sufficiently to eliminate user training and provide speaker-independent speech recognition. Grammars are defined with XML Data Type Definitions (DTDs).

RG interacts with its users and people in the environment through speech and audio icons. An audio icon is a non-verbal sound that can be readily associated with a specific object, e.g., the sound of water bubbles associated with a water cooler. For example, when RG is passing a water cooler, it can either say “water cooler” or play an audio file with sounds of water bubbles. We added audio icons to the system because, as recent research findings indicate (Tran, Letowski, & Abouchacra 2000), speech perception can be slow and prone to block ambient sounds from the environment. On the other hand, associating objects and events with non-speech audio messages requires training or the presence of a universally accepted mapping between events and objects and sounds. Since no such mapping is currently available, our assumption is that the user can quickly create such a mapping.

Pilot Experiments

We deployed our system for a total of approximately fifty hours in two indoor environments: the Assistive Technology Laboratory (ATL) of the Utah State University (USU) Center for Persons with Disabilities and the USU Computer Science Department. The ATL occupies part of a floor in a building on the USU North Campus. The floor has an area of approximately 14,000 square feet. The floor contains 6 laboratories, two bathrooms, two staircases, and an elevator. The CS Department occupies an entire floor in a multi-floor building. The floor's area is 21,600 square feet. The floor contains 23 offices, 7 laboratories, a conference room, a student lounge, a tutor room, two elevators, several bathrooms, and two staircases.

Forty RFID tags were deployed at the ATL and one hundred tags were deployed at the CS Department. Once the destinations were known, it took one person 20 minutes to deploy the tags and about 10 minutes to remove them at the ATL. The same measurements at the CS Department were 30 and 20 minutes, respectively. The tags were attached to objects with regular scotch tape. The creation of the knowledge base took one hour at the ATL and about 2 hours at the CS Department. In both environments, one administrator first walked around the area with a laptop and recorded tag-destination associations. Then the administrator associates specific robotic actions with tags. RG was first deployed and tested by the project's team at the ATL, the smaller of the two environments, and then deployed at the CS Department for the pilot experiments with visually impaired participants.

Our pilot experiments involved five visually impaired participants, one participant at a time, over a period of two months. The five participants were selected from the target population: three participants were completely blind and two participants could perceive only light. The participants had no speech impediments, hearing problems, or cognitive disabilities. Two participants were dog users; the other three used white canes. The participants were asked to use RG to navigate to three distinct locations (an office, a lounge, and a bathroom) at the USU Computer Science Department. All participants were new to the environment and had to navigate approximately 40 meters to get to all destinations. In communication and audio perception experiments, we tested the participants' ability to communicate with RG through speech and the participants' audio perception preferences, i.e., whether they preferred to be notified of events and objects in the environment through speech or audio icons. The exit interviews showed that the participants really liked the system and thought it was a very useful navigation aid in unfamiliar environments. They especially liked the idea that they did not have to give up their white canes and guide dogs to use RG.

Navigation

All five participants reached the assigned destinations. However, they expressed concerns about RG's slow speed and movement jerkiness. RG's speed was 0.5 m/s, which is below normal walking speeds of 1.2-1.5 m/s. Since at the time of the pilot study, the laser range finder had not yet been integrated into the WT, RG's navigation was based on sonar sensors, which are sometimes unreliable due to specular reflection and cross talk. In addition, since the effective range of the sonars is 2.5 meters, we had to adjust the speed of the robot so that it could stop in time to avoid bumping into people in hallways. The jerky movements occurred when the sonars either underestimated or overestimated the distances to the closest objects. For example, once when the robot followed a right wall and two of its right sonars told it that the wall was 50 centimeters (cm) further than it actually was, the robot turned right to better align itself with the wall, which caused jerkiness.

Another structural improvement suggested by the visually impaired participants in our pilot study was to replace the dog leash with a more static plastic pole similar to a white cane. This suggestion was made by the white cane users. It is interesting that the two dog users liked the dog leash, but said they would not mind a more static holder.

Speech-based communication

Our first human-robot communication experiment tested the feasibility of using speech as a means of input for humans to communicate with the robot. Each participant was asked to speak approximately sixty phrases while wearing a headset that consisted of a microphone and one headphone. The phrase list was a list of standard phrases that a person may say to a robotic guide in an unfamiliar environment, e.g., "go to the bathroom," "where am I?" etc. Each phrase was encoded as a context-free command and control grammar rule in SAPI's XML-based grammar formalism. Each participant was positioned in front of a computer running SAPI. The test program was written to use SAPI's text-to-speech engine to read the phrases to the participant one by one, wait for the participant to repeat a phrase, and record a recognition result (speech recognized vs. speech not recognized) in a database. This experiment was repeated in two environments: noise-free and noisy. The noise-free environment did not have any ambient sounds other than the usual sounds of a typical office. To simulate a noisy environment, a long audio file of a busy bus station was played on another computer in the office. All five participants were native English speakers and did not train SAPI's speech recognition engine on sample texts.

We found that the average percentage of phrases recognized by the system in the noise-free environment was 38%, while the average percentage of recognized phrases in the noisy environment was 40.2%. While the level of ambient noise in the environment did not seem to affect the system's speech recognition, in both environments fewer than 50% of phrases were correctly recognized. Even worse, on average, 20% of spoken phrases were incorrectly recognized by the system. For example, when one participant made two throat clearing sounds, the system recognized the sound sequence as the phrase "men's room."

The statistics were far better for the participants understanding phrases spoken by the computer. The average percentage of speech understood in the noise-free environment was 83.3%, while the average percentage of phrases understood in the noisy environment was 93.5%. Clearly, in the second trial (the noisy environment), the participants were more used to SAPI's speech synthesis patterns. These results suggest that speech appears to be a better output medium than input.

Another problem with speech recognition occurs when the person guided by RG stops and engages in conversation with someone. Since speech recognition runs continuously, some phrases said by the person are erroneously recognized as route directives, which causes RG to start moving. For example, once RG erroneously recognized a directive and started pulling its user away from his interlocutor until the user's stop command pacified it. In another situation, RG managed to run a few meters away from its user, because the user hung the leash on the PCV pole when he stopped to talk to a friend of his in a hallway. Thus, after saying “Stop,” the user had to grope his way along a wall to RG, standing a few meters away.

As was argued elsewhere (Kulyukin 2004), it is unlikely that these problems can be solved on the software level until there is a substantial improvement in the state-of-the-art speech recognition. Of course, one could add yes-no route change confirmation interactions. However, since unintended speech recognition is frequent, such interactions could become annoying to the user. Therefore, we decided to seek a wearable hardware solution. Specifically, we are exploring human-robot interaction through a wearable keyboard. Many wearable keyboards now fit in the palm of one's hand or can be worn as badges. We are currently experimenting with a small Belkin© keypad that directly interfaces to the WT laptop. When a guided person stops to talk to someone, one button push disables the speech recognition process for the duration of the conversation. Similarly, when the guided person clears her throat and RG misinterprets it as a command, one button push can tell RG to ignore the command and stay on the route. Potentially, the wearable keyboard may replace speech recognition altogether. The obvious advantage is that keyboard-based interaction eliminates the input ambiguity problems of speech recognition. One potential disadvantage is the learning curve required of a human participant to master the necessary key combinations. Another potential disadvantage is restrictions on the quality and quantity of interactions due to the small number of keys. Additional experiments with human participants are needed to determine the validity of these speculations.

Audio perception

We conducted audio perception experiments with all five participants to test whether they preferred speech to audio icons, e.g., a sound of water bubbles, to signify different objects and events in the environment and how well participants remembered their audio icon selections. The participants used the tool to associate events and objects, e.g., water cooler to the right, approaching left turn, etc., with three audio messages: one speech message and two audio icons. There were seven different objects, e.g., elevator, vending machine, bathroom, office, water cooler, left turn, and right turn. This small number was chosen to eliminate steep learning curves.

Each object was associated with two different events: *at* and *approaching*. For example, one can be at the elevator or approaching the elevator. The audio icons available for each event were played to each participant at selection time. The following statistics were gathered: 1) percentage of accurately recognized icons; 2) percentage of objects/events associated with speech; 3) percentage of objects/events associated with audio icons; 4) percentage of objects/events associated with both. The averages for these experiments were: 1) accurately recognized icons (93.3%); 2) objects/events associated with speech (55.8%); 3) objects/events associated with icons (32.6%); 4) objects/events associated with both (11.4%). The analysis of the audio perception experiments showed that two participants were choosing audio preferences essentially at random, while the other three tended to follow a pattern: they chose speech messages for at events and audio icons for approaching events or vice versa. The experiments also showed that the participants tended to go either with speech or with audio icons, but rarely with both. The experiments did not give a clear answer as to whether visually impaired individuals prefer to be notified of objects/events via speech or audio icons. Further work is needed on a larger sample to answer this question on a statistically significant level.

Conclusion

We presented an assisted navigation system for the visually impaired. The system consists of a mobile robotic guide and small RFID sensors embedded in the environment. The system allows visually impaired individuals to navigate in unfamiliar indoor environments and interact with the robotic guide via speech, sound, and wearable keyboard.

The main question addressed by our research is the feasibility of robot-assisted navigation in indoor environments instrumented with inexpensive passive sensors. So, is indoor robot-assisted navigation feasible? While our experiments show that the technology has promise, the answer to this question cannot be given either negatively or affirmatively at this point. The benefits of autonomous systems, such as RG, increase with longer deployments. Only through longer deployments can one answer how easy it is to maintain the system over extended time periods and whether the target environment accepts the technology sociologically. Our deployments have not been sufficiently long to answer either question.

A major barrier to long-term deployment is the unwillingness of real target environments, e.g., airports, to agree to the exploratory long-term deployments (two to three months) of assistive technologies. Only after this barrier is overcome, both sociologically and technologically, will we be able to render a definite verdict on the feasibility of indoor robot-assisted navigation.

References

Addlesee, M.; Curwen, R.; Hodges S.; Newman J.; Steggles P.; and Ward A. 2001. Implementing a Sentient Computing System. *IEEE Computer*, August: 2-8.

Benjamin, J. M.; Ali, N. A.; and Schepis, A. F. 1973. A Laser Cane for the Blind. In *Proceedings of San Diego Medical Symposium*, San Diego, CA.

Bissit, D. and Heyes, A. 1980. An Application of Biofeedback in the Rehabilitation of the Blind. *Applied Ergonomics*, 11(1):31-33.

Borenstein, J. and Ulrich, I. 1994. The GuideCane – A Computerized Travel Guide for the Active Guidance of Blind Pedestrians. In *Proceedings of IEEE International Conference on Robotics and Automation.* San Diego, CA.

Burgard, W.; Cremers, A.; Fox, D.; Hahnel, D; Lakemeyer G.; Schulz, D.; Steiner, W.; and Thrun, S. 1999. Experiences with an Interactive Museum Tour-Guide Robot. *Artificial Intelligence* (114):3-55.

Horswill, I. 1993. Polly: A Vision-Based Artificial Agent. In *Proceedings of the 11th Conference of the American Association for Artificial Intelligence (AAAI-93).* Washington, DC.

Kulyukin, V. and Blair, M. 2003. Distributed Tracking and Guidance in Indoor Environments. In *Proceedings of the Conference of Rehabilitation Engineering and Assistive Technology Society of North America (RESNA-2003).* Atlanta, GA.

Kulyukin, V.; Gharpure, C.; and De Graw, N. 2004. A Robotic Guide for the Visually Impaired. In *Proceedings of the AAAI Spring Symposium on Interaction between Humans and Autonomous Systems over Extended Operation*. Palo Alto, CA.

Kulyukin, V. 2004. Human-Robot Interaction through Gesture-Free Spoken Dialogue. *Autonomous Robots*, 16(3), to appear.

Kupiers, B. 2000. The Spatial Semantic Hierarchy. Artificial Intelligence, 119:191-233.

LaPlante, M. and Carlson, D. 2000. *Disability in the United States: Prevalence and Causes.* U.S. Department of Education, National Institute of Disability and Rehabilitation Research.

Murphy, R. 2000. *Introduction to AI Robotics*. Cambridge, MA: MIT Press.

Ross, D. and Blasch, B. 2002. Development of a Wearable Computer Orientation System. *IEEE Personal and Ubiquitous Computing*, 6:49-63.

Ross, D. 2001. Implementing Assistive Technology on Wearable Computers. IEEE Intelligent Systems, May:2-8.

Shoval, S.; Borenstein, J.; and Koren, Y. 1994. Mobile Robot Obstacle Avoidance in a Computerized Travel for the Blind. In *Proceedings of the IEEE International Conference on Robotics and Automation*. San Diego, CA.

Thrun, S.; Bennewitz, M.; Burgard, W.; Cremers, A.; Dellaert, F.; Fox, D.; Hahnel, D.; Rosenberg, C.; Roby, N.; Schutle, J.; and Schultz, D. 1999. Minerva: A Second Generation Mobile Tour-Guide Robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA-99).* Antwerp, Belgium.

Tran, T.; Letowski, T.; and Abouchacra, K. 2000. Evaluation of Acoustic Beacon Characteristics for Navigation Tasks. *Ergonomics*, 43(6):807-827.