MAGELLAN: An Integrated Adaptive Architecture for Mobile Robotics

BRIAN YAMAUCHI[1,2] (yamauchi@aic.nrl.navy.mil)
PAT LANGLEY[1] (langley@newatlantis.isle.org)
ALAN C. SCHULTZ[2] (schultz@aic.nrl.navy.mil)
JOHN GREFENSTETTE[2] (gref@aic.nrl.navy.mil)
WILLIAM ADAMS[2] (adams@aic.nrl.navy.mil)

[1]Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, California 94306

[2]Navy Center for Applied Research in Artificial Intelligence
Naval Research Laboratory, Washington, DC 20375-5337

Abstract

In this paper we describe MAGELLAN, an integrated architecture for mobile robotics. The system represents its spatial knowledge in terms of a topological network that connects a set of distinct places, each represented by evidence grids that contain probabilistic descriptions of occupancy. MAGELLAN includes a module for place recognition that determines its initial location and when it has reached a goal, a module for continuous localization that maintains accurate estimates of the robot's position, and a module for navigation that generates path plans and executes them using reactive behaviors. Experiments in two laboratories with different characteristics suggest that the system can operate robustly across a range of environments, including ones that involve dynamic changes.

This work was funded by the Office of Naval Research under Grant No. N00014-94-1-0505. We thank Hans Moravec for sharing his software, and we thank Nils Nilsson and Barbara Hayes-Roth for providing access to their Nomad mobile robot.

1. Introduction

Mobile robots require many different capabilities in order to perform useful tasks in real-world domains. Two of these capabilities are localization, determining the robot's own location, and adaptive navigation, the ability to move robustly from one location to another. Localization can be further divided into the tasks of place recognition and continuous localization. Place recognition is like waking up in a hotel room, looking out the window, and trying to determine what city one is in; in contrast, continuous localization is like trying to avoid getting lost while driving downtown. For a mobile robot, place recognition is necessary whenever the robot is turned on in an unknown initial location, whereas continuous localization is needed as the robot moves through the world. Adaptive navigation refers to methods for moving around the world that are robust to dynamic changes, such as moving people or rearranged obstacles.

In previous research, we have developed methods for place recognition (Langley & Pfleger, 1995; Langley, Pfleger, & Sahami, 1997; Yamauchi & Langley, 1997), continuous localization (Schultz, Adams, & Grefenstette, 1996; Graves, Schultz, & Adams, 1997), and adaptive navigation (Yamauchi & Beer, 1996). In this paper we report on the Magellan project, in which the primary goal was to integrate these different techniques into a complete system that, given a map of its environment, can localize itself from an initially unknown position and then navigate through the world while maintaining an accurate position estimate.

We start with an overview of the Magellan system architecture and its representation of spatial knowledge. After this, we describe place recognition, continuous localization, and adaptive navigation in detail. Next we present results from experiments in which we tested Magellan in two different indoor environments. Finally, we survey related work, present our conclusions, and consider directions for future research.

2. The Magellan System

Magellan is an integrated architecture for mobile robotics. In this section, we describe the robot hardware on which it operates and present an overview of the architecture's components. We then examine to the underlying representation of spatial knowledge that these modules use in their processing. After laying this groundwork, we describe each of the architecture's three components in some detail.

2.1 Robot Hardware and System Architecture

We have implemented Magellan on a number of Nomad 200 mobile robots, shown in Figure 1, at both the Navy Center for Applied Research in Artificial Intelligence and at the Stanford University Robotics Laboratory. The Nomad 200 is a wheeled robot with a zero-turning radius that has 16 sonar range sensors and 16 infrared range sensors spaced evenly around the base. The robot also has a planar laser rangefinder that can collect high-resolution range data in a ten degree arc. All of these sensors are mounted on a turret that can rotate independently of the base.

As noted earlier, the Magellan architecture integrates three capabilities that seem necessary for a complete mobile robot: place recognition, continuous localization, and navigation. Figure 2 shows the high-level architecture for the system, including these major components, their lines of communication, and the type of information they pass to one another.

As we detail in Section 2.3, the place recognition system provides an initial estimate of the robot's position by building a description of the robot's current location and finding the best match among the descriptions associated with places in the world. The identity of the room that contains the current place is then passed to the continuous localization system, along with the robot's inferred Cartesian position.

Figure 1. The Nomad 200 mobile robot.

The continuous localization system retrieves the description associated with the current room and builds new descriptions as the robot moves through the world, registering with the room description to update the robot's position estimate. In this way, the robot avoids the cumulative position error that usually accompanies position estimation based on dead reckoning. Continuous localization directly sets the robot's encoders, so no direct communication between localization and navigation is necessary.

The navigation system guides the robot to its destination. This system carries out heuristic search through a topological network of places to generate a path plan and uses reactive behaviors to execute that plan. The navigation module is adaptive in that successes and failures in execution lead to changes in the topological network. Thus, it can take advantage of stable structures in the environment but also deal with aspects that vary over time.

In the remainder of the section, we describe each of MAGELLAN's component systems in more detail. However, one cannot characterize mechanisms without establishing the structures on which they operate. Thus, we will first consider the system's underlying representation of spatial knowledge.

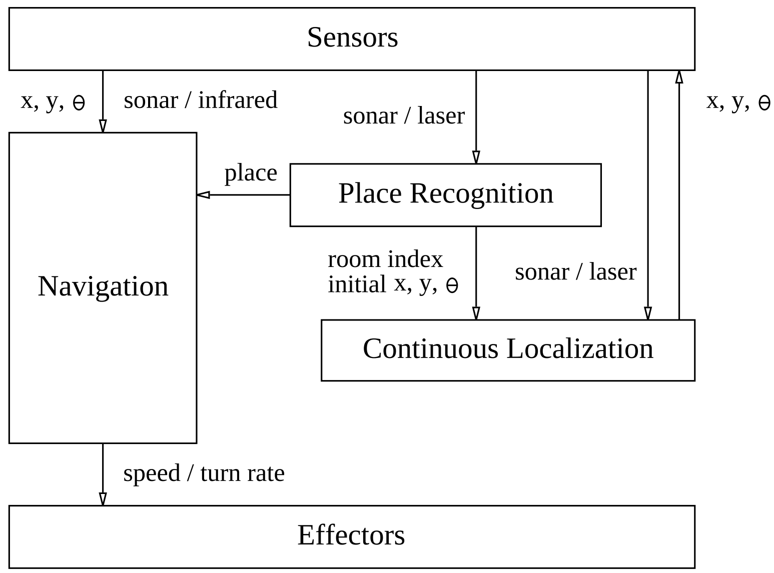

Figure 2. The Magellan system architecture.

2.2 Spatial Representation

Two representational schemes have predominated in work on mobile robots. The first, *topological maps*, represent the general structure of the environment in terms of place nodes and links that indicate adjacency between those nodes. The second, *occupancy grids*, represent the detailed structure of a region by dividing it into a Cartesian grid and indicating whether each cell is occupied. Topological maps have advantages for rapid navigation, whereas occupancy grids provide better spatial information for use in localization.

Magellan incorporates both topological and grid representations. The system uses a topological map called an *adaptive place network* (Yamauchi & Beer, 1996) for navigation. These networks differ from traditional topological maps in that they support changes in their connections to reflect changes encountered in the world. Magellan uses occupancy grids to represent the local structure of each place (for place recognition) and the larger structure of room regions (for continuous localization). By integrating these two representations, the architecture obtains the advantages of both topological maps and occupancy grids.

2.2.1 Evidence Grids

The Magellan system represents its local spatial knowledge in terms of *evidence grids*, a particular type of occupancy grid developed by Moravec and Elfes (1985). These consist of Cartesian grids in which each cell has a certain probability of being occupied. Initially, each cell probability is set to the estimated prior probability of cell occupancy. For example, if one quarter of the space in a given area is occupied, one might set the prior probability to 0.25. In practice, evidence grids tend to be insensitive to variations in the prior probability, and we have found that an estimate of 0.5 generally works well.

Each time the robot receives a sensor input, the evidence grid is updated using the corresponding sensor model, which describes the probability that cells are occupied given the reading received. This model depends on the characteristics of the individual sensor. One of the major advantages of the evidence grid representation is its ability to fuse sensor information from different sources. Any number of sensor readings from any number of sensors can be combined, as long as models exist for each sensor type.

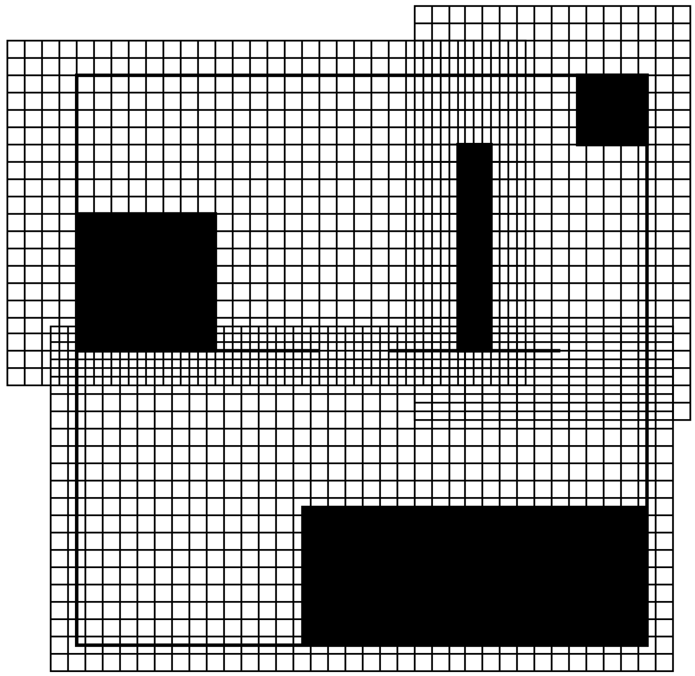

Figure 3. Room grids for an example environment.

Accumulating multiple readings over time is an effective method of filtering out transient changes. Consider a person walking past the robot as it maps a particular region of space. This person's path will cover many grid cells, but each only for a brief moment. Each sonar reading that reflects from the person will increase the occupancy probability of the corresponding cells, but each cell will only be occupied briefly, so all of the subsequent sonar readings incident on the cell will reduce its occupancy probability. As a result, the cells along this path will have a low probability despite the person's passage.

In addition to providing an effective method for combining data from multiple sensor readings, evidence grids have two other advantages for use in dynamic environments. First, they can be updated quickly; using the equations described by Moravec (1988), each cell update can be computed with a single addition. Second, small changes in the environment tend to result in small changes to the corresponding grid representation, which is important for dealing with lasting changes in the environment.

One exception to the second property concerns specular reflection, which occurs when a sonar pulse hits a flat surface and reflects away from (rather than back to) the sensor. As a result, the sensor registers a range that is substantially larger than the actual range. In such cases, a small change in the angle of a surface can result in a substantial change to the evidence grid. One response is to rotate the sonar sensors through a range of angles equivalent to the width of the sonar arc. If both specular and non-specular reflections are possible from a given viewpoint, then both will be incorporated into the evidence grid.

2.2.2 Integrated Spatial Representation

Magellan's integrated spatial representation combines evidence grids with the topological framework of the adaptive place network. The system stores a single room grid for each room in the environment, along with one or more place grids that describe regions of that room. Each room grid contains its own local coordinate frame but also specifies transformations that map the local frame to the frames of adjacent rooms. Each place grid stores the coordinates of its center in the room coordinate frame, as well as its relative orientation. The

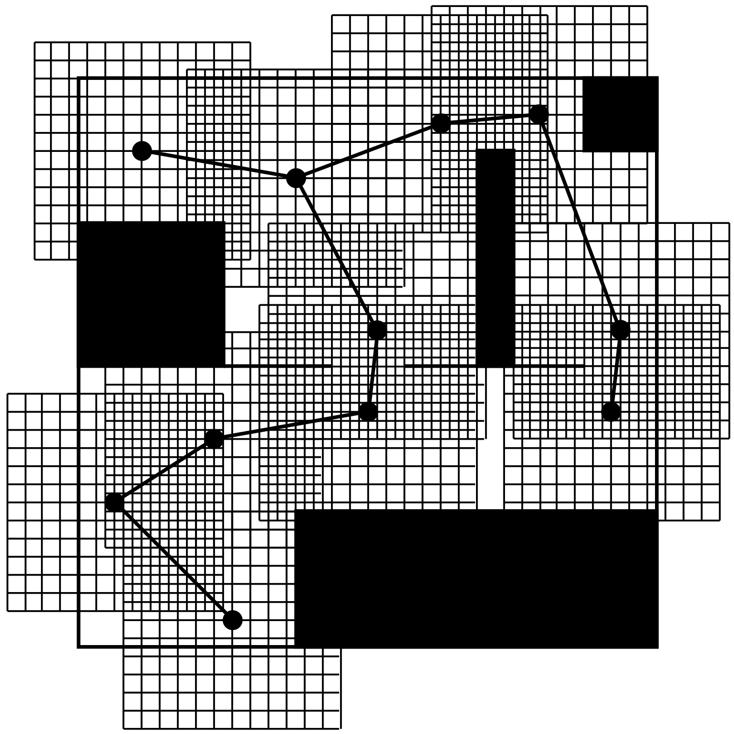

Figure 4. An adaptive place network and place grids for an example environment.

topological network takes the form of directional links between places, each of which specifies the probability that the robot can traverse the link. Taken together, the evidence grids and the place network describe both the details of local places and the large-scale structure of the environment.

Figure 3 shows the room grids for an example environment that consists of three rooms connected by three doorways. Although these evidence grids are separate, they can overlap to a greater or lesser degree, as the figure depicts. Naturally, a given room grid can be more or less complete, depending on the sensory data that Magellan uses for its construction. In the studies reported later, in each room we positioned the robot to collect readings from multiple viewpoints that we selected to guarantee reasonable coverage of that room.

Figure 4 shows an adaptive place network and place grids for the same environment, with place units denoted by black circles and place links by the lines connecting them. In this case, for each gateway between adjacent rooms, Magellan has stored place units on both sides of that gateway, along with a link that connects these places. Other place units, connected by additional place links, occur elsewhere in the rooms. In these figures, the room and place grids have the same orientation, but the system can also handle environments in which they differ. In our experiments, we had Magellan base its place grids on sensor readings taken from one (manually selected) viewpoint, so that each grid describes the view from the center of its associated region.

2.3 Place Recognition

Whenever a robot is turned on in an unknown location, it needs some to way to determine that location. If the robot has a map of the world, it can accomplish this taskn by comparing its current perceptions to those predicted by the map. We will refer to this generic process as *place recognition*, by analogy with that activity in humans.

Magellan's module for place recognition builds a new evidence grid at the robot's current location (the *recognition* grid) and matches it against all grids that have been previously associated with places in the world (the *stored* grids). The system translates and rotates the recognition grid to find the best match with each stored grid, using a hill-climbing algorithm to search the space of possible transformations. This search process lets Magellan recognize a place despite variations in pose.

Each step in the search algorithm relies on rotation and translation processes. The origin of the coordinate frame is located at the center of each grid, corresponding to the robot's position when the grid was constructed. Rotation involves turning the center of each cell in the recognition grid about that grid's origin, whereas translation shifts each center relative to the rotated coordinate frame. The system then compares each recognition cell to the cell in the stored grid, computing a match score to evaluate the transformed grid.

The hill-climbing algorithm applies this process iteratively to find the best transformation between the recognition grid and each stored grid. The system halves the hill-climbing step size when it reaches a local maximum, in order to more precisely locate this maximum. When Magellan reaches a local maximum using the minimum step size, it halts the search and uses the score for the current transformation as the overall match score for the stored grid. The system repeats this process for each of the stored grids and selects the grid with the highest match score as the winner.

The match score that directs the search process is computed in steps, one for each pair of corresponding cells in the recognition and stored grids. This match metric is given by

$$s(i,j) = \begin{cases} 1 & \text{if} \quad p(i) > p_0 \quad \text{and} \quad p(j) > p_0 \\ 1 & \text{if} \quad p(i) < p_0 \quad \text{and} \quad p(j) < p_0 \\ 1 & \text{if} \quad p(i) = p_0 \quad \text{and} \quad p(j) = p_0 \\ 0 & \text{otherwise} \end{cases}$$

where $s(i,j)$ is the match score for corresponding cells i and j, $p(i)$ is the probability that cell i is occupied, $p(j)$ is the probability that cell j is occupied, and p_0 is the prior probability that any cell is occupied. The sum of this measure over all of the corresponding cells gives the total match score for the stored grid with the current transformation.

We developed this match metric, which differs from the one reported in our earlier work (Langley & Pfleger, 1995; Langley, Pfleger, & Sahami, 1997), to deal with the problem of nonindependent sensor readings. Since the sonar cones overlap, their sensor readings are not independent. As a result, the occupancy probabilities in the evidence grid do not accurately reflect the precise probability that each cell will be occupied. However, what is reliable is whether each cell is more likely or less likely to be occupied than the prior probability (or whether it has not been sensed at all, in which case it will be equal to the prior probability). Thus, the match metric increases the match score whenever two corresponding cells are both more likely to be occupied, less likely to be occupied, or unsensed in both the recognition grid and the stored grid.

In addition to providing the robot's initial location, Magellan also uses the place recognition module to confirm the robot's location during the navigation process, which we will discuss shortly. Whenever the robot moves to a new place along its path to the goal, the system invokes place recognition to confirm that the robot has successfully arrived at the desired place.

2.4 Continuous Localization

The place recognition system provides Magellan with an accurate estimate of the robot's initial position, but, as the robot moves through the world, wheel slippage introduces errors into dead reckoning. As a result, any localization system that relies exclusively upon dead reckoning for pose updates will become increasingly

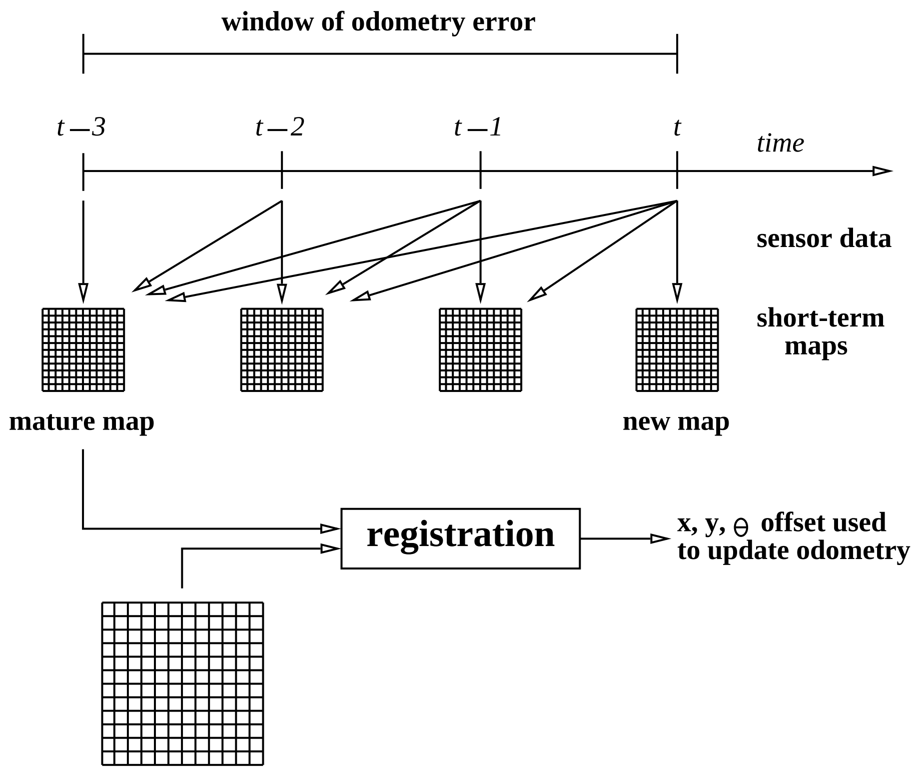

Figure 5. Magellan's procedure for continuous localization.

inaccurate over time. The method for place recognition described above works well for initialization, but it requires that the robot stop to construct a detailed map of its surroundings, which makes it impractical for continuous localization.

In order to support localization without stopping, Magellan incorporates a method for correcting pose estimates incrementally (Schultz et al., 1996). Using this technique, the system builds a series of *short-term* grids based on localized sensor readings and the current odometry, and then registers the short-term and long-term grids. The registration process invokes the hill-climbing algorithm described earlier but uses a slightly different match function in which matching against unknown cells does not add to the match score. This registration produces an offset that the system uses to correct the odometry.

The short-term evidence grid represents the immediate temporal and spatial environment of the robot. Only recent sensor readings of the robot contribute to this structure. Several short-term grids of the robot's environment may exist at the same time, each with a different amount of sensor data contributing to the grid's "maturity", as depicted in Figure 5. Magellan considers a short-term perceptual grid to be mature after the robot has traveled 96 inches. The expected dead reckoning error accumulated over this distance corresponds to the maximum pose error allowed in the grid.

After a short-term grid has matured, the system registers it to correct the pose error and then discards the grid. The registration process carries out a hill-climbing search to align the mature short-term grid and the global grid. This process gives an offset in position and rotation that Magellan uses to update the odometry of the robot. Since the relocalization occurs incrementally and regularly, about once every ten seconds, the search involved in the registration can be constrained to ± 6 inches and ± 2 degrees, and therefore is rapid.

Unlike place recognition, continuous localization cannot determine the robot's position from a completely unknown initial pose; however, once the system has determined a pose, it can maintain the accuracy of that pose as the robot moves through the world. In recent experiments, MAGELLAN accumulated an average of only five inches of translational error during indefinite operation. Moreover, continuous localization sets the encoder positions directly, so its operation is transparent to other processes controlling the robot (like navigation and place recognition).

2.5 The Navigation Process

MAGELLAN's navigation system incorporates both topological path planning, to determine the general shape of the robot's course, and reactive behaviors, to avoid collisions and navigate around unexpected obstacles. The first process operates once for each navigation problem, whereas the second operates once on every time step. However, as we explain below, navigation may be reinvoked if the system encounters difficulties during its execution of the path plan.

The system refers to the adaptive place network when planning a path to its destination. In the network, each link between adjacent places includes a confidence value that represents the robot's estimate of how likely that link is to be traversable. The cost of each link is inversely proportional to this confidence value. Using these cost assignments, MAGELLAN uses Dijkstra's algorithm (Bondy & Murty, 1976) to find the path with least cost from the current place C to the desired place D.

Each path consists of a series of places connected by place links. At each step along the path, reactive behaviors attempt to move the robot toward the center of the next place N, while also attempting to avoid collisions with any nearby obstacles. If MAGELLAN thinks the robot has arrived at place N, it halts and invokes the place recognition module to confirm its location. If this process determines that the robot has succeeded at moving from C to N, the system increases confidence on the link from C to N, which may alter navigation decisions in the future.

If MAGELLAN thinks the robot has failed to move from place C to place N, or if the traversal process exceeds a specified time limit, it also stops and invokes place recognition to determine its location. If the system finds that, indeed, it has not reached N, then it decreases confidence on the link from C to N and reinvokes the search algorithm to find a path from its revised place to the desired place D. Again, this adaptive process may change the system's future route choices.

Of course, if the robot is still in place C, search may again return N as the next place along the path to D, in which case MAGELLAN will try again to move from C to N. If the failure was due to transient obstacles or similar problems, then this attempt may well succeed. But if the failure resulted instead from long-term changes, such as the introduction of a permanent obstacle, then repeated attempts will eventually reduce confidence from C to N enough that the search algorithm will select an alternative path from C to destination D.

Another important type of event occurs when the robot crosses into a new room. When place recognition determines that this has happened, MAGELLAN activates the corresponding room grid for use in continuous localization and transforms the Cartesian location to the new room's coordinate frame. This lets the continuous localization process maintain an accurate pose estimate as the robot moves from room to room, even though there is a separate grid for each room.

Thus integrated with place recognition and continuous localization, the navigation system lets the robot start at an unknown position, recognize its current location, and navigate to anywhere in its global map while keeping track of its location. This integrated set of abilities makes MAGELLAN considerably more adaptive than many existing systems for mobile robotics.

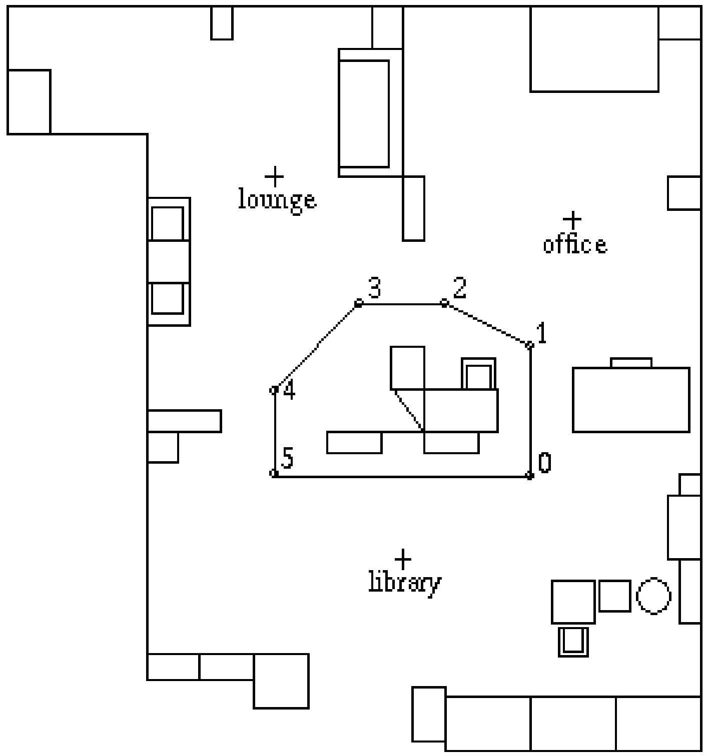

Figure 6. Layout of the NCARAI robotics laboratory.

3. Experimental Results

Most mobile robot systems are tested within a single environment or a single set of closely related environments. For Magellan, we wanted to demonstrate a more general capability, so we tested the system in two completely different settings: the Navy Center for Applied Research in Artificial Intelligence (NCARAI) in Washington, DC, and the Robotics Laboratory at Stanford University. In this way, we could test whether Magellan's capabilities would be robust to changes in the shapes of rooms and halls, the number and types of obstacles, and the number and frequency of humans in the vicinity.

3.1 Navy Center for Applied Research in Artificial Intelligence

In the first study, performed at NCARAI, we arranged the furniture in the robotics lab to form three distinct rooms, dubbed the library, the office, and the lounge. Figure 6 shows the layout of the NCARAI robotics lab. Each of the numbers in the figure labels a place unit, whereas the lines indicate the place links that connect the place units.

Figure 7 shows the room grids that Magellan constructed for the NCARAI robotics laboratory. Each of these grids consists of 128×128 cells covering a 50 foot $\times$ 50 foot area. The system constructed these grids by taking sonar and laser readings at a number of different viewpoints (four to six), which we selected to ensure that the robot mapped the entire room. Eleven sets of 16 sonar readings were taken at each viewpoint, at two degree intervals over a 22 degree range (the approximate width of each sonar cone). Some 36 sets of laser readings were taken at each viewpoint, at ten degree intervals over a 360 degree range. The amount of time required to gather sensor information from each viewpoint was approximately 90 seconds.

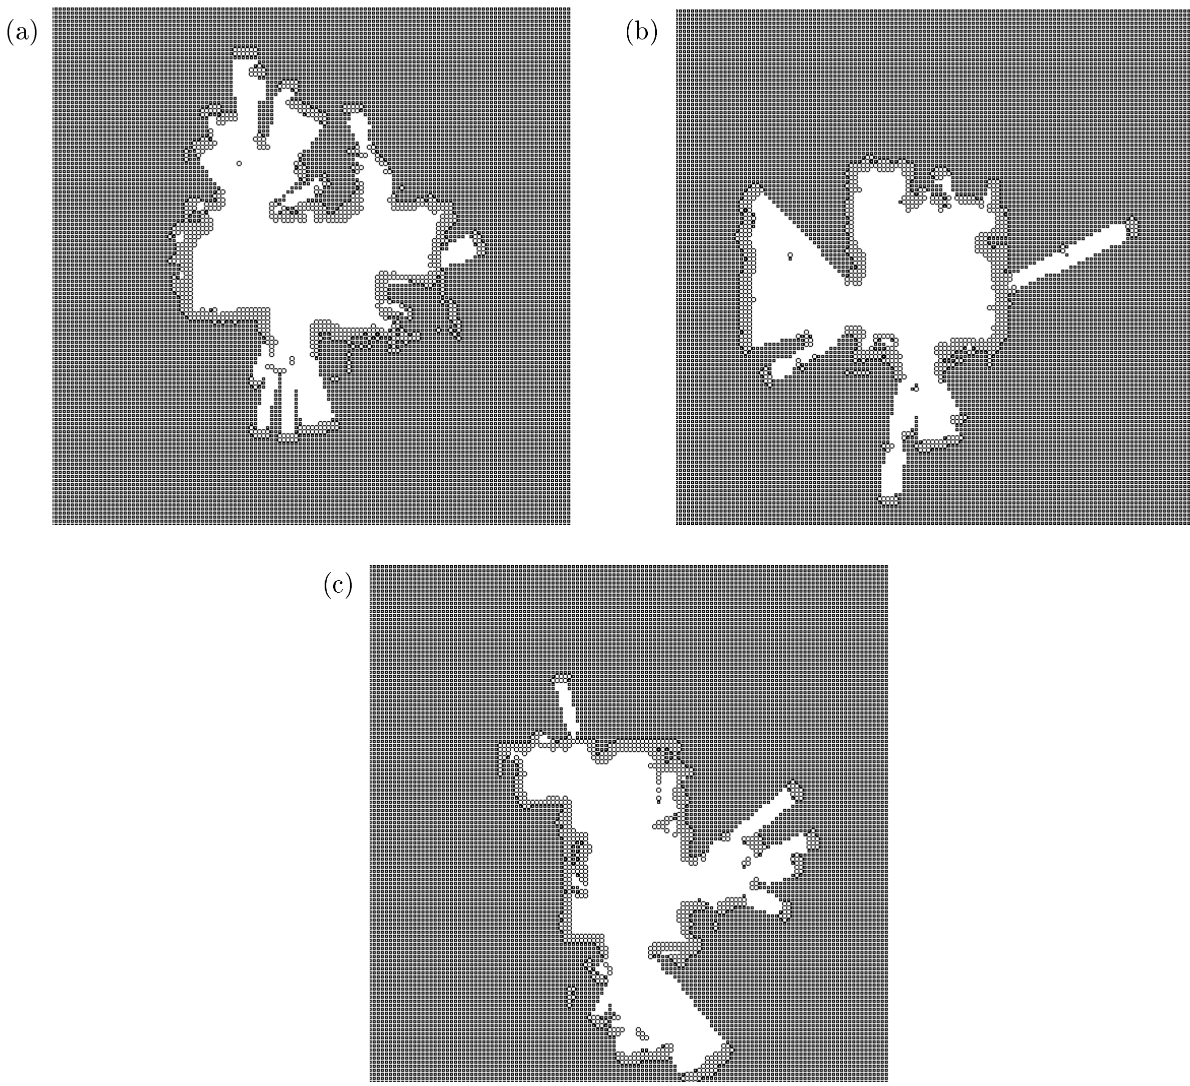

Figure 7. Room grids constructed for the NCARAI robotics laboratory (a) lounge, (b) office, and (c) library.

Figure 8 shows the place grids for the six places we defined in the environment. Magellan constructed each of these grids from a single viewpoint at the center of the corresponding place, to which we moved the robot manually. These grids are the same size (50 feet $\times$ 50 feet) and resolution (128 $\times$ 128) as the room grids. The procedure for obtaining sensor information was the same as for the room grids (11 sets of sonar readings at two degree intervals, 36 sets of laser readings at ten degree intervals), but taken from a single viewpoint, so the total amount of time required to build each place grid was approximately 90 seconds.

The six places in this environment define 15 distinct navigation tasks that involve moving from one place to another. We ran Magellan on all of these tasks, which ranged from traversing a single room to moving through all three rooms. There were no failures in either place recognition or navigation on any of these tasks, including a demonstration trial that had several observers stationed throughout all three rooms. These results suggest that our approach has the robustness we had anticipated.

(a)

(b)

(c)

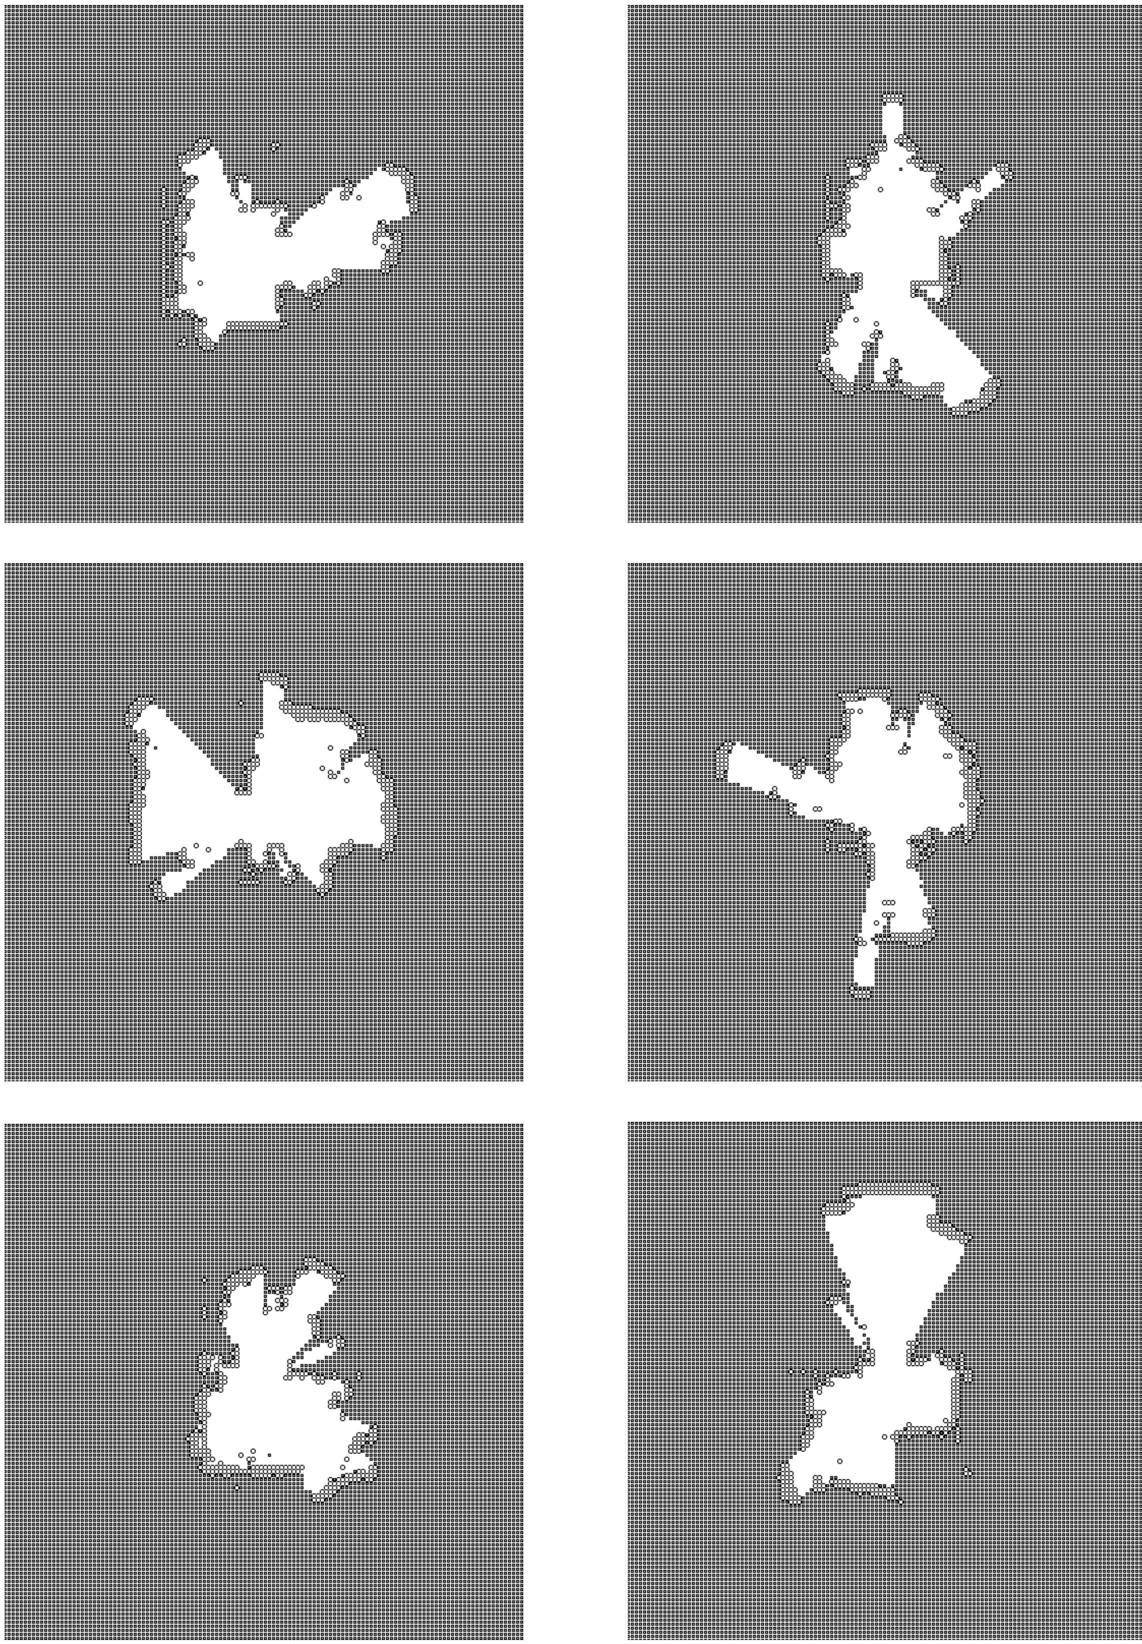

Figure 8. Place grids constructed for the NCARAI robotics laboratory (a) lounge, (b) office, and (c) library.

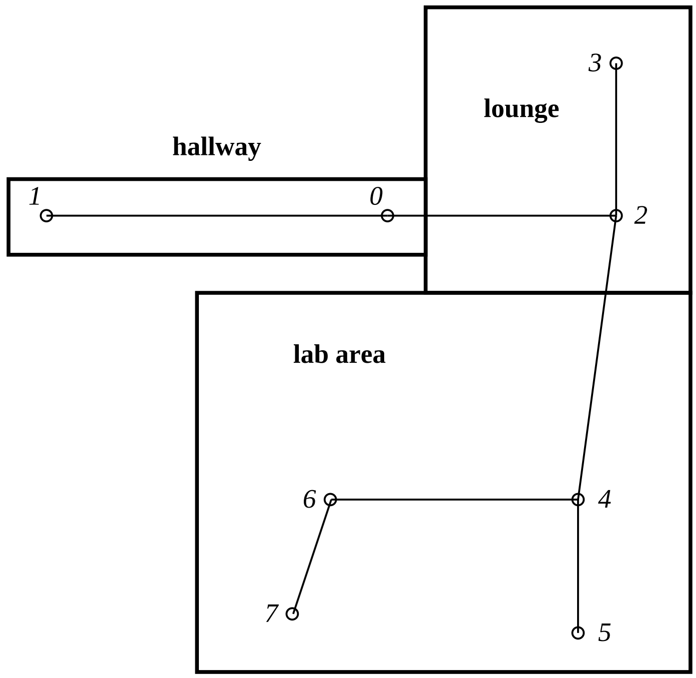

Figure 9. Layout of the Robotics Laboratory at Stanford University.

Continuous localization did not play a prominent role in the NCARAI runs because it requires a certain amount of motion within a room before appreciable pose error accrues and before enough sensor data accumulates to perform localization. For most trials, the robot did not travel enough in any single room for continuous localization to take effect. In a few cases, we gave the system extra motion tasks and for these cases continuous localization worked well, limiting the robot's pose error to about six inches and one degree.

Although this study demonstrated that Magellan could operate robustly in an office environment in the NCARAI robotics laboratory, we wanted to test the generality of our approach to integrating place recognition, continuous localization, and navigation. To this end, we conducted another study, to which we will now turn.

3.2 The Stanford Robotics Laboratory

We performed a second set of runs in the Robotics Laboratory at Stanford University. Unlike many laboratories, which are isolated from outside interference, the Stanford Laboratory is a public space in the center of an academic building. The rooms in this area are subject to constant human traffic, as researchers and students wander from offices to classrooms, often stopping to chat in the hallways or directly in front of the robot. In addition, the lab is used by several different robotics groups, each of which frequently rearranges the furniture to suit its own needs. As such, this lab provides a particularly challenging environment for robot localization and navigation.

We defined three rooms within this environment: the main lab area, the adjoining lounge, and an adjacent hallway. Magellan learned seven places within these three rooms and connected these places to form the adaptive place network shown in Figure 9. The corresponding room grids appear in Figure 10. Like the grids used at NCARAI, these consist of 128×128 cells covering an area 50 feet $\times$ 50 feet. As in the previous study, we positioned the robot at a number of different viewpoints (four to five), which we selected to support viewing the entire room using sonar and laser sensors.

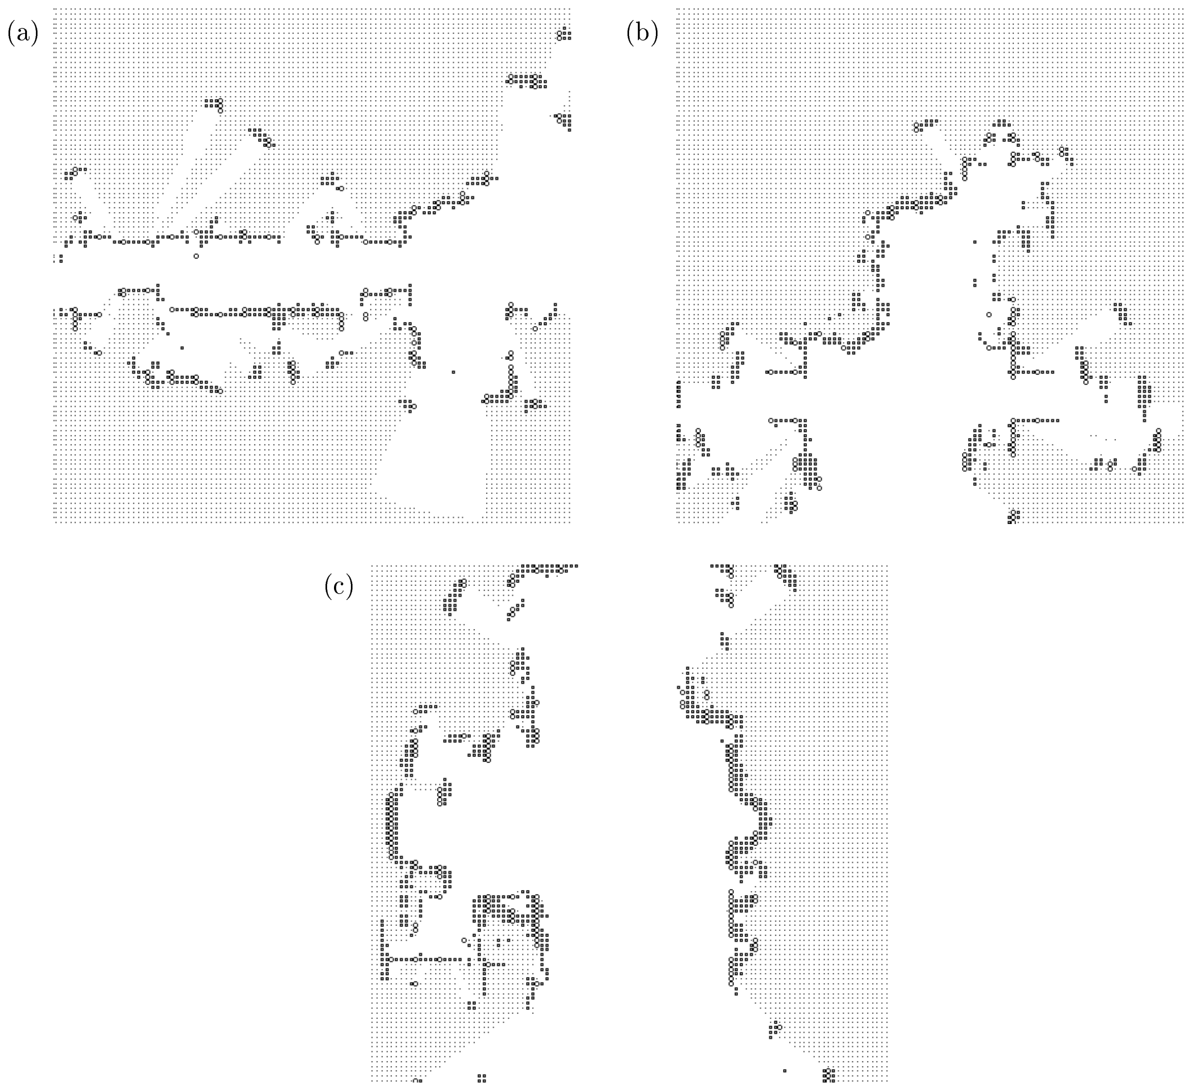

Figure 10. Room grids for the (a) hallway, (b) lounge, and (c) main lab in Stanford University Robotics Laboratory.

Figure 11 (a) shows the two place grids created for the hallway, Figure 11 (b) depicts the two place grids constructed for the lounge, and Figure 12 gives the four place grids for the main lab. As in the previous study at NCARAI, these grids are the same size and resolution as the room grids, and Magellan constructed them from 11 sets of sonar readings and 36 sets of laser readings, which it collected after we moved the robot to the center of each place.

As before, we selected 15 navigation tasks, some of which required traversal of a single room and others involving movement to other rooms. In each of these trials, Magellan recognized its initial place correctly and it navigated successfully to its specified destination, despite the presence of moving people and rearranged furniture. For each task, we initially put the robot near the center of one of the places and gave it an approximate (plus or minus five degrees) estimate of its heading.[3]

3. We provided this information because the current hill-climbing search used during place recognition has problems handling rotations larger than five degrees. However, we could reduce this problem, at greater computational expense, by starting multiple searches from different initial rotations.

(a)

(b)

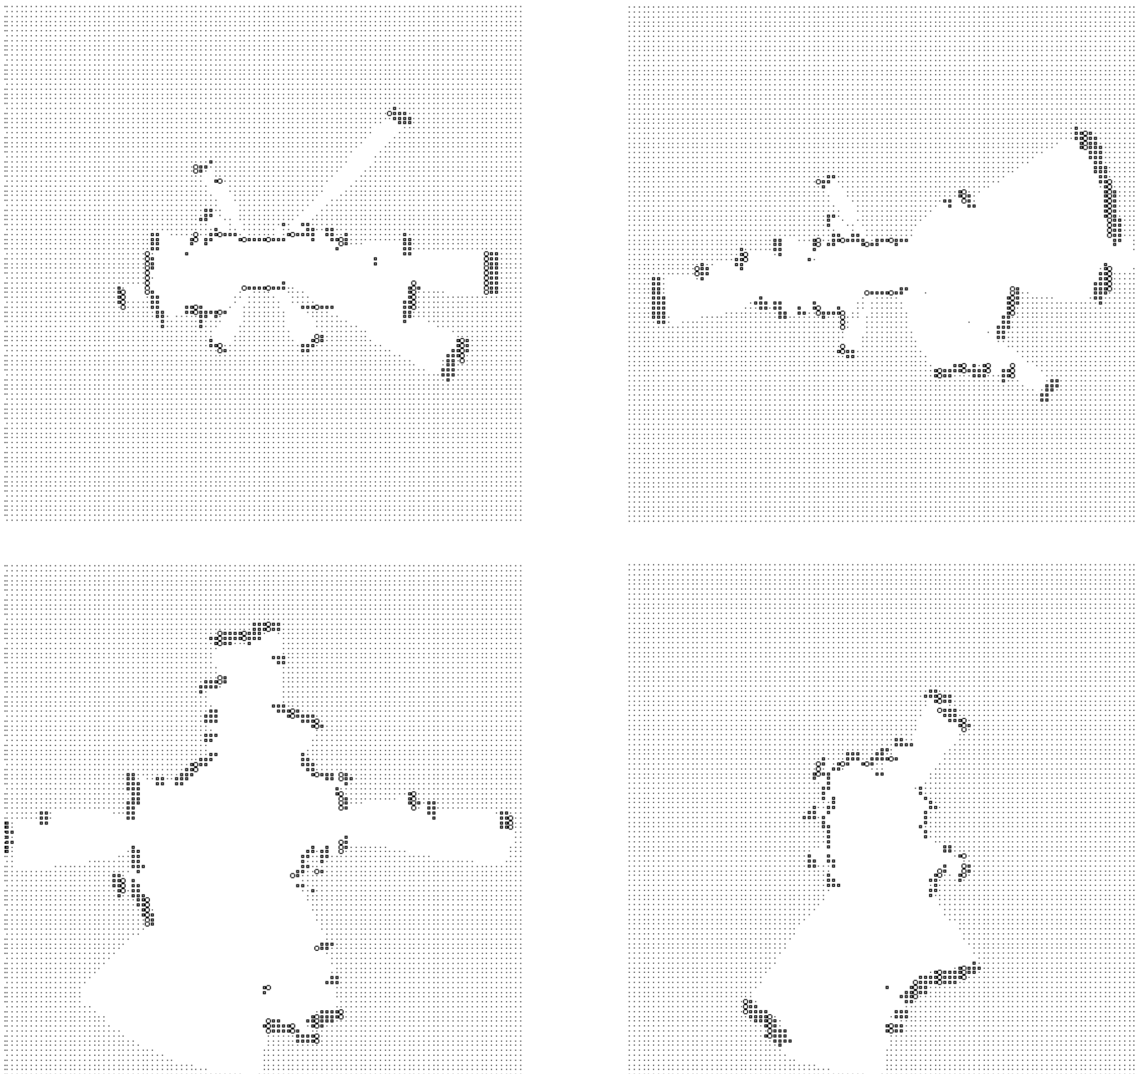

Figure 11. Place grids for the Robotics Laboratory (a) hallway and (b) lounge at Stanford University.

In these trials, place recognition typically determined the robot's position to within three inches and two degrees. The recognition process required approximately two minutes on a Silicon Graphics Indy workstation, including 15 seconds for the sonar sweep, 45 seconds for the laser sweep, and 60 seconds for the grid matching.

Continuous localization was rarely invoked, due to the relatively short distances between adjacent places. When the robot moved without place recognition, continuous localization determined the robot's position to within ten inches and four degrees. This error was slightly larger than in the NCARAI study, probably due to the more dynamic nature of the Stanford environment.

Magellan's successful performance in the Stanford Robotics Laboratory showed that the system's competence was not limited to a single environment. Moreover, this study demonstrated the additional ability to recognize familiar places and to navigate robustly even when the robot was situated within a highly dynamic environment that was substantially different from the NCARAI office environment, where we initially developed the Magellan architecture.

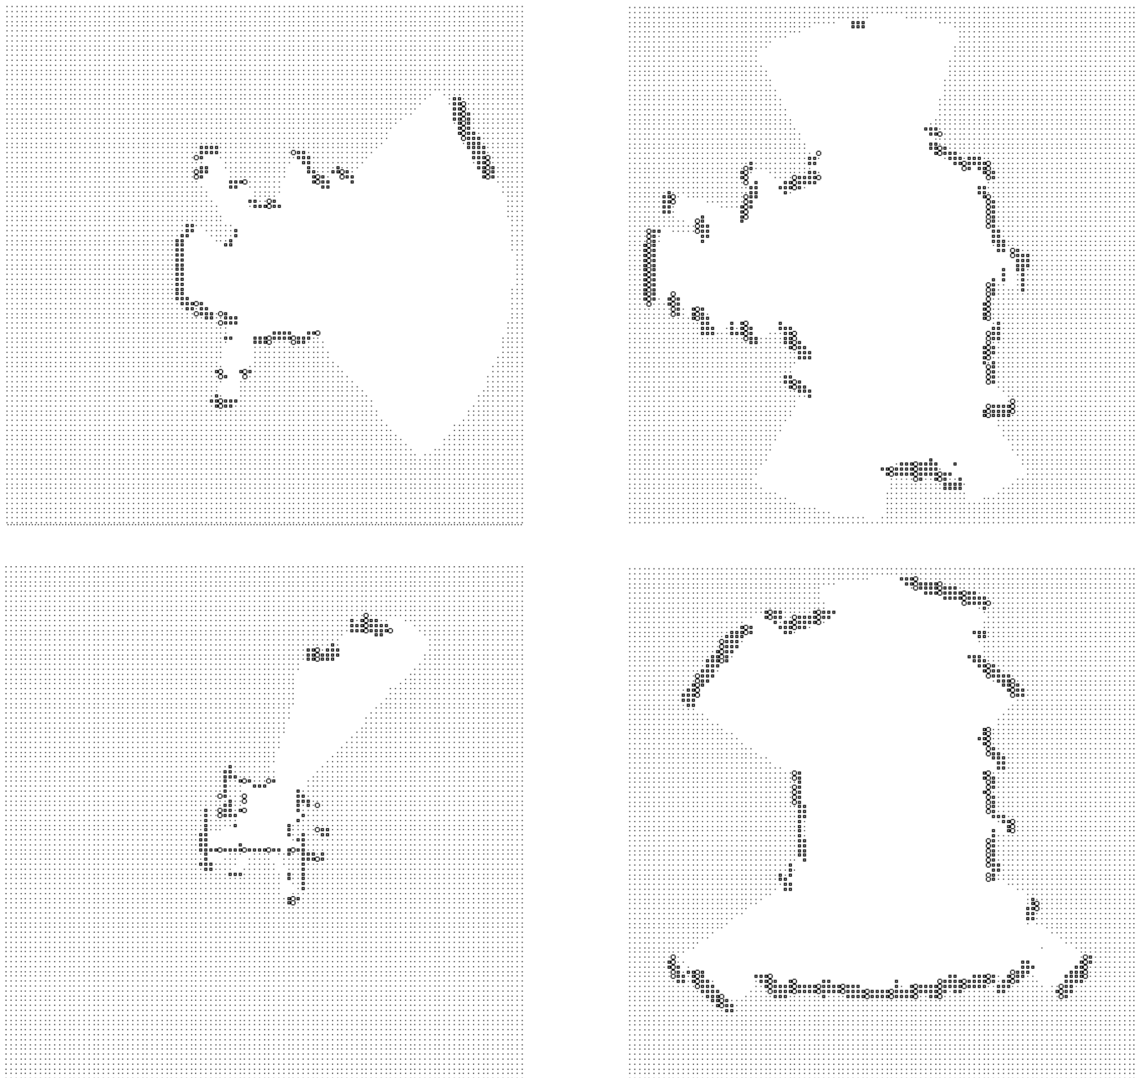

Figure 12. Place grids for the main lab area at the Stanford University Robotics Laboratory.

4. Related Work

There has been considerable research on localization and navigation for mobile robots. Our work on Magellan differs from most of these efforts in integrating techniques for place recognition, continuous localization, and adaptive navigation – all of which are robust to changes encountered in dynamic, real-world environments. However, we should briefly consider this earlier work on mobile robotics and its relation to the current architecture.

Moravec and Elfes (1988) have used evidence grids for mobile robot navigation, but their approach differs widely from our own. Their systems use evidence grids as the only representation of space, and they draw on traditional search methods (such as A* and relaxation) to find paths through the environment at the level of individual grid cells. Position uncertainty is handled by blurring the robot's sensor readings over time, with the amount of blur corresponding to the uncertainty in the robot's position. In contrast, Magellan plans paths through its topological network, leaving the low-level details of navigation to its reactive behaviors. Moreover, our system can maintain an accurate position estimate using continuous localization and it can localize from a completely unknown position using place recognition.

Kuipers and Byun (1993) have developed a spatial learning system that identifies distinctive places, as defined by a set of pre-defined criteria (e.g., equal range readings in three directions), and links these places with edges specifying transit behaviors that take the robot from one place to another. Like Magellan, it generates path plans using search through this topological network and invokes its transit behaviors to execute them. One difference is that their system describes its 'places' in terms of sensor readings, whereas ours uses inferred evidence grids, which lets it represent the detailed structure of each place. We believe this provides a more general approach to place recognition that, as we have shown elsewhere (Yamauchi & Langley, 1996), is robust even to substantial changes in the environment. On the other hand, Magellan currently relies on humans to identify useful places, whereas Kuipers and Byun's system invokes heuristics to learn places in an unsupervised manner.

Mataric (1992) reports a learning system for mobile robotics which combines a reactive controller with a distributed representation that reflects the topological structure of the environment. The system also uses the reactive controller to deal with transient changes, a feature that it shares with Magellan. One significant difference is that Mataric's approach represents places in terms of a fixed set of simple landmark types, whereas our system relies on the richer formalism of evidence grids to support place recognition and continuous localization. Magellan can also respond to topological changes by modifying the confidence values on links that connect the place units. In addition, it uses a more flexible method for arbitrating behaviors, similar to the one in Langer, Rosenblatt, and Hebert's (1994) DAMN architecture, that allows a compromise between commands from different behaviors.

Schiele and Crowley (1994) take yet another approach to position estimation. Their methods involve matching surfaces in a stored geometric map against line segments extracted from evidence grids using Hough transforms and Kalman filtering. Clearly, our work on Magellan differs by matching evidence grids against one another directly, which should work well in a broader range of indoor settings. Another distinction is that, to date, Schiele and Crowley's research has focused on static environments, whereas we have shown our approach is robust in dynamic environments.

Burgard, Fox, Hennig, and Schmidt (1996) have developed a localization technique that compares current sensor readings to a stored evidence grid. Their approach generates a probability distribution over all possible positions, rather than a single position estimate. However, their method requires a full minute to localize in a small office, so the time required to localize in a large environment may be prohibitive due to its reliance on a single global grid. In contrast, our place recognition system uses a separate grid for each place and has scaled effectively to large environments (Yamauchi & Langley, 1996).

Thrun and Bücken (1996) have integrated topological and grid representations for use in localization and navigation, but their approach assumes that walls will only be parallel or perpendicular to each other. Although this may be true in most indoor environments, obstacles can make it difficult to determine the actual orientation of walls. Our approach differs in using all detected features of the environment during localization, and by making no a priori assumptions about the structure of the world. Also, Thrun and Bücken's technique generates a topological map from a single global evidence grid, whereas Magellan uses a separate grid for each node in the topological map.

Koenig and Simmons (1996) have used partially observable Markov decision processes for localization in a dynamic, real-world environment. Their system uses a simple set of perceptual features (walls and openings to either side of the robot) and localizes by determining the sequence of positions most likely to have generated those features. Our approach concentrates on the detailed spatial structure of the environment, whereas theirs focuses on information obtained from sequences of simple spatial features observed over time as the robot moves through the world. As a result, Magellan should be more useful in environments with complex local features, such as offices containing furniture, where there is substantial information in the immediate view of the surroundings, whereas their method should fare better in environments that contain few distinctive local features (such as empty hallways) but more high-level structure (such as patterns of interconnecting corridors).

5. Concluding Remarks

In summary, Magellan is an integrated architecture for mobile robotics with the ability to recognize places and navigate through the world while maintaining an accurate estimate of its position. The place recognition module builds an evidence grid – a probabilistic representation of occupancy – that describes the robot's current surroundings and matches this structure against grids stored for previously visited places. A hill-climbing registration process translates and rotates the current grid to find the best match with a grid stored in memory.

The system passes the current location to the module for continuous localization, which repeatedly updates the robot's position as it moves through the world. Continuous localization incrementally modifies the robot's encoder position to reflect the best match between recent perceptions, stored in short-term evidence grids, and the room layout, represented in the long-term map. This process differs from place recognition in that, already knowing the robot's approximate location, it can operate in real time.

Adaptive navigation lets the robot move to a specified destination within a topological network that represents the environment. This process combines topological route planning to provide high-level direction with reactive behaviors to provide robust execution. Based on the success or failure of the executed plan, Magellan increments or decrements the confidences on links in its topological network, which eventually leads the system to plan alternative paths to its goal if the original path is blocked by unexpected obstacles.

We have tested Magellan in two environments with substantially different characteristics. The first was a simulated office setting within a research laboratory, relatively isolated from human interference; the second environment was a large area within an academic building, subject with constant human traffic and frequently rearranged furniture. In both cases, Magellan was able to robustly recognize places, continuously localize, and navigate to its destinations even when the environment was in flux.

We are currently developing Magellan's successor, ARIEL (Autonomous Robot for Integrated Exploration and Localization) (Yamauchi et al., 1997), which incorporates both the continuous localization system described in this paper and a new method for autonomously exploring and mapping unknown environments. This technique – frontier-based exploration (Yamauchi, 1997) – directs exploration toward the boundaries between known open space and unknown territory. The initial version of ARIEL has successfully explored real-world office environments while using continuous localization to maintain an accurate estimate of its position. In the future, we plan to integrate the place recognition system from Magellan into ARIEL, thus creating a system that can localize itself within an existing map or build a map from scratch, and then navigate within that environment while exploring new frontiers.

Acknowledgements

This work was funded by the Office of Naval Research under Grant No. N00014-94-1-0505. We thank Hans Moravec for sharing his software, and we thank Nils Nilsson and Barbara Hayes-Roth for providing access to their Nomad mobile robot.

References

Bondy, J. A., & Murty, U. S. R. (1976). *Graph theory with applications*. New York: Elsevier.

Burgard, W., Fox, D., Hennig, D., & Schmidt, T. (1996). Estimating the absolute position of a mobile robot using position probability grids. *Proceedings of the Thirteenth National Conference on Artificial Intelligence* (pp. 896–901). Portland, OR: AAAI Press.

Graves, K., Schultz, A., and Adams, W. (1997). Continuous localization in changing environments. *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*. Monterey, CA.

Koenig, S., & Simmons, R. (1996). Unsupervised learning of probabilistic models for robot navigation. *Proceedings of the 1996 IEEE International Conference on Robotics and Automation* (pp. 2301–2308). Minneapolis, MN.

Kuipers, B., & Byun, Y. (1993). A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, *8*, 47–63.

Langer, D., Rosenblatt, J., & Hebert, M. (1994). A behavior-based system for off-road navigation. *IEEE Journal of Robotics and Automation*, *10*, 776–782.

Langley, P., & Pfleger, K. (1995). Case-based acquisition of place knowledge. *Proceedings of the Twelfth International Conference on Machine Learning* (pp. 244–352). Lake Tahoe, CA: Morgan Kaufmann.

Langley, P., Pfleger, K., & Sahami, M. (1997). Lazy acquisition of place knowledge. *Artificial Intelligence Review*, *11*, 315–342.

Mataric, M. (1992). Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation*, *8*, 304–312.

Moravec, H. (1988). Sensor fusion in certainty grids for mobile robots. *AI Magazine*, *9*, 61–74.

Moravec, H., & Elfes, A. (1985). High resolution maps from wide angle sonar. *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 116–121). St. Louis, MO.

Schiele, B., & Crowley, J. L. (1994). A comparison of position estimation techniques using occupancy grids. *Robotics and Autonomous Systems*, *12*, 163–171.

Schultz, A., Adams, W., & Grefenstette, J. (1996). *Continuous localization using evidence grids* (NCARAI Technical Report AIC-96-007). Naval Research Laboratory, Washington, DC.

Thrun, S., & Bücken, A. (1996). Integrating grid-based and topological maps for mobile robot navigation. *Proceedings of the Thirteenth National Conference on Artificial Intelligence* (pp. 944–950). Portland, OR: AAAI Press.

Yamauchi, B. (1997). A frontier-based approach to autonomous exploration. *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*. Monterey, CA.

Yamauchi, B., & Beer, R. (1996). Spatial learning for navigation in dynamic environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics, Special Issue on Robot Learning*, *26*, pp. 496–505.

Yamauchi, B., & Langley, P. (1996). Place learning in dynamic real-world environments. *Proceedings of RoboLearn 96: An International Workshop on Learning for Autonomous Robots* (pp. 123–129). Key West, FL.

Yamauchi, B., & Langley, P. (1997). Place recognition in dynamic environments. *Journal of Robotic Systems*, *14*, 107–120.

Yamauchi, B., Schultz, A., Adams, W., Grefenstette, J., & Perzanowski, D. (1997). ARIEL: Autonomous robot for integrated exploration and localization. *Proceedings of the Fourteenth National Conference on Artificial Intelligence* (pp. 804–805). Providence, RI: AAAI Press.