Place Learning in Dynamic Real-World Environments

Brian Yamauchi[1] (yamauchi@robotics.stanford.edu)
Pat Langley[2] (langley@robotics.stanford.edu)
Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306

Abstract

In this paper, we present an approach for mobile robot localization designed for use in dynamic environments. Our approach integrates evidence grids within a topological/metric network that can be used for navigation. Place learning consists of associating evidence grids with places in the topological network. Place recognition consists of building an evidence grid at the current location and using a registration procedure based on hill climbing to find the best match between the current grid and the grids associated with places in the network. This approach has been implemented on a real mobile robot and has been tested in a real-world office environment containing multiple forms of dynamic change. In these experiments, this approach demonstrated robust localization in the presence of transient changes (such as moving people) and lasting changes (such as rearranged furniture) in the environment.

1. Introduction

A central issue in mobile robotics is localization: how can a robot use its sensors to determine its location? A particular example is the "kidnaped robot problem" (Engelson, 1994) where, once a robot has explored its environment, it is turned off and transported to an unknown location. At the new location, the robot must use its sensors to survey its surroundings and determine its position based upon the knowledge previously acquired about the environment.

Many localization strategies combine place learning with place recognition. Place learning consists of associating perceptions with the locations visited by the robot. Place recognition localizes the robot by finding the best match between current perceptions and those associated with each place. While much research has been done in place learning and place recognition, most of this work has been limited to static environments that do not change over time. In contrast, human environments are constantly changing: people walk down hallways and across rooms; they move chairs and desks; they open and close doors. Any localization strategy that relies upon a static world is likely to fail in an environment that contains human beings.

The goal of our research is to develop methods of place learning and place recognition that are robust to the types of change typically encountered in human environments. In particular, this includes both transient changes (such as those caused by people walking past the robot) and lasting changes (such as rearranged furniture). In this paper, we describe such a method, which uses evidence grids for place learning and a grid-matching procedure based on hill climbing for place recognition. This method has been successfully implemented on a real mobile robot in a dynamic, real-world office environment, and we present initial results from experiments with this system.

2. Place Learning

2.1 Learning a Topological Map

The localization system described in this paper is the newest component of ELDEN (Exploration and Learning in Dynamic ENvironments), an integrated mobile robot system developed for exploration, learning, and navigation in dynamic, real-world environments. An earlier version of this system, without place learning and place recognition capabilities, is described by Yamauchi and Beer (in press).

Place learning consists of building an evidence grid for a region in space and associating it with a place in the environment. Each place is represented as a node within a topological/metric map, and each node stores the Cartesian location of the corresponding place. The topological component is included for navigation purposes, even though it is not used for place recognition.

[1] Also affiliated with the Center for the Study of Language and Information, Stanford University.
[2] Also affiliated with the Robotics Laboratory, Computer Science Department, Stanford University.

Each place corresponds to a region five feet in diameter. We selected this place size so that the topological map could represent the traversable paths through the environment for a robot that is roughly two feet in diameter. The maximum reliable range for the sonar sensors is roughly fifteen feet, so each grid covers an area with a fifteen foot radius, and there is substantial overlap between adjacent grids.

Initially the robot starts with an empty map. The robot's starting location becomes the first place in the new map. As the robot moves through the world, the system creates a new place whenever the robot moves out of the space contained in the existing place regions, and creates a topological link between the new place and the place corresponding to the robot's previous location. The system also builds a new grid, which it associates with the place unit. Currently the paths taken to explore the world are determined manually, but we plan to automate this exploration procedure in the future.

2.2 Constructing Evidence Grids

2.2.1 Prior Probabilities and Sensor Models

Evidence grids are a spatial representation developed by Moravec and Elfes (1985) that describe space as a Cartesian grid where each cell has a certain probability of being occupied. Initially, each of these cell probabilities 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 errors in the prior probability, and 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. Each sensor model describes the probability that cells are occupied given the reading received. This model depends on the characteristics of the individual sensor. For example, a sonar sensor emits a pulse of sound in a cone that expands as it gets farther from the transducer. When this pulse hits an obstacle, it is reflected back to the sensor. A range reading of R indicates that an obstacle has been detected somewhere along the sonar arc at range R, so the occupancy probability of all cells along this arc should be increased.

At the same time, this reading indicates that no obstacle was detected at a range closer than R, so all cells within the sonar cone at ranges less than R should have their occupancy probability reduced. Cells at ranges beyond R are not affected, since the obstacle at range R prevented the sonar from obtaining any information about them. In practice, we use a similar, but more realistic, sonar model that also considers sonar attenuation at the edges of the sensor cone and reduced likelihood of sonar returns at longer ranges.

One of the major advantages of the evidence-grid representation is its ability to fuse sensor information. Any number of sensor readings from any number of sensors can be combined as long as models exist for each sensor type.

2.2.2 Updating Evidence Grids

Formally, evidence grids provide a means for combining information from sensor readings in an elegant way (Moravec, 1988). If X represents information such as a sensor reading, then $p(o|X)$ is the probability that a cell is occupied given X, and $p(\neg o|X)$ is the probability that this cell is not occupied given X. Thus, from Bayes' theorem:

$$\frac{p(o|X)}{p(\neg o|X)} = \frac{p(X|o)}{p(X|\neg o)} \times \frac{p(o)}{p(\neg o)}$$

where $p(X|o)$ is the probability of receiving information X given that this cell is occupied, $p(X|\neg o)$ is the probability of receiving information X given that this cell is not occupied, $p(o)$ is the prior probability that any given cell is occupied, and $p(\neg o)$ is the prior probability that any given cell is unoccupied, where $p(\neg o) = 1 - p(o)$.

If A represents the current grid state and B represents the information from a new sensor reading, then cell occupancy probabilities can be combined using:

$$\frac{p(o|A \cap B)}{p(\neg o|A \cap B)} = \frac{p(o|A)}{p(\neg o|A)} \times \frac{p(o|B)}{p(\neg o|B)} \ .$$

This expression assumes that A and B represent independent information, which is not true when a particular point can be sensed more than once (by the same or different sensors). In practice, this approximation means that the overall occupancy results tend to be accurate, but the numerical occupancy probabilities are not reliable. For example, if the sonar cones overlap for two sensor readings, the cells in the overlap will have their probabilities increased or decreased twice, as if the two sensor readings provided independent information about the structure within these region.

Konolige (1995) presents one approach to dealing with this problem. In this scheme, pose information is stored with each cell, indicating the incident direction of each sonar reading. The method considers only the first sonar reading from a particular direction for each cell – it ignores subsequent readings. This approach works well in static environments, but is not well-suited to dynamic environments, since the early state of the world will become "frozen" into the grid, and the grid will not be updated to reflect future changes that occur in the world. Instead, our approach accumulates multiple sensor readings over time, using the standard evidence grid formulation, and then we design our grid matching function (described in Section 3) to be tolerant to the uncertainty in cell occupancy probabilities.

2.2.3 Advantages of the Approach

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. However, each cell will only be occupied briefly, so all of the other sonar readings incident on this cell will reduce its occupancy probability. As a result, the cells along this path will have a low occupancy 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 a logarithmic transformation of the equations described above, each cell update can be computed with a single addition. Second, small changes in the environment tend to produce small changes to the corresponding grid representation. This property is important for handling lasting changes in the environment.

One exception to the second property is the case of specular reflections, 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. Because of this, a small change in the angle of a surface could potentially result in a substantial change to the evidence grid. Konolige (1995) also suggests a method for dealing with specular reflections by ignoring all sonar readings if they would imply that previously occupied cells are unoccupied (as would occur if a specular reflection were to overlap an obstacle). However, this would not work for dynamic environments, since a previously occupied space may actually have become unoccupied due to changes in the world. Instead, during the construction of each evidence grid, our method rotates the sonar sensors through a range of angles equivalent to the width of the sonar arc. As a result, if both specular and non-specular reflections are possible from a given viewpoint, then both will be incorporated into the evidence grid.

We used a Nomad 200 mobile robot, shown in Figure 1, in our research. This robot is equipped with sixteen sonar sensors, evenly spaced around the base at 22.5 degree intervals. In order to build each evidence grid, the robot remains at the center of the place region and takes eleven sets of sixteen sonar readings at two degree intervals (for a total of 176 sonar readings for each grid).

Figure 1: The Nomad 200 mobile robot.

3. Place Recognition

Place recognition in ELDEN consists of building a new evidence grid at the current location and matching this grid (the recognition grid) against all of the grids that have been previously associated with places in the world (the learned grids). The system translates and rotates the recognition grid to find the best fit with each of the learned grids.

ELDEN uses a multiple resolution hill-climbing algorithm to search the space of possible translations and rotations. Figure 2 illustrates the translation operators considered during a single step of this search process. The hill-climbing algorithm starts with the null translation and rotation, then takes steps in the space of possible translations and rotations in order to maximize the match between the recognition grids and each learned grid.

Translations and rotations are defined over evidence grids in a straightforward way. The origin of the coordinate frame is located at the center of each grid, corresponding to the robot's position when it constructed the grid. Each cell in the recognition grid is translated by displacing the point corresponding to the center of each cell and determining into which cell the new point would fall in the learned grid. Each cell in the recognition grid is rotated by computing the vector from the origin to the center of the cell, then rotating this

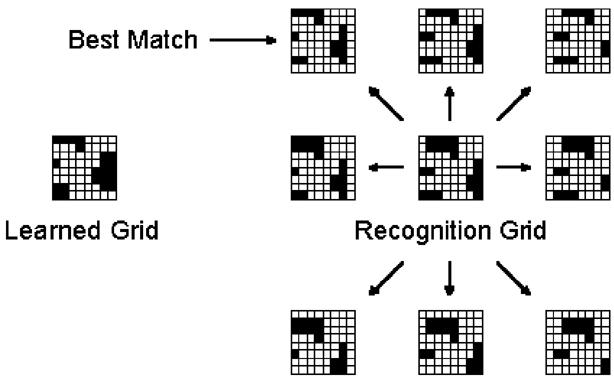

Figure 2: Translating and rotating a recognition grid to align it with a learned grid during place recognition.

vector around the origin, and determining into which cell the new vector would fall in the learned grid.

The system computes a match score for each pair of corresponding cells in the recognition grid and in the learned grid. The match metric is given by:

$$s_{ij} = \begin{cases} 1 & \text{if } p_i > p_0 \text{ and } p_j > p_0 \\ 1 & \text{if } p_i < p_0 \text{ and } p_j < p_0 \\ 1 & \text{if } p_i = p_0 \text{ and } p_j = p_0 \\ 0 & \text{otherwise} \end{cases}$$

where s_{ij} 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. This score is summed over all of the corresponding cells, and the total is the match score for the learned grid given the current transformation.

We developed this match metric to deal with the problem of non-independent sensor readings. The multiple sonar readings taken to filter transient changes and deal with specular reflections 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 either both more likely to be occupied, less likely to be occupied, or unsensed in both the recognition grid and the learned grid.

The hill-climbing algorithm applies this process iteratively to find the best transformation between the recognition grid and each learned grid. The step size for hill climbing is initially set to 7.5 inches and 5 degrees and is halved when a local maximum is reached, in order to more precisely locate this maximum. The system repeats this process twice. When it reaches a local maximum using the minimum step size (1.875 inches and 1.25 degrees), it stops the search and uses the score for the current transformation as the overall match score for the learned grid. The search also stops if the translation exceeds half of the place region radius (1.25 feet) or the rotation exceeds the maximum expected compass error (10 degrees).

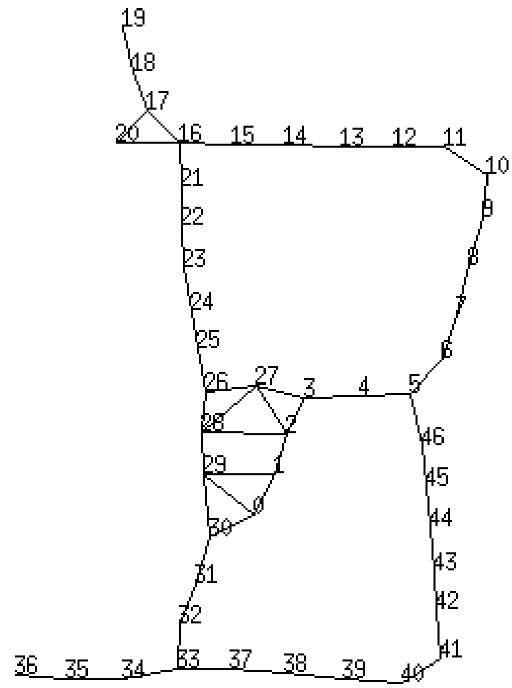

Figure 3: Topological/metric map of learned places.

The system repeats this process for each of the learned grids, and selects the grid with the maximum match score as the winner. The localization system concludes that the robot is currently at the place corresponding to the winning grid. In addition, the grid-matching procedure outputs the best transformation between the recognition grid and the winning grid; in combination with the stored Cartesian location of the corresponding place, this information can be used to determine the robot's precise location.

4. Experimental Results

We conducted experiments using a real mobile robot in an unmodified office environment, where obstacles included chairs, tables, desks, bookshelves, workstations, bicycles, and a copy machine. Dynamic change was present in both transient and lasting forms. Transient changes were caused by people moving through the environment, during both place learning and place recognition. Lasting changes occurred when people rearranged chairs, added and removed bicycles, and opened and closed doors.

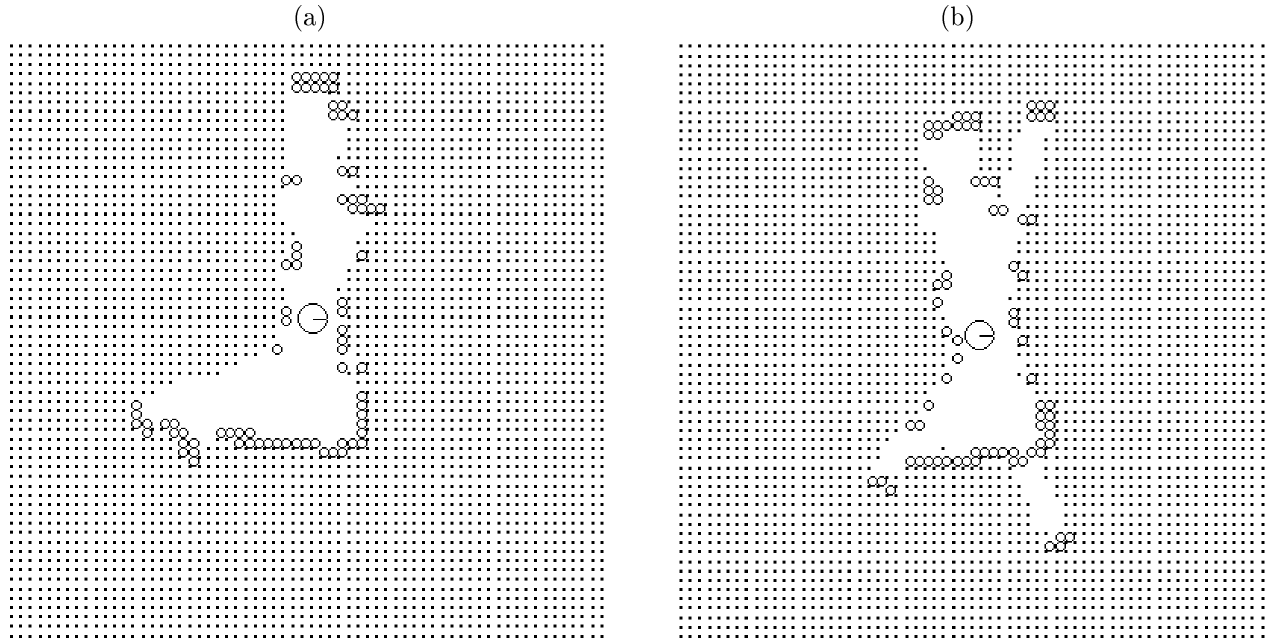

Figure 4: The learned (a) and recognition (b) grids for place 41.

Figure 3 shows the topological/metric map constructed as the robot moved through the environment, including the place locations and the topological links connecting these places. The system learned a total of 47 places, each with an associated evidence grid. The time required to build each evidence grid was approximately thirty seconds, with most of the time spent in sensorimotor control (performing the ten two degree rotations necessary to collect eleven sets of sonar readings and triggering the sixteen sonar sensors at each position). The time required for each place recognition was approximately five minutes (including the time required to build the recognition grid) using a DecStation 3100, with most of the time spent in the grid-matching procedure. We plan to transfer this system to faster hardware in the near future, which should greatly reduce the time required for place recognition.

Figure 4 (a) shows the evidence grid learned for place 41, which corresponds to a corner where two walkways intersect. The robot's position is indicated by the large circle, with the line giving the robot's heading. In all of these figures, the robot is facing due north, as determined by the onboard compass. White space indicates cells with occupancy probabilities less than the prior probability of occupancy, while circles denote cells with occupancy probabilities greater than the prior probability. Dots mark cells with occupancy probability equal to the prior probability (meaning that they are not visible from the robot's position).

Figure 4 (b) shows the recognition grid constructed by the robot at place 41 during testing. The system correctly matched this grid to the learned grid for place 41 despite a number of differences. First, the robot was positioned at slightly different locations, so that its views of the lower corridor extended different lengths. Second, variability in the compass caused the robot's perception of "due north" to be slightly different in the two cases, resulting in shifted orientations. The rotational component of the transformation space searched during grid matching compensates for this angular uncertainty. Third, two specular reflections are present in the recognition grid that were not present in the learned grid. However, these specular reflections were not always present, so by taking multiple readings at different rotation angles, the system could perceive the surfaces from which the sonar beams reflected.

Figure 5 shows the learned grid and recognition grid for a more complex place area. On the left side of this area is a wall containing open doorways leading to offices. On the right side is a large open area containing chairs, desks, and workstations. The clear area in the lower-left corner of this area is actually a (permanent) specular reflection caused by a whiteboard. This surface is sufficiently smooth that it acts as a mirror for the sonar, consistently reflecting all of the beams originating near the center of this area. In this case, these reflections can actually be useful as a distinguishing feature of this place, but only if the place regions are sufficiently small that the angle of reflection is similar during learning and recognition.

People walked past the robot during both place learning and place recognition, but the use of multiple sensor readings allowed the corresponding transient changes to be filtered out of these grids. The chairs

Figure 5: The learned (a) and recognition (b) grids for place 26.

on the right side of the room were rearranged between the times that these two grids were constructed and, in addition, a bicycle (not present in the learned grid) was present in the upper-central region of the area during place recognition. Despite these changes, the localization system was able to correctly match the recognition grid with the learned grid.

In order to measure the effects of larger lasting changes, we removed the whiteboard that was causing the specular reflection in the learned grid, apparent in Figure 5 (a). As a result, the robot detected the wall itself rather than a specular reflection, giving the recognition grid Figure 5 (c). In spite of the substantial difference between the learned grid and the new recognition grid, the place recognition system still successfully identified the robot's location.

Overall, the robot was able to localize itself accurately throughout the environment. In most places the robot localized itself with 100% accuracy, always determining the correct place. In the remaining 10% of the places, the robot localized itself correctly three out of four times. These results are preliminary, but promising, and we plan more precise, quantitative performance measurements in the near future.

5. Related Work

Much research has been conducted on place recognition for mobile robots using a variety of techniques and knowledge structures, including distinctive places (Kuipers & Byun, 1993), gateways (Kortenkamp, 1993), image signatures (Engelson, 1994), and landmarks (Greiner & Isukapalli, 1996). However, none of this research has addressed the issue of place recognition in dynamic environments, where the appearance of places may change over time.

In our previous research, we have described a system that builds evidence grids for different places and uses case-based techniques for place recognition (Langley & Pfleger, 1995). However, we tested this approach only in static environments. We have also conducted research on robot localization in dynamic environments using evidence grids (Yamauchi, 1996), but that work was aimed at correcting dead-reckoning errors using a grid constructed for a single home location. The approach we have described in this paper combines the capability for place learning and place recognition with robustness to dynamic changes, both transient and lasting.

Schiele and Crowley (1994) have done work on position estimation based on matching line segments extracted from evidence grids using Hough transforms and Kalman filtering. However, their research has only dealt with static environments, and it is unclear how robust these techniques would be in dynamic ones.

Schultz and Grefenstette (1995) have reported a method for continuous localization using evidence grids. In their work, local grids constructed by the robot are continuously registered with a global grid to determine the robot's position. Our approach differs in using evidence grid localization as a part of a topological exploration and navigation system, rather than as a means for building a global evidence grid.

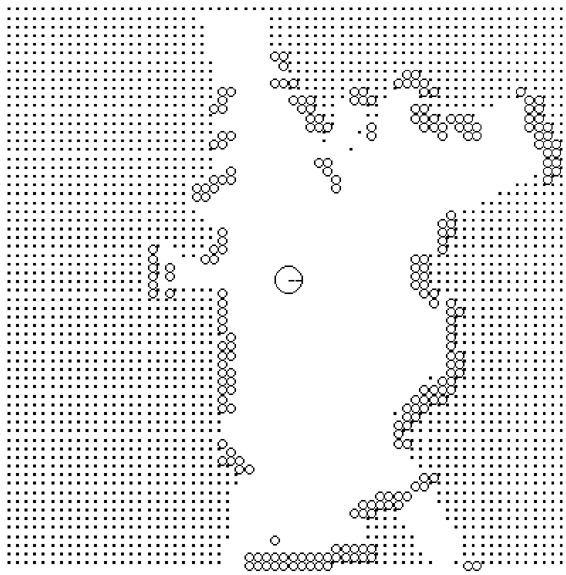

Figure 6: The recognition grid for place 26 after removing the source of specular reflection.

6. Conclusions

In this paper, we have presented a method for place learning and place recognition in dynamic environments. Place learning consists of associating evidence grids with places corresponding to nodes within a topological/metric map. Place recognition consists of building an evidence grid for the robot's current location and using a hill-climbing search to find the best match between this grid and the previously learned grids. By taking multiple sets of sonar readings, transient changes can be filtered out of each grid. By rotating the robot through a small angle between each set of readings, specular reflections can be represented consistently within each grid.

We have implemented this method on a real mobile robot and tested it in a real-world office environment. In these tests, the robot was able to localize itself successfully, despite the presence of people moving past the robot during place learning and place recognition, and despite changes in the arrangement of furniture occurring in the interval between place learning and place recognition. We plan additional experiments to provide a more precise quantitative measure of the system's performance and its robustness to transient and lasting changes.

Acknowledgements

We owe thanks to Hans Moravec for making his evidence grid software available for use, to Alan Schultz and Bill Adams for many useful discussions and for sharing their software, and to Barbara Hayes-Roth and Nils Nilsson for providing access to a Nomadics robot. This research was supported in part by Grant Number N00014-94-1-0505 from the Office of Naval Research.

References

Engelson, S. (1994). *Passive map learning and visual place recognition.* Doctoral dissertation, Department of Computer Science, Yale University, New Haven, CT.

Greiner, R., & Isukapalli, R. (in press). Learning to select useful landmarks. *IEEE Transactions on Systems, Man, and Cybernetics.*

Konolige, K. (1995). A refined method for occupancy grid interpretation. *Proceedings of the International Workshop on Uncertainty in Robotics.* Amsterdam, Netherlands.

Kortenkamp, D. (1993). *Cognitive maps for mobile robots: A representation for mapping and navigation.* Doctoral dissertation, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor.

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

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.

Moravec, H. P. (1988). Sensor fusion in certainty grids for mobile robots. *AI Magazine*, *16*, 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). Saint Louis, MO.

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

Schultz, A., & Grefenstette, J. (1995). *Continuous localization using evidence grids* (Technical Report AIC-95-024). Washington, DC: Navy Center for Applied Research on Artificial Intelligence, Naval Research Laboratory.

Yamauchi, B. (1996). Mobile robot localization in dynamic environments using dead reckoning and evidence grids. *Proceedings of the 1996 IEEE International Conference on Robotics and Automation.* Minneapolis, MN.

Yamauchi, B., & Beer, R. (in press). Spatial learning for navigation in dynamic environments. *IEEE Transactions on Systems, Man, and Cybernetics.*