Case-Based Acquisition of Place Knowledge

Pat Langley* (Langley@cs.stanford.edu)
Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306 USA

Karl Pfleger (KPfleger@hpp.stanford.edu)
Knowledge Systems Laboratory, Computer Science Dept.
Stanford University, Stanford, CA 94305 USA

Abstract

In this paper we define the task of place learning and describe one approach to this problem. The framework represents distinct places using evidence grids, a probabilistic description of occupancy. Place recognition relies on case-based classification, augmented by a registration process to correct for translations. The learning mechanism is also similar to that in case-based systems, involving the simple storage of inferred evidence grids. Experimental studies with both physical and simulated robots suggest that this approach improves place recognition with experience, that it can handle significant sensor noise, and that it scales well to increasing numbers of places. Previous researchers have studied evidence grids and place learning, but they have not combined these two powerful concepts, nor have they used the experimental methods of machine learning to evaluate their methods' abilities.

1. Introduction and Basic Concepts

A physical agent exists in an environment, and knowledge about that environment can aid its achievement of goals. Research on the representation, use, and acquisition of spatial knowledge has occupied an important role in robotics. In this paper, we consider a novel approach to this area that combines ideas from robotics and machine learning.

Let us begin with some definitions of concepts and tasks that appear central to spatial reasoning. Consider a physical agent, say a robot, that is situated in the world. We can say that:

Definition 1 *The* POSITION *of an agent is a coordinate in 2D or 3D space.*

Position corresponds to ground truth, giving the actual location of the agent in some established coordinate system. We can also define the related concept of agent *orientation*, but in this paper we will assume that the agent has a 360 degree field of view, making this notion less central.

A physical agent does not typically have direct access to knowledge of its position, but it does have indirect information.

Definition 2 *A* SENSOR READING *is a description of the environment around the agent's position that has been filtered through its sensors.*

The information in sensor readings may be imperfect in various ways. For example, it may be incomplete in that it describes only certain characteristics of the local environment, and it may be noisy in that sensor readings for the same position may produce different results at different times.

Nevertheless, the agent must find some way to use this information to make useful inferences. This suggests a natural task for a physical agent:

Definition 3 LOCALIZATION *involves the determination of an agent's position in the environment from a set of sensor readings.*

Other tasks, such as navigating from position A to position B, are certainly possible. However, note that an agent cannot begin to carry out such a task without first knowing A and without knowing when it has achieved B. Thus, localization seems more basic than navigation, and we will focus our attention on it here.

However, in many cases humans seem to care less about their exact position in space than about more abstract spatial regions. This suggests another, somewhat different, concept:

Definition 4 *A* PLACE *is a contiguous set of positions in 2D or 3D space.*

Robotics researchers have paid relatively little attention to the notion of place, but its central role in human spatial reasoning suggests that it deserves a closer

*Also affiliated with the Robotics Laboratory, Computer Science Department, Stanford University.

look. Naturally, this new concept lets us define an associated performance task by analogy with the localization task:

Definition 5 PLACE RECOGNITION *involves determining the place in which the agent currently resides from a set of sensor readings.*

At least in principle, the place recognition task seems more tractable than localization, in that it transforms a problem of numeric prediction into one involving discrete classification. One can also carry out localization within the context of a given place, but this in turn may be easier than global localization. Navigation between two places may also be simpler than navigation between two positions.

Of course, reliance on places rather than positions also introduces a problem: one must specify some descriptions in memory that let the agent map sensor readings onto place names. One might attempt to enter such descriptions manually, but it seems desirable to automate this process, suggesting a final task:

Definition 6 PLACE LEARNING *involves the induction of descriptions, from the sensor readings and place names for a set of training positions, that let the agent accurately recognize the places of novel positions.*

Note that this task formulation makes minimal demands on the teacher, who does not have to give the agent information about its actual positions. Rather, the agent collects its own sensor readings, and the teacher must only label each reading as an instance of one place or another. This formulation assumes supervised training data, but unsupervised versions, in which the agent decides on its own place names, are possible as well. We will touch briefly on unsupervised place learning in Section 7, but we will focus on the supervised version in this paper.

In the pages that follow, we present one approach to dealing with knowledge about places. First we describe a representational formalism for storing place knowledge – evidence grids – and then examine a method for place recognition that operates on this representation. After this, we consider a simple learning process that lets one acquire and refine knowledge of places. Next we present some hypotheses about our approach, along with some experimental tests of those hypotheses. Finally, we review related work on spatial learning and discuss some directions for future work.

One important difference between our approach and earlier robotics work on spatial knowledge lies in our incorporation of ideas from machine learning. In particular, we view place recognition as a classification task and we view place learning as a supervised concept induction task. This suggests not only certain learning methods, but also the use of experimental methods prevalent in machine learning to evaluate our technique. However, the tasks of place recognition and learning introduce some difficulties not usually present in machine learning research, in particular the pervasive presence of significant sensor noise. Our approach to representing, using, and learning place knowledge is designed with this in mind.

2. The Evidence Grid Representation

Robotics researchers have explored a variety of formalisms for representing spatial knowledge. One approach relies on using geometric primitives to describe the edges or surfaces of obstacles in the environment. For example, one can use a set of lines to approximate the walls of an office and the furniture it contains. Such representations are precise, but Schiele and Crowley (1994) note that they can be difficult to use when sensors are noisy.

Another common scheme involves dividing the environment into a rectangular grid of mutually exclusive cells, each corresponding to a distinct position in space. In this framework, each cell is specified as either occupied (containing an obstacle) or open (containing none). This approach is well suited to navigation tasks in which one already knows the structure of the environment (i.e., which cells are occupied) and the position of the agent within the grid. However, this scheme is not designed to handle the uncertainty that arises when the position is unknown or when the agent has yet to learn the structure of the environment.

An alternative framework uses the *evidence grid* (Elfes, 1989; Moravec & Blackwell, 1992), a data structure that is specifically designed to deal with uncertainty. In this approach, each cell C has an associated probability that C is occupied by some tangible object that would block the agent's path if it tried to move through the cell. These probabilities range from near zero (nearly certain a cell is open) to near one (nearly certain a cell is occupied), with the middle corresponding to cells for which little information is available (e.g., behind a wall or inside an object). We will adopt this framework in the current paper.

Figure 1 shows the position of an agent in a room within a larger office environment, similar in structure to an actual area at Stanford University. Figure 2 depicts evidence grids constructed from simulated sensor readings taken from a position (a) in the top left room and (b) in the lower left room from the same orientation. Note that open regions within the agent's view have low probability of being occupied (lighter shades) and that edges of obstacles and walls within view have high probability (darker shades). However, areas that are occluded, such as those behind obstacles and walls, have probabilities around $\frac{1}{2}$ (empty regions), since the agent's sensors provide no information about them.

Previous work with evidence grids has emphasized their use in representing single rooms over a relatively

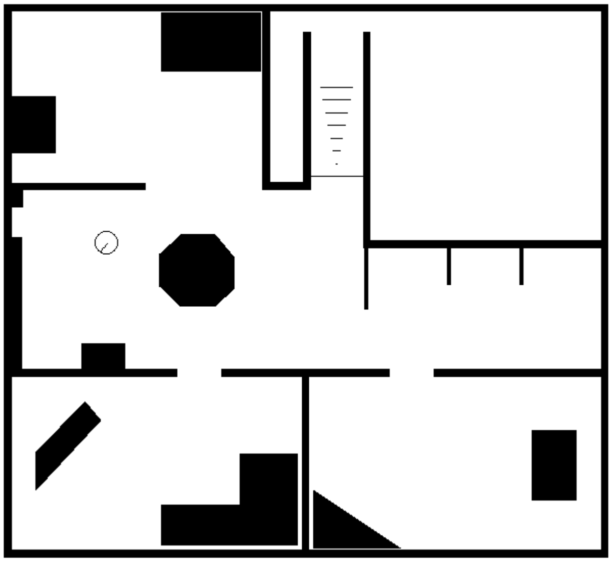

Figure 1: A simulated office environment with a number of distinct places.

short period. However, they also have potential for handling large-scale spatial knowledge over longer time spans. An agent could store its knowledge about an entire building or even a city in a single, large evidence grid. But this scheme seems impractical from a computational viewpoint, and odometry errors could cause grid cells to become increasingly uncertain over time.

A more tractable approach to representing large-scale spatial knowledge, which we take here, involves storage of separate evidence grids for each distinct place. For example, one might use one or a few grids to encode each room in a building. This knowledge can be augmented by geometric relations among places, which would support navigation planning, but we will not address that aspect here. The retention of place descriptions in memory has much in common with the storage of a *case library* in work on case-based reasoning (Aha, Kibler, & Albert, 1991; Kolodner, 1993). In both frameworks, the stored items represent alternative situations in which the agent can find itself, and which suggest different inferences.

3. Case-Based Recognition of Places

Now that we have described the nature of evidence grids, we can examine their use in place recognition. Let us assume the agent has a stored place library, with each place described as an $n \times m$ evidence grid with an associated place name. Our approach to place recognition relies on a three-step process that is closely akin to case-based reasoning.

First, the agent constructs a temporary or short-term evidence grid for its current position from a set of sensor readings. This involves transforming each sensor reading into a probability of occupancy for each cell. Following Moravec and Blackwell (1992), we assume a sensor model that specifies this mapping. The result is an initial temporary evidence grid, based on the sensor reading, that characterizes the region in the vicinity of the agent. The agent repeats this process a specified number of times, in each case incorporating the result into the temporary grid using a Bayesian updating scheme. We will not describe this updating process in depth, but readers can find details in Elfes (1989) and in Moravec and Blackwell (1992).

Next, the agent matches the short-term evidence grid against each of the grids stored in the place library. The evaluation function used in this comparison process measures the degree of match between two grids. Specifically, if $S_{r,c}$ is the probability associated with the rth of R rows and the cth of C columns for the short-term grid, and if $L_{r,c}$ is the analogous probability for the stored, long-term grid, then

$$M = \sum_{r}^{R} \sum_{c}^{C} F(S_{r,c}, L_{r,c})$$

computes the similarity between the short-term and stored grids. One can instantiate the function F in many ways, provided they satisfy certain properties: two cells should be treated as similar if they are confident in the same direction, as dissimilar if they are confident in opposite directions, and generally ignored if either is uncertain.

Moravec and Blackwell (1992) implement this cell-to-cell component of grid similarity $F(S_{r,c}, L_{r,c})$ as

$$log\,[S_{r,c}L_{r,c} + (1 - S_{r,c})(1 - L_{r,c})] + 1 \quad ,$$

which varies from one (a perfect match) to negative infinity (the worst possible match). Reflection suggests that this scheme might give very low match scores to reasonably similar grids if even a few cells are confident in opposite directions. For this reason, we decided to use an alternative definition of $F(S_{r,c}, L_{r,c})$:

$$\begin{array}{rlcclccl} 1 & \text{if} & S_{r,c} & > & \frac{2}{3} & \text{and} & L_{r,c} & > \frac{2}{3} \\ 1 & \text{if} & S_{r,c} & < & -\frac{2}{3} & \text{and} & L_{r,c} & < -\frac{2}{3} \\ -1 & \text{if} & S_{r,c} & > & \frac{2}{3} & \text{and} & L_{r,c} & < -\frac{2}{3} \\ -1 & \text{if} & S_{r,c} & < & -\frac{2}{3} & \text{and} & L_{r,c} & > \frac{2}{3} \\ 0 & \text{otherwise .} & & & & & & \end{array}$$

We felt this measure would be less sensitive to situations in which disagreements arise between cells having high certainty, and initial experiments (Langley & Pfleger, 1995) suggest that our version fares much better than the Moravec/Blackwell measure.

The above metrics assume that the stored and temporary grids are described in the same coordinate system. One can plausibly assume the presence of a reasonably accurate compass to determine the relative rotations, but possible differences in translation require

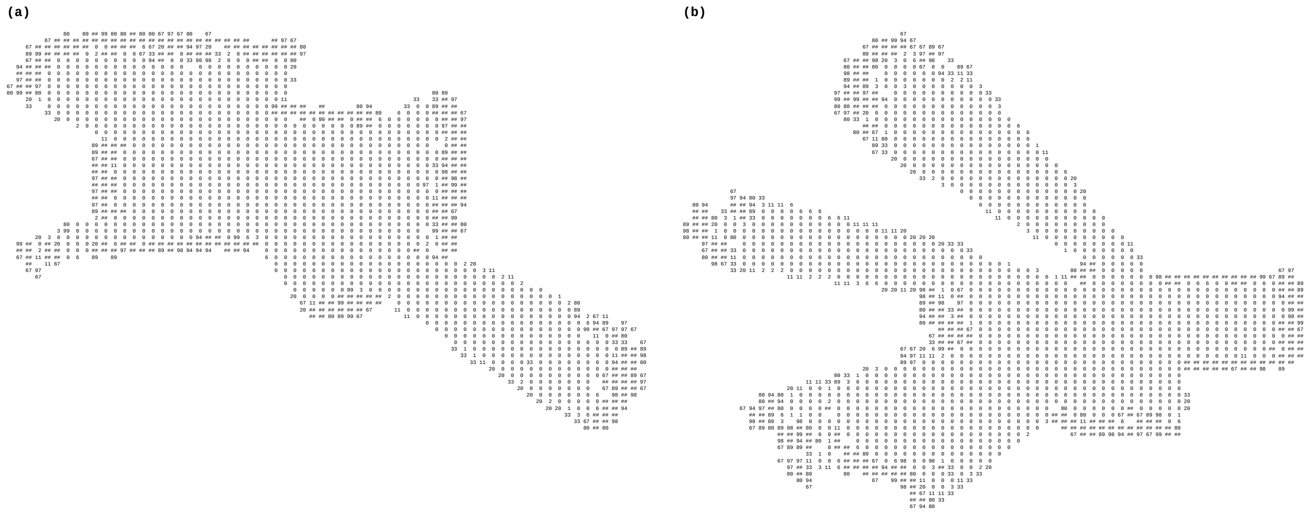

Figure 2: Evidence grids generated from simulated sonar readings for (a) the top left room in Figure 1 and (b) the lower left room in the same figure.

some form of registration that coerces the temporary evidence grid into the same coordinate system as the stored place. To this end, our system carries out an exhaustive search using operators that modify the position by one grid row or column, evaluating each alternative using the metric M defined above[1] and selecting the translation that gives the highest score. If the compass is not accurate, one can extend this approach to correct for small offsets in rotation estimates.

Finally, the agent compares the match scores for the various registered grids and selects the best of these competitors. This strategy provides both the place name associated with the selected evidence grid and the estimated position within that place description. Because adjacent evidence grids may cover overlapping regions, this scheme has some potential for misclassifying a place based on its outlying rather than its central cells. However, this issue has not been a problem in our studies to date.

As we noted earlier, this approach has much in common with methods for case-based reasoning. Here the evidence grids in the place library correspond to stored cases, whereas the short-term grid maps on to the test case for which one wants to make a prediction. The match function corresponds to the similarity metric that determines the nearness of the test case to each stored case in an $R \times C$ dimensional space, and the final classification step is similar to that used in the nearest neighbor method, perhaps the simplest case-based technique. The fact that each evidence grid may be a probabilistic summary, computed from a set of sensor readings, differs from the prototypical case-based system, but abstract cases are not that unusual. A more intriguing difference concerns the registration process. Many case-based systems incorporate some *adaptation* method, but usually this occurs after case selection, whereas here adaptation (registration) takes place during the evaluation (match) process itself.

[1] When translation causes two grids to overlap on only $R' \times C'$ cells, the metric uses only these cells in its summation. This creates a bias toward stored grids that share more cells with the temporary grid, which seems reasonable, but it does not actively punish a stored grid for having only partial overlap.

4. Case-Based Learning of Places

Now let us consider an approach to the acquisition of place knowledge that is stored in evidence grids. We would like an incremental learning process, since the agent encounters its environment sequentially. However, we are not concerned here with the task of effectively exploring an unknown world, so we will assume that the agent is led to a position, given time to observe its surroundings, given a place name, led to another position, and so forth.

Given our commitment to a place library and to a case-based method for place recognition, we naturally assume a case-based learning scheme as well. In particular, at each position to which it is led, the agent collects a specified number of sensor readings and constructs a short-term evidence grid S using the method described in the previous section. The system then simply adds the new grid to the place library, along with the specified place name. The same place name may be associated with multiple evidence grids, but this seems appropriate for places that appear different from different positions.

At first glance, this approach to place learning sounds guaranteed to work, in that one simply stores a description for each place, after which recognition will be perfect. However, this view ignores the central feature of the task – uncertainty. Even with noise-free sensors, the same place typically looks different from different positions, if only because certain regions are occluded. Moreover, standard robotic sensors such as sonar are notoriously noisy, and will produce different sensor readings, and thus different evidence grids, even when repeated from the same position. Thus, the adequacy of this approach remains an open question that is best answered by experiment.

5. Experiments with Place Learning

In Section 1 we formulated the place learning task in terms similar to those used to describe other induction problems. Thus, we can use the experimental methods developed for machine learning to evaluate the robustness of our framework. In this section we present a number of hypotheses about the system's behavior, followed by experimental tests of those hypotheses. Our primary measure of performance is recognition accuracy for places in a test set of evidence grids that differ from those in the training set.

The experiments we designed to evaluate the abilities of our approach relied on both a physical robot – a Nomad 200 with a 16-sensor sonar ring – and a high-fidelity simulation of this machine. The physical environment was a suite of offices and common areas at Stanford University, and the simulated environment was an idealized layout of a similar suite, depicted earlier in Figure 1. We used the physical Nomad to ensure realism in our results, while the simulation gave us experimental control over device parameters not possible with the actual robot.

We generated each training or test case by placing the physical or simulated robot in a position, collecting readings from the sonar ring to construct an initial evidence grid, rotating and/or moving the robot (as described below), collecting new sonar readings and updating the evidence grid, and repeating this process many times. For the simulated robot, we generated six different grids for each of six distinct places,[2] giving 36 total evidence grids. For the physical robot, we produced only three grids for each place (because the process took longer), giving 18 total grids.

The Nomad simulator incorporates a number parameters that affect the quality of sonar information. For example, the `error` parameter controls random variation in the distance returned by the sonar sensors, `critical` controls the angle of incidence at which specular reflection occurs, and `halfcone` controls the angular width of each sonar signal. Unless otherwise specified, we set `error` to 0.15, which was our best estimate of the error encountered by the physical robot, and we left all other parameters at their default values, which produce a 25 degree field of view for each sensor and specular reflection at angles of incidence with the sensed surface of 30 degrees or less.

For each experimental condition with the simulated environment, we ran the learning system 400 times with different random partitions of the evidence grids into 33 training and three test cases, randomly ordering the storage of training cases. For the physical environment, we partitioned the grids into 17 training cases and one test case, again averaging over 400 runs for each condition.

Following Kibler and Langley (1988), we can divide the factors that affect the learner's behavior into two broad types, those involving characteristics of the environment and those involving features of the learner. The most basic environmental characteristic is the number of training cases available. Naturally, we hypothesized that the accuracy of place recognition would improve as the agent encounters more positions. However, the literature sometimes reports actual decreases in performance, so we needed to test this expectation.

Figure 3 (a) shows the learning curve, giving 95% confidence intervals, for the physical Nomad robot when each training and test grid was based on 45 sets of sonar readings, taken from a single position but with successive orientations incremented by one degree. As expected, the system's ability to recognize places gradually increases as it observes and stores more training cases. However, the shape of the curve suggests that the learning task is not trivial, in that multiple cases for each place are needed to achieve even 70% accuracy. The curve has not yet leveled off at 17 instances, so presumably additional cases would further improve recognition.

Figure 3 (a) also shows an analogous learning curve for the simulated robot. The general shape of the curve is very similar to that for the physical device, but the rate of learning is somewhat higher. Although a few errors still occur even after 35 training cases, the performance component generally assigns the correct place name to the test cases.

We were interested not only in our method's ability to recognize places, but in its ability to identify the precise position of the robot within a given place. Thus, we also measured the absolute difference between the actual robot position in each test case and the estimated position as computed during the registration process. Figure 3 (b) shows the learning curves for the physical robot, as well as similar results for the simulated one. In this case, since we are measuring error rather than accuracy, the quantities start high and gradually decrease with experience. Again, behavior in the simulated environment generally mimics that in the physical world, though the system fares somewhat better on the former.

[2] We defined these places so they corresponded to the lower left, lower right, middle right, and upper left rooms in Figure 1, and to the areas to the left and right of the octagonal table in the figure.

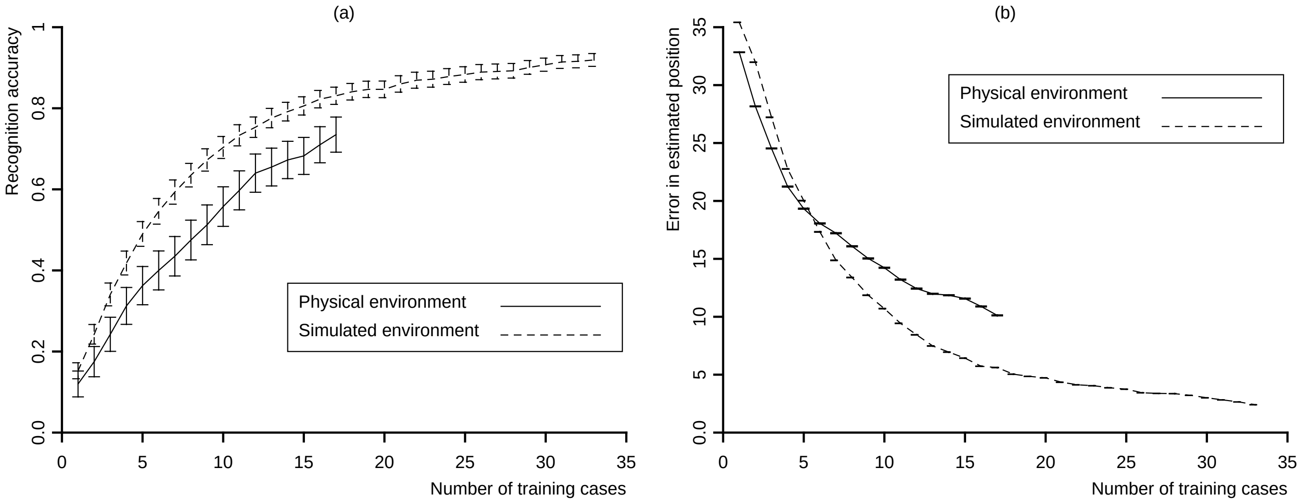

Figure 3: Learning curves with 95% confidence intervals for the case-based place learning system for a physical Nomad robot and a simulated robot in a similar office environment, (a) using recognition accuracy as the performance measure and (b) using error in estimated position.

The amount of sensor noise constitutes a more interesting environmental factor. We would not expect increased noise to affect the asymptotic accuracy, but it should decrease the rate of place learning, that is, the number of training cases needed to reach a given accuracy level. Fortunately, our reliance on evidence grids suggested a natural response to noisy sense data. Because each stored case can be based on multiple sensor signals, we can attempt to improve the *quality* of these cases by increasing the number of the signals used to generate them. We hypothesized that place descriptions based on more sensor readings would be less affected by increases in sensor noise. Thus, we predicted an interaction between these two independent variables, specifically one that affects learning rate but not asymptotic accuracy.

To test this hypothesis, we used the Nomad simulator to produce four different levels of sensor noise, in which the `error` parameter was set to 0.0, 0.15, 0.30, and 0.45, respectively. We also attempted to vary the quality of the stored cases by using two different sensing strategies. In one, we based each evidence grid (both training and test cases) on 45 sonar readings collected from a single position but produced at orientations one degree apart, as used to generate the results in Figure 3. In the other, we based each grid on 90 readings, produced by repeating this strategy in two nearby, randomly selected positions within the same room.

Figure 4 (a) shows the learning curves that result for the zero and 0.45 noise levels using the one-position sensing strategy, whereas Figure 4 (b) presents analogous results for the two-position strategy. (The results for the 0.15 and 0.3 settings fell between these extremes; we have omitted them for the sake of clarity.) The two-position scheme clearly fares better than the simpler strategy, but the curves diverge somewhat from our predictions. The rate of learning for the two-position method is much higher than for the one-position method, even when no sensor noise is present. Also, the introduction of sensor noise clearly affects both strategies, but it alters only the learning rate for the more sophisticated scheme, while it actually appears to reduce the asymptotic recognition accuracy for the simpler one.

Some real-world environments contain many distinct places, and we hypothesized that our learning method would scale well as the number of places increased. We obtained preliminary results along these lines by examining our algorithm's behavior with different subsets of the places available in our environment. Figure 5 (a) shows the learning curves that result for two through six places, with each case based on 45 simulated sonar readings from one position. Each reported accuracy is averaged over 400 runs for each possible subset of k out of six places, using 35 randomly selected training cases and one test case. Thus, when $k = 2$ we carried out $\binom{6}{2} \times 400 = 6000$ runs, and when $k = 3$ we carried out $\binom{6}{3} \times 400 = 8000$ runs. We have not reported confidence intervals here, since the accuracies are averages of averages.

Naturally, increasing the number of places decreases the speed of learning, but we can also examine the rate of this decrease. Note that the figure also shows where each learning curve crosses the level of 90 percent accuracy. These crossover points produce the *scaling* curve in Figure 5 (b), which maps the number of distinct places against the number of training cases needed to reach this accuracy level.

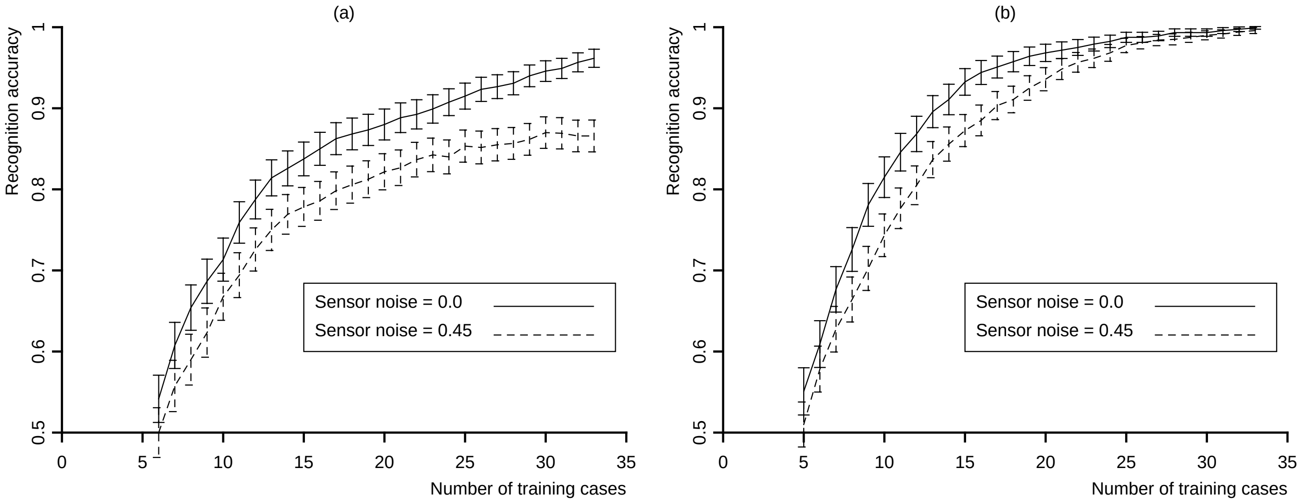

Figure 4: Learning curves for the case-based place learning system for two levels of sensor noise when evidence grids are based on (a) 45 readings from one position and (b) 90 readings from two nearby positions.

This higher-order curve seems to be either linear or quadratic, but the analogous scaling curve for the two-position sensing strategy definitely appears linear. These results suggest that our approach requires, more or less, a fixed number of training cases per place, independent of the total number of places. This encourages us to believe that the method will scale well to domains that involve many more different places than the six we have examined, though ultimately we should test this prediction using larger environments.

Clearly, there exist many other factors that could influence the behavior of our place-learning method. These include the resolution of the evidence grids, the distinctiveness of the places to be learned, and the complexity of these places in terms of the number of separate grids needed to describe them adequately. However, we will reserve these issues for future studies, as the current experiments have been sufficient to show that our approach is a promising one.

6. Related Work on Spatial Learning

Our research on the acquisition of spatial knowledge is certainly not the first in this area. Clearly, our work owes a strong intellectual debt to Elfes (1989), Moravec and Blackwell (1992), and other developers of the evidence grid framework. Our basic representation and our performance system directly employ techniques developed by these researchers. However, most research in this framework has focused on the construction of a single global map, rather than a collection of evidence grids for distinct places. Although such approaches clearly acquire spatial knowledge, they do not involve induction in the sense of using training instances to improve performance on novel test cases, whereas our work on place learning fits easily into this paradigm. Thrun (1993) has used reinforcement learning to improve sensor interpretation for evidence-grid construction, but his goal was to construct a global map. Mahadevan (1992) describes a method that forms generalizations expressed as evidence grids, but his aim was to learn not places but action models.

Nevertheless, some work outside the evidence grid formalism has focused on place learning. For example, Yamauchi and Beer (1994) describe ELDEN, a system that represents places in terms of means and variances of direct sensor readings, rather than inferred grid occupancies. Their place descriptions also include features for the robot's position as estimated through dead reckoning and connections to recently visited places. Place recognition involves passing each attribute's value through Gaussian functions associated with each place, then selecting the competitor with the highest sum. Learning consists of updating the means and variances for recognized places, creating new places when no existing ones match well enough, and adding predictive connections between places. Yamauchi and Beer's reliance on a Gaussian distance metric makes their method similar to our case-based approach, though ELDEN differs in its use of instance averaging, its use of raw sensor data, and the unsupervised nature of the learning processes.[3]

Lin, Hanson, and Judd (1994) have taken a similar approach to representing and using spatial knowledge. Their system also describes places (which they call *landmarks*) as means and variances of sonar readings and uses a Gaussian metric to determine the degree of match against the current sensor signals. However, their learning mechanisms include not only the cre-

[3] Yamauchi and Beer's system also incorporates an evidence grid representation, but it constructs a global map and uses this map for correcting errors in dead reckoning rather than for place recognition.

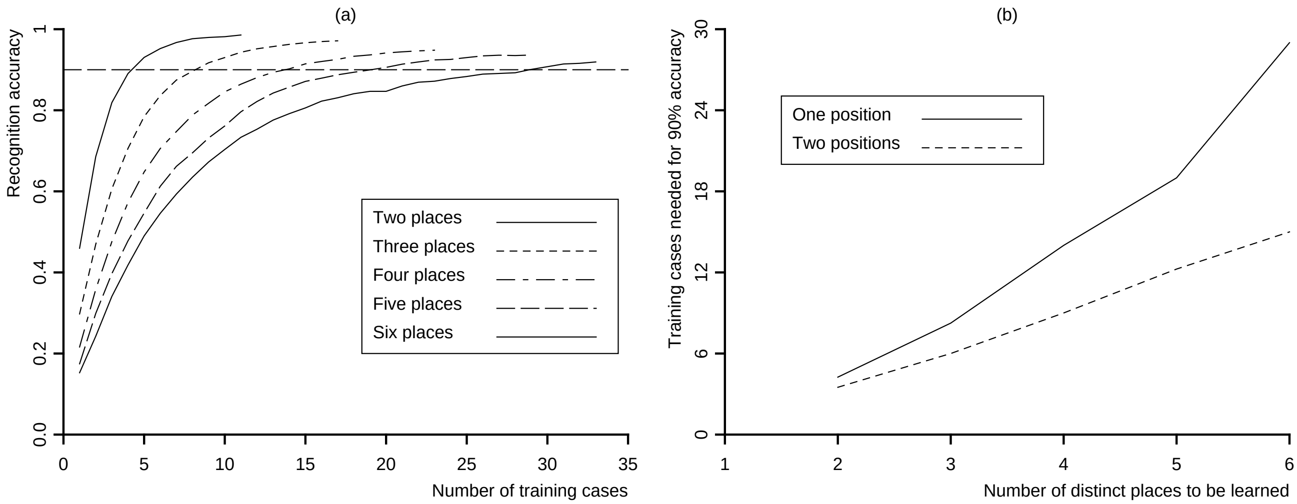

Figure 5: (a) Learning curves for different numbers of distinct places, based on 45 sonar readings from one position. (b) Scaling curves that map, for two different sensing strategies, the number of training cases needed to achieve 90% accuracy as a function of the number of places.

ation and updating of place descriptions, but also a reinforcement process designed to improve estimates of the robot's location. This latter technique can lead the learner to add a new place or remove an existing one if these actions reduce errors in location estimates.

Kuipers and Byun's (1988) NX system also operates on direct sensory readings, but it stores only places that are distinctive in terms of optimizing certain measures. For example, NX defines the central point in a hallway corner as being symmetrical and being equidistant from the walls, in addition to containing information about the angles and distances to obstacles. The system also describes edges, which connect distinctive places, in terms of length, width, and similar characteristics. Whenever NX encounters a local optimum L on one of its measures, it compares the sensor readings to each known place P stored in memory; if the descriptions for L and P are similar, and if their locations are metrically or topologically close, the system classifies L as the known place P. Otherwise, NX creates a new place based on L's description and stores this in memory, along with its connections to other places. Mataric (1991) describes a similar scheme, though the details of place creation are different.

In methodological terms, Kortencamp and Weymouth's (1994) work is perhaps the most similar to our own. Their approach emphasizes *gateways* such as doors that connect two regions, but their system represents these locations using a grid structure and they evaluate its behavior in terms of recognition accuracy. However, their scheme uses hand-coded descriptions for a few gateway types to recognize candidate places and create new ones, rather than actual supervised training data, and they compare a number of different recognition strategies, including one that combines evidence from sonar and visual sensors.

On another dimension, our approach is most similar to Yeap's (1988) work on spatial reasoning. His framework also posits the storage of distinct places, the descriptions of which are not direct sensory readings but inferred summaries. However, his "absolute space representation" does not take the form of evidence grids but rather consists of a connected sequence of line segments that, except for occasional openings, enclose an area. Yeap does not describe a performance element that uses these descriptions in place recognition but, as in our own framework, learning involves the simple storage of the inferred place descriptions, which suggests the use of a case-based method.

7. Directions for Future Work

Although our experimental studies of place learning have revealed some insight into our approach, clearly more work remains to be done. The most immediate extension would replace the current supervised learning method with an unsupervised one that can identify distinctive places on its own. To this end, we plan to employ a technique similar to that used by Yamauchi and Beer (1994), but adapted to operate on evidence grids rather than direct sensor descriptions. As the agent moves through the environment, it would regularly stop and construct a short-term evidence grid, merging this with the previous place description if the match was high enough and using the short-term grid as the basis for a new place otherwise. Discontinuities caused by passage through doors and past obstacles should be enough to identify distinguishable places.

Unlike the current system, this approach could produce exactly one evidence grid for each place.

Most methods for place learning, including those discussed above, also construct topological maps that connect different places. Clearly, this is another important direction in which to extend our approach. We expect that storing rough estimates of the direction of movement between one place and its successor will be sufficient for many navigation tasks. Upon executing a navigation plan, the agent would still need to register its location upon entering each place along the path, but expectations about the next place and its rough translation should greatly increase the speed of the registration process.[4] Storing recently visited places with each grid could also aid recognition in domains with perceptually similar places.

In future work, we also hope to develop methods for recognizing places that change over time, as occurs in rooms with moveable furniture. We plan to store two levels of evidence grid, one that averages over all encounters with a given place and another that describes individual encounters. The former generalized grids would contain fewer cells with near-certain scores, but they could be used to generate priors for known places before they are entered, and also used to focus attentional resources toward regions of uncertainty.

In summary, we have presented a framework for representing, using, and learning knowledge about places in which evidence grids play a central role. Our approach draws on earlier work for updating these probabilistic summaries, but diverges from previous schemes by storing a set of local grids in a case library, then retrieving and matching them for use in place recognition. Experimental studies adapted from the machine learning literature indicate that this approach improves recognition accuracy with experience, that the quality of stored cases can offset the effects of sensor noise, and that the method scales well to increased numbers of places. Many other environmental and system factors remain to be examined, but the basic approach appears promising and suggests many natural extensions.

Acknowledgements

We owe thanks to Hans Moravec for making his evidence grid software available for use, and to Brian Yamauchi for integrating that code with the Nomad 200 interface. Alan Schultz, Brian Yamauchi, and Bill Adams made many useful suggestions, and Barbara Hayes-Roth and Nils Nilsson provided access to a Nomadics robot. This research was supported in part by Grant Number N00014-94-1-0505 from the Office of Naval Research.

[4] We also plan to extend the registration process to handle minor errors in orientation estimates due to compass imperfections.

References

Aha, D. W., Kibler, D., & Albert, M. K. (1991). Instance-based learning algorithms. *Machine Learning*, *6*, 37–66.

Elfes, A. (1989). Using occupancy grids for mobile robot perception and navigation. *IEEE Computer Magazine*, June, 46–58.

Kibler, D., & Langley, P. (1988). Machine learning as an experimental science. *Proceedings of the Third European Working Session on Learning* (pp. 81–92). Glasgow: Pittman.

Kolodner, J. L. (1993). *Case-based reasoning*. San Mateo, CA: Morgan Kaufmann.

Kortencamp, D., & Weymouth, T. (1994). Topological mapping for mobile robots using a combination of sonar and vision sensing. *Proceedings of the Twelfth National Conference on Artificial Intelligence* (pp. 979–984). Seattle, WA: AAAI Press.

Kuipers, B., & Byun, Y. T. (1988). *Proceedings of the Eighth National Conference on Artificial Intelligence* (pp. 774–779). St. Paul, MN.

Langley, P., & Pfleger, K. (1995). *Acquisition of place knowledge through case-based learning* (Technical Report 95-2). Palo Alto, CA: Institute for the Study of Learning and Expertise.

Lin, L., Hanson, S. J., & Judd, J. S. (1994). *On-line learning for landmark-based navigation* (Tech. Report No. SCR-94-TR-472). Princeton, NJ: Siemens Corporate Research, Learning Systems Department.

Mahadevan, S. (1992). Enhancing transfer in reinforcement learning by building stochastic models of robot actions. *Proceedings of the Ninth International Conference on Machine Learning* (pp. 290–299). Aberdeen: Morgan Kaufmann.

Mataric, M. J. (1991). Behavioral synergy without explicit integration. *Sigart Bulletin*, *2*, 130–133.

Moravec, H., & Blackwell, M. (1992). Learning sensor models for evidence grids. *Robotics Institute Research Review*. Pittsburgh, PA: Carnegie Mellon University.

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

Thrun, S. B. (1993). Exploration and model building in mobile robot domains. *Proceedings of the IEEE International Conference on Neural Networks* (pp. 175–180). San Francisco: IEEE.

Yamauchi, B., & Beer, R. (1994). *Exploration and spatial learning in dynamic real-world environments* (Technical Report No. CES-95-01). Cleveland, OH: Case Western Reserve University, Department of Computer Engineering and Science.

Yeap, Y. K. (1988). Towards a computational theory of cognitive maps. *Artificial Intelligence*, *34*, 297–360.