Coping With Uncertainty in Map Learning

Kenneth Basye **Thomas Dean**[*] **Jeffrey Scott Vitter**[†]
Department of Computer Science, Brown University
Box 1910, Providence, RI 02912

Abstract

In many applications in mobile robotics, it is important for a robot to explore its environment in order to construct a representation of space useful for guiding movement. We refer to such a representation as a *map*, and the process of constructing a map from a set of measurements as *map learning*. In this paper, we develop a framework for describing map-learning problems in which the measurements taken by the robot are subject to known errors. We investigate two approaches to learning maps under such conditions: one based on Valiant's *probably approximately correct* learning model, and a second based on Rivest & Sloan's *reliable and probably nearly almost always useful* learning model. Both methods deal with the problem of accumulated error in combining local measurements to make global inferences. In the first approach, the effects of accumulated error are eliminated by the use of reliable and probably useful methods for discerning the local properties of space. In the second, the effects of accumulated error are reduced to acceptable levels by repeated exploration of the area to be learned. Finally, we suggest some insights into why certain existing techniques for map learning perform as well as they do.

1 Introduction

Many of the problems faced by robots navigating in the environment can be facilitated by using expectations in the form of explicit models of objects and the spaces that they occupy. We use the term *map* to refer to any model of large-scale space used for purposes of navigation. *Map learning* involves exploring the environment, making observations, and then using the observations to construct a map. The construction of useful maps is complicated by the fact that observations involving the position, orientation, and identification of spatially remote objects are invariably error prone. In this paper, we explore a number of problems involved in constructing useful maps from measurements taken with sensors subject to known errors.

In previous work [Dean, 1988], we have looked at various optimization problems related to constructing maps (*e.g.,* construct the most accurate map consistent with a set of measurements). Even in cases involving only a single dimension, such optimization problems can turn out to be NP-hard [Yemini, 1979]. In this paper, rather than look at problems that involve doing the best with what you have, we consider problems that involve going out and getting what you need to generate useful representations. In particular, we consider a form of *reliable and probably almost always useful* learning [Rivest and Sloan, 1988] in which the robot gathers information to ensure that it nearly always (with probability $1-\delta$) can provide a guaranteed perfect path from one location to another. A prerequisite to this sort of learning is that the robot, in moving around in its environment, can discern the local properties of space with absolute certainty with high probability having expended an amount of effort polynomial in $\frac{1}{\delta}$ and n, where n is some measure of the size of the environment.

By eliminating local uncertainty, small errors incurred in making local measurements are not allowed to propagate rendering global queries unacceptably inaccurate. In general, local uncertainty accumulates as the product of the distance in generating global estimates. One way to avoid this sort of accumulation is to establish strategies such that the robot can discern properties of its environment with certainty. Most existing map learning schemes exploit this sort of certainty in one way or another (see Section 4). The rehearsal strategies of Kuipers [1988] are one example of how a robot might plan to eliminate uncertainty. Once we have a method for eliminating uncertainty, the problem then reduces to one of planning out and executing the necessary experiments to extract certain information about the environment.

In situations in which it is not possible to eliminate local uncertainty completely, it is still possible to reduce

[*]This work was supported in part by the National Science Foundation under grant IRI-8612644 and by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under Contract No. F49620-88-C-0132.

[†]This work was supported in part by a National Science Foundation Presidential Young Investigator Award CCR–8846714 with matching funds from IBM, and by National Science Foundation research grant CCR–8403613.

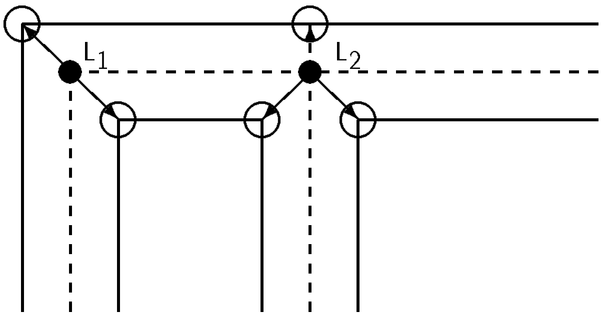

Figure 1: Identifying distinguished locations

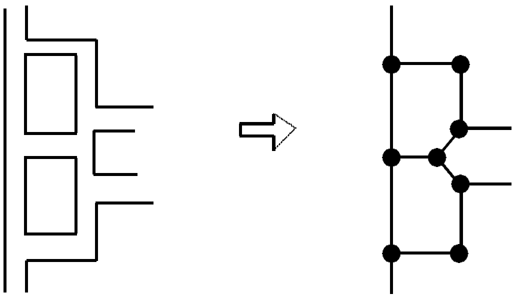

Figure 2: The induced graph of a building

the effects of accumulated errors to acceptable levels by performing repeated experiments. To support this claim, we describe a map-learning technique based on Valiant's *probably approximately correct* learning model [Valiant, 1984] that, given small $\delta > 0$, constructs a map to answer global queries such that the answer provided in response to any given query is correct with probability $1 - \delta$. The techniques presented apply to a wide range of map-learning problems of which the specific problems addressed in this paper are meant to be merely illustrative.

2 Spatial Representation

We model the world, for the purposes of studying map learning, as a graph with labels on the edges at each vertex. In practice, a graph will be induced from a set of measurements by identifying a set of distinctive locations in the world, and by noting their connectivity. For example, we might model a city by considering intersections of streets to be distinguished locations, and this will induce a grid-like graph. Kuipers [1988] develops a mapping based on locations distinguished by sensed features like those found in buildings (see Figure 1). Figure 2 shows a portion of a building and the graph that might be induced from it. Levitt [1987] develops a mapping based on locations in the world distinguished by the visibility of landmarks at a distance.

In general, different mappings result in graphs with different characteristics, but there are some properties common to most mappings. For example, if the mapping is built for the purpose of navigating on a surface, the graph induced will almost certainly be planar and cyclic. Other properties may include regularity or bounded degree. In what follows, we will always assume that the graphs induced are connected and undirected; any other properties will be explicitly noted.

Following [Aleliunas *et al.*, 1979], a graph model consists of a graph, $G = (V, E)$, a set L of labels, and a labeling, $\phi : \{V \times E\} \rightarrow L$, where we may assume that L has a null element $\perp$ which is the label of any pair $(v \in V, e \in E)$ where e is not an edge from v. We will frequently use the word *direction* to refer to an edge and its associated label from a given vertex. With this notation, we can describe a path in the graph as a sequence of labels indicating the edges to be taken at each vertex. We can describe a procedure to follow as a function from $V \rightarrow L$ indicating the preferred direction at each location.

If the graph is a regular tessellation, we may assume that the labeling of the edges at each vertex is consistent, *i.e.,* there is a global scheme for labeling the edges and the labels conform to this scheme at every vertex. For example, in a grid tessellation, it is natural to label the edges at each vertex as North, South, East, and West. In general, we do not require a labeling scheme that is globally consistent. You can think of the labels on edges emanating from a given vertex as local directions. Such local directions might correspond to the robot having a compass that is locally consistent but globally inaccurate, or local directions might correspond to locally distinctive features visible from intersections in learning the map of a city.

In the following, we identify three sources of uncertainty in map learning. First, there may be uncertainty in the movement of the robot. In particular, the robot may occasionally move in an unintended direction. We refer to this as *directional* uncertainty, and we model this type of uncertainty by introducing a probabilistic movement function from $\{V \times L\} \rightarrow V$. The intuition behind this function is that for any location, one may specify a desired edge to traverse, and the function gives the location reached when the move is executed. For example, if G is a grid with the labeling given above, and we associate the vertices of G with points (i, j) in the plane, we might define a movement function as follows:

$$\psi(i,j,l) = \begin{cases} (i, j+1) & 70\% \text{ of the time if } (l = North) \\ (i+1, j) & 10\% \text{ of the time if } (l = North) \\ (i-1, j) & 10\% \text{ of the time if } (l = North) \\ (i, j-1) & 10\% \text{ of the time if } (l = North) \\ \ldots & \end{cases}$$

where the "..." indicate the distribution governing movement in the other three directions. The probabilities associated with each direction sum to 1. If all directions are equally likely regardless of the intended direction, then the movement function is said to be *random*. Throughout this paper, we will assume that movement in the intended direction takes place with probability better than chance.

A second source of uncertainty involves sensors, and in particular recognizing locations that have been seen before. The robot's sensors have some error, and this can cause error in the recognition of places previously visited; the robot might either fail to recognize some previously visited location, or it might err by mistaking some new location for one seen in the past. We refer to this type of uncertainty as *recognition* uncertainty, and model it by partitioning the set of vertices into equivalence classes. We assume that the robot is unable to distinguish between elements of a given class using only its sensors.

A third source of error involves another manifestation of sensor error. In representing the world using a graph, some mapping must be established from a set of distinguished locations in the world to V. Error in the sensors could cause the robot to fail to notice a distinguished location some of the time. For example, a robot taxi might use intersections as distinguished locations, leading to a grid-like graph. But if sensor error causes the robot not to notice that he is passing through an intersection, his map will become flawed. In exploring an office environment, the point in a hallway in front of a door may correspond to a vertex in the induced graph. If the door is closed, there is some chance that the robot will not recognize the vertex in traversing the hall. We model this type of uncertainty by introducing a probabilistic movement function that can skip over vertices. We refer to this type of movement function as *discontinuous* and to the type of uncertainty modeled as *continuity* uncertainty.

Apparently, the three types of uncertainty described above are orthogonal in the sense that none implies or precludes the others. The issues involved in modeling and reasoning about continuity uncertainty are complex and will not be treated further in this paper. In the following, we are concerned with directional and recognition uncertainty.

3 Map Learning

For our purposes, a map is a data structure that facilitates queries concerning connectivity, both local and global. Answers to queries involving global connectivity will generally rely on information concerning local connectivity, and hence we regard the fundamental unit of information to be a connection between two nearby locations (*i.e.,* an edge between two vertices in the induced undirected graph). We say that a graph has been *learned completely* if for every location we know all of its neighbors and the directions in which they lie (*i.e.,* we know every triple of the form $\langle u, l, v\rangle$ where u and v are vertices and l is the label at u of an edge in G from u to v). We assume that the information used to construct the map will come from exploring the environment, and we identify two different procedures involved in learning maps: *exploration* and *assimilation*. Exploration involves moving about in the world gathering information, and assimilation involves using that information to construct a useful representation of space. Exploration and assimilation are generally handled in parallel, with assimilation performed incrementally as new information becomes available during exploration. In this section, we are concerned with the conditions under which a graph can be completely learned, and how much time will be required for the exploration and assimilation.

3.1 Tessellation Graphs

It's not hard to see that any connected, undirected graph can be completely learned easily if there is no uncertainty; [Kuipers and Byun, 1988] describes a way of doing this by building up an agenda consisting of unexplored paths leading out of locations and then moving about so as to eventually explore all such paths. Nothing about the graph need be known before the exploration begins. Introducing the kinds of uncertainty described in Section 2 complicates things considerably. If, however, the graph has additional structure, then that structure can often be exploited to eliminate uncertainty. In the following, we sketch a proof that it is possible to efficiently learn maps that correspond to regular tessellations with boundaries. It turns out that the exploration component of learning regular tessellations is quite simple; random walks suffice for polynomial-time performance. In the longer version of this paper, we describe an efficient incremental assimilation procedure that is called whenever the robot encounters a location during exploration, and then prove the following[1].

Lemma 1 *The assimilation algorithm provided will learn a finite tessellation completely if the exploration tour traverses every edge in the graph. The overall cost of assimilation is $O(m)$ where m is the length of the tour.*

We now have to ensure that during exploration the robot traverses each edge in the graph at least once with high probability. The following two lemmas establish that, for any connected, regular, undirected graph G and any $\delta > 0$, a random walk of length polynomial in $\frac{1}{\delta}$ and the size of G is sufficient for traversing every edge in G with probability $1 - \delta$.

Lemma 2 *For any $d > 1$, there exists a polynomial $p(d, \frac{1}{\delta})$ of order $O(d \log \frac{d}{\delta})$ such that with probability $1-\delta$, p visits to a vertex of order d result in traversing all edges out of the vertex at least once.*

Lemma 3 *For any connected, regular, undirected graph $G = (V, E)$ with order d, any $\delta > 0$, and any $m \geq 1$, there exists a polynomial $p(|E|, m, \frac{1}{\delta})$ such that with probability $1 - \delta$, a random tour on G of length p visits every vertex in V at least m times.*

In most cases, we can do better than random exploration. If the robot moves in the direction it is pointing with probability better than chance, then the robot can traverse every edge in the graph with high probability in time linear in the size of the graph. Using the above three lemmas it is easy to prove the following.

Theorem 1 *Any finite regular tessellation $G = (V, E)$ can be reliably, probably almost always usefully learned.*

The lemmas and form of the proof described above provide a framework for proving that other kinds of graphs can be reliably probably almost always usefully learned in a polynomial number of steps. In general, all

[1] To meet the submission length requirements, all proofs have been omitted. The longer version of the paper, including all proofs[Basye *et al.*, 1989], is available upon request.

we require is that a polynomial number of visits to every vertex provides enough information to learn the graph. Perhaps, the most important lesson to extract from this exercise is that the effects of multiplicative error in learning maps of large-scale space can be eliminated if there is a reliable method for eliminating local uncertainty that works with high probability. The above approach to map learning was inspired by Rivest's model of learning [Rivest and Sloan, 1988], in which complex problems are broken down into simple subproblems that can be learned independently. In order to learn a useful representation of the global structure of its environment, it is sufficient that a robot have reliable and usually effective methods for sensing the local structure of its environment and a method for composing the local structure to generate an accurate global structure. The sensing methods need not always provide useful answers; they need only guarantee that the answer returned is not wrong. The problem then becomes largely one of determining a sequence of sensing and movement tasks that will provide useful answers with high probability. There are situations, however, in which reliable sensing methods are not available, and it is still possible to learn useful maps of large-scale space.

3.2 General Graphs

The next problem we look at involves both recognition and directional uncertainty with general undirected graphs. We show that a form of Valiant's probably approximately correct learning is possible when applied to learning maps. In this section, we consider the case in which movement in the intended direction takes place with probability better than chance, and that, upon entering a vertex, the robot knows with certainty the local name of the edge upon which it entered. We call the latter requirement *reverse movement certainty.* Results for related models are summarized in the next section.

At any point in time, the robot is facing in a direction defined by the label of a particular edge/vertex pair—the vertex being the location of the robot and the edge being one of the edges emanating from that vertex. We assume that the robot can turn to face in the direction of any of the edges emanating from the robot's location. We also assume that upon entering a vertex the robot can determine with certainty the direction in which it entered. Directional uncertainty arises when the robot attempts to move in the direction it is pointing. Let $\gamma > 0.5$ be the probability that the robot moves in the direction it is currently pointing. More than 50% of the time, the robot ends up at the other end of the edge defining its current direction, but some percentage of the time it ends up at the other end of some other edge emanating from its starting vertex. While the robot won't know that it has ended up at some unintended location, it will know the direction to follow in trying to return to its previous location.

To model recognition uncertainty, we assume that the vertices V are partitioned into two sets, the distinguishable vertices D and the indistinguishable vertices I. We are able to distinguish only vertices in D. We refer to the vertices in D as *landmarks* and to the graph as a *landmark graph.* We define the *landmark distribution parameter*, r, to be the maximum distance from any vertex in I to its nearest landmark (if $r = 0$, then I is empty and all vertices are landmarks). We say that a procedure learns the *local connectivity within radius* r of some $v \in D$ if it can provide the shortest path between v and any other vertex in D within a radius r of v. We say that a procedure learns the *global connectivity of a graph* G *within a constant factor* if, for any two vertices u and v in D, it can provide a path between u and v whose length is within a constant factor of the length of the shortest path between u and v in G.

We begin by showing that the multiplicative error incurred in trying to answer global path queries can be kept low if the local error can be kept low, that the transition from a local uncertainty measure to a global uncertainty measure does not increase the complexity by more than a polynomial factor, and that it is possible to build a procedure that directs exploration and map building so as to answer global path queries that are accurate and within a small constant factor of optimal with high probability.

Lemma 4 *Let G be a landmark graph with distribution parameter r, and let c be some integer > 2. Given a procedure that, for any $\delta_l > 0$, learns the local connectivity within cr of any landmark in G in time polynomial in $\frac{1}{\delta_l}$ with probability $1 - \delta_l$, there is a procedure that learns the global connectivity of G with probability $1 - \delta_g$ for any $\delta_g > 0$ in time polynomial in $\frac{1}{\delta_g}$ and the size of the graph. Any global path returned as a result will be at most $\frac{c}{c-2}$ times the length of the optimal path.*

The procedure presented in the proof of Lemma 4 searches outward from a vertex $v \in D$ to a distance cr, and then uses the edges found while entering vertices on the outward path to attempt to return to v. The directions used on the way out form an expectation for the labels observed on the way back. When these expectations are not met, the traversal is said to have failed, and the procedure tries again. The procedure keeps track of the edge/vertex labels associated with vertices visited during exploration in order to ensure that it explores all paths of length cr or less emanating from each vertex in D with high probability.

There is a possibility that some combination of movement errors could result in false positive or false negative tests. But we show by exploiting reverse certainty that we can statistically distinguish between the true and false test results. By attempting enough traversals, the procedure can ensure with high probability that the most frequently occurring sets of directions corresponding to perceived traversals actually correspond to paths in G. What is required, then, is for the learning procedure to do enough exploration to identify all paths of length cr or less in G with high probability.

Lemma 5 *There exists a procedure that, for any $\delta_l > 0$, learns the local connectivity within cr of a vertex in any landmark graph with probability $1-\delta_l$ in time polynomial in $\frac{1}{\delta_l}$, $\frac{1}{1-2\gamma}$ and the size of G, and exponential in r.*

Theorem 2 *It is possible to learn the global connectivity of any landmark graph with probability* $1-\delta$ *in time polynomial in* $\frac{1}{\delta}$, $\frac{1}{1-2\gamma}$, *and the size of* G, *and exponential in* r.

Theorem 2 is a simple consequence of Lemma 4 and 5. It has an immediate application to the problem of learning the global connectivity of a graph where all the vertices are landmarks. In this case, the parameter $r = 0$, and we need only explore paths of length 1 in order to establish the global connectivity of the graph. This process works even if there is no reverse certainty.

Corollary 1 *It is possible to learn the connectivity of a graph* G *with only distinguishable locations with probability* $1-\delta$ *in time polynomial in* $\frac{1}{\delta}$, $\frac{1}{1-2\gamma}$, *and the size of* G, *even if there is reverse uncertainty.*

Given the notion of global connectivity defined above, no attempt is made to *completely learn* the graph (*i.e.*, to recover the structure of the entire graph). It is assumed that the indistinguishable vertices are of interest only in so far as they provide directions necessary to traverse a direct path between two landmarks. But it is easy to imagine situations where the indistinguishable vertices and the paths between them are of interest. For instance, the indistinguishable vertices might be partitioned further into equivalence classes so that one could uniquely designate a vertex by specifying its equivalence class and some radius from a particular global landmark (*e.g.,* the bookstore just across the street from the Chrysler building).

We shall apply our above approach and try to completely learn the graph by first completely learning local neighborhoods of each landmark. Let us define $G_d(v)$ to be the subgraph of G consisting of all vertices and edges within radius d of v.

Lemma 6 *Let* G *be a landmark graph with distribution parameter* r. *Given a procedure that, for any* $\delta_l > 0$, *completely learns* G_{2r+1} *in time polynomial in* $\frac{1}{\delta_l}$ *with probability* $1-\delta_l$, *there is a procedure that completely learns* G *with probability* $1-\delta_g$ *for any* $\delta_g > 0$ *in time polynomial in* $\frac{1}{\delta_g}$ *and the size of* G.

The above algorithms used for determining the local connectivity of landmarks can be thought of as building a search tree emanating from each landmark, in which each indistinguishable node in G may correspond to several nodes in the search tree. In order to completely learn $G_{2r+1}(v)$ (as opposed to just learning the local connectivity), we must avoid this redundant representation.

It turns out that we can extend our methods to completely learn $G_{2r+1}(v)$. The algorithm builds G_{2r+1} via an incremental breadth-first search in which each vertex encountered is tested via repeated walks from v to determine with high probability if it has already been added to G_{2r+1}. The proof requires a careful examination of the probabilities of true and false test results.

Lemma 7 *There exists a procedure that, for any* $\delta_l > 0$, *completely learns* $G_{2r+1}(v)$ *for any landmark* v *in a landmark graph with probability* $1-\delta_l$ *in time polynomial in* $\frac{1}{\delta_l}$, $\frac{1}{1-2\gamma}$ *and the size of* G_{2r+1}, *and exponential in* r.

Theorem 3 *It is possible to completely learn any landmark graph with probability* $1-\delta$ *in time polynomial in* $\frac{1}{\delta}$, $\frac{1}{1-2\gamma}$, *and the size of* G, *and exponential in* r.

3.3 Related Models

We can get the same results as in the last section if we allow movement uncertainty in the reverse direction, but demand forward movement certainty. The algorithms are similar, the justifications different. In this case, the graph can be reliably navigated by the same agent that did the map learning.

We are also investigating ways to remove the requirement of either reverse certainty or forward certainty. Reverse certainty is used in the last section to help distinguish probabilistically between true and false results in our testing procedures. We can show, for example, that if $r(1-\gamma)$ is bounded by a small constant, then efficient map learning is possible without either the reverse certainty or forward certainty requirement. Another way around this restriction is to allow the exploring agent to drop pebbles or beacons to remember where it has been.

4 Related Work

There have been many approaches to dealing with uncertainty in spatial reasoning [Brooks, 1984, Davis, 1986, Durrant-Whyte, 1988, Kuipers, 1978, Lozano-Perez, 1983, McDermott and Davis, 1982, Moravec and Elfes, 1985, Smith and Cheeseman, 1986], but most of these methods suffer from the effects of multiplicative error in estimating relative position and orientation. This paper is concerned with eliminating the effects of multiplicative error by either eliminating local uncertainty altogether or by taking enough measurements to ensure that such effects are reduced to tolerable levels. In this section, we consider two related approaches.

Kuipers defines the notion of "place" in terms of a set of related visual events [Kuipers, 1978]. This notion provides a basis for inducing graphs from measurements. In Kuipers' framework [1988], locations are arranged in an unrestricted planar graph. There is recognition uncertainty, but there is no directional uncertainty (if a robot tries to traverse a particular hall, then it will actually traverse that hall; it may not be able to measure exactly how long the hall is, but it will not mistakenly move down the wrong hall). Kuipers goes to some length to deal with recognition uncertainty. To ensure correctness, he has to assume that there is some reference location that is distinguishable from all other locations. Since there is no directional uncertainty, any two locations can be distinguished by traversing paths to the reference location. Given a procedure that is guaranteed to uniquely identify a location if it succeeds, and succeeds with high probability, we can show that a Kuipers-style map can be reliably probably almost always usefully learned using an analysis similar to that of Section 3. In fact, we do not require that the edges emanating from each vertex be labeled, just that they are cyclically ordered.

Dudek *et al* [1988] consider the problem of learning a graph in which all vertices are indistinguishable and upon entering a vertex the robot can leave by any arc indexed from the one it entered on. The robot can always retrace its steps if it remembers the directions it took at each point during exploration. The authors show that the problem is unsolvable in general, but that by providing the robot with a number of distinct markers ($k \geq 1$) the robot can learn the graph in time polynomial in the graph's size. In order to place a marker on a particular vertex, the robot must visit that vertex; in order to recover the marker at later time, the robot must return to the vertex. A vertex with a marker on it acts as a temporary landmark. No assumption is made regarding the planarity of the graph. The problem with a single marker that can be placed once but not recovered is also unsolvable, but, if you allow a compass in addition, the problem can be solved in polynomial time.

Levitt *et al* [1987] describe an approach to spatial reasoning that avoids multiplicative error by introducing local coordinate systems based on landmarks. Landmarks correspond to environmental features that can be acquired and, more importantly, reacquired in exploring the environment. Given that landmarks can be uniquely identified, one can induce a graph whose vertices correspond to regions of space defined by the landmarks visible in that region. The resulting problem involves neither recognition nor movement uncertainty. Our results in Section 3 bear directly on any extension of Levitt's work that involves either recognition or movement uncertainty.

5 Conclusion

This paper examines the role of uncertainty in map learning. We assume an environmental model that provides for a finite set of distinctive locations that can be reliably detected and repeatedly found. Under this assumption, the problem of map learning reduces to one of extracting the structure of a graph through a process of exploration in which only small parts of the structure can be sensed at a time and sensing is subject to error. We are particularly interested in showing that cumulative errors in reasoning about the global properties of the environment based on local measurements can be reduced to acceptable levels using a polynomial (in the size of the graph) amount of exploration. The results in this paper shed light on several existing approaches to map learning by showing how they might be extended to handle various types of uncertainty. Our basic framework is general enough to be applied to a wide variety of map learning problems. We have identified one particular source of uncertainty, namely continuity uncertainty (see Section 2), that we believe of particular interest in learning maps of buildings and other environments possessing an easily discernable structure.

References

[Aleliunas *et al.*, 1979] Romas Aleliunas, M. Karp, Richard, Richard J. Lipton, Laszlo Lovasz, and Charles Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proceedings of the 20th Symposium on the Foundations of Computer Science*, pages 218–223, 1979.

[Basye *et al.*, 1989] Kenneth Basye, Thomas Dean, and Jeffrey Scott Vitter. Coping with uncertainty in map learning. Technical Report CS-89-27, Brown University Department of Computer Science, 1989.

[Brooks, 1984] Rodney A. Brooks. Aspects of mobile robot visual map making. In H. Hanafusa and H. Inoue, editors, *Second International Symposium on Robotics Research*, pages 325–331, Cambridge, Massachusetts, 1984. MIT Press.

[Davis, 1986] Ernest Davis. *Representing and Acquiring Geographic Knowledge*. Morgan-Kaufman, Los Altos, California, 1986.

[Dean, 1988] Thomas Dean. On the complexity of integrating spatial measurements. In *Proceedings of the SPIE Conference on Advances in Intelligent Robotic Systems*. SPIE, 1988.

[Dudek *et al.*, 1988] Gregory Dudek, Michael Jenkins, Evangelos Milios, and David Wilkes. Robotic exploration as graph construction. Technical Report RBCV-TR-88-23, University of Toronto, 1988.

[Durrant-Whyte, 1988] Hugh F. Durrant-Whyte. *Integration, Coordination and Control of Multi-Sensor Robot Systems*. Kluwer Academic Publishers, 1988.

[Kuipers and Byun, 1988] Benjamin J. Kuipers and Yung-Tai Byun. A robust, qualitative method for robot spatial reasoning. In *Proceedings AAAI-88*, pages 774–779. AAAI, 1988.

[Kuipers, 1978] Benjamin Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978.

[Levitt *et al.*, 1987] Tod S. Levitt, Daryl T. Lawton, David M. Chelberg, and Philip C. Nelson. Qualitative landmark-based path planning and following. In *Proceedings AAAI-87*, pages 689–694. AAAI, 1987.

[Lozano-Perez, 1983] Tomas Lozano-Perez. Spatial planning: A configuration space approach. *IEEE Transactions on Computers*, 32:108–120, 1983.

[McDermott and Davis, 1982] Drew V. McDermott and Ernest Davis. Planning routes through uncertain territory. *Artificial Intelligence*, 22:107–156, 1982.

[Moravec and Elfes, 1985] H. P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *IEEE International Conference on Robotics and Automation*, pages 138–145, March 1985.

[Rivest and Sloan, 1988] Ronald L. Rivest and Robert Sloan. Learning complicated concepts reliably and usefully. In *Proceedings AAAI-88*, pages 635–640. AAAI, 1988.

[Smith and Cheeseman, 1986] Randall Smith and Peter Cheeseman. On the representation and estimation of spatial uncertainty. *The International Journal of Robotics Research*, 5:56–68, 1986.

[Valiant, 1984] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27:1134–1142, 1984.

[Yemini, 1979] Yechiam Yemini. Some theoretical aspects of position-location problems. In *Proceedings of the 20th Symposium on the Foundations of Computer Science*, pages 1–7, 1979.