Robotics and Autonomous Systems 8 (1981) 47-63
North-Holland

A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations *

Benjamin Kuipers and Yung-Tai Byun

Department of Computer Sciences, University of Texas at Austin, Austin, Texas 78712, USA

Abstract

Kuipers, B. and Byun, Y.T., A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations, Robotics and Autonomous Systems, 8 (1991) 47-63.

We have developed a robust qualitative method for robot exploration, mapping, and navigation in large-scale spatial environments. Experiments with a simulated robot in a variety of complex 2D environments have demonstrated that our qualitative method can build an accurate map of a previously unknown environment in spite of substantial random and systematic sensorimotor error.

Most current approaches to robot exploration and mapping analyze sensor input to build a geometrically precise map of the environment, then extract topological structure from the geometric description. Our approach recognizes and exploits qualitative properties of large-scale space before relatively error-prone geometrical properties.

[sensorimotor ↔ control] → topology → geometry

At the control level, distinctive places and distinctive travel edges are identified based on the interaction between the robot's control strategies, its sensorimotor system, and the world. A distinctive place is defined as the local maximum of a distinctiveness measure appropriate to its immediate neighborhood, and is found by a hill-climbing control strategy. A distinctive travel edge, similarly, is defined by a suitable measure and a path-following control strategy. The topological network description is created by linking the distinctive places and travel edges. Metrical information is then incrementally assimilated into local geometric descriptions of places and edges, and finally merged into a global geometric map. Topological ambiguity arising from sensorily indistinguishable places can be resolved at the topological level by the exploration strategy. With this representation, successful navigation is not critically dependent on the accuracy, or even the existence, of the geometrical description.

We present examples demonstrating the process by which the robot explores and builds a map of a complex environment, including the effect of sensory errors. We also discuss new research directions that are suggested by this approach.

Keywords: Distinctive place; Topological map; Cognitive map; Large-scale space; Robot exploration; Environmental mapping; Spatial reasoning.

1. Introduction and overview

We have developed a robust qualitative method for robot exploration, mapping, and navigation in large-scale spatial environments. An environment is *large-scale* if its spatial structure is at a significantly larger scale than the sensory horizon of the observer.

Experiments with a simulated robot in a variety of 2-D environments have demonstrated that our method can build an accurate map of an unknown environment in spite of substantial random and systematic sensorimotor error.

Most current approaches to robot exploration and mapping analyze sensor input to build a geo-

* This work has taken place in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Qualitative Reasoning Group is supported in part by NSF grants IRI-8602665, IRI-8905494, and IRI-8904454, and by NASA grants NAG 2-507 and NAG 9-200.

metric map of the environment, then extract topological structure from the geometric description.

sensors → geometry → topology.

In our qualitative method, location-specific control algorithms are dynamically selected to control the robot's interaction with its environment. These algorithms define *distinctive places and paths*, which are linked to form a topological network description. Finally, geometric knowledge is assimilated onto the elements of the network (*Fig. 1*):

[*sensorimotor* ↔ *control*] → *topology* → *geometry*.

This relationship is an instance of the *spatial semantic hierarchy* defined and discussed by Kuipers and Levitt [25].

(1) *The control level.* Distinctive places and path are defined in terms of the control strategies and sensory measures (called *distinctiveness measures*, or *d-measures*) which support convergence to them from anywhere within a local neighborhood. A distinctive place is defined as the local maximum found by a hill-climbing control strategy, given an appropriate distinctiveness measure. A distinctive path is defined by the distinctiveness measure and control strategy (e.g. follow-the-midline or follow-left-wall), which allows the robot to follow it.

(2) *The topological level.* A topological network description of the global environment is created *before* the global geometric map, by identifying and linking distinctive places and distinctive paths in the environment.

(3) *The geometric level.* Once a topological map is in place, the geometric map can be incrementally created by accumulating, first, local geometric information about places and paths, then global metrical relations among these elements within a common frame of reference.

Our approach, based on the spatial semantic hierarchy, provides a coherent framework for exploiting the strengths of a variety of powerful spatial reasoning methods while minimizing the robot's vulnerability to their weaknesses.

- *Cumulative location error* is essentially eliminated while traveling among distinctive places in the topological network by alternating between path-following and hill-climbing control algorithms.
- *Feedback-guided motion control* can draw on the full range of control algorithms and performance analysis methods in the fields of control engineering and control theory (e.g. [8]) to mitigate the effects of sensor and motor uncertainty on navigation ability.
- *Successful navigation is not dependent on geometric accuracy*, since the control and topology levels do not depend on the geometric description. However, when geometric information is available, it can be used to optimize route-planning or to resolve topological ambiguities.
- *Geometric sensor fusion methods* [6,11,33,40] can be naturally incorporated as methods for acquiring local geometric descriptions of places and paths in the topological network. (A global geometric description can be derived by global relaxation of local metrical relations into a single frame of reference.)
- *Indistinguishable places* – i.e. places with identical local sensory characteristics – can be identified correctly, except in the most pathological environments, using a topological matching pro-

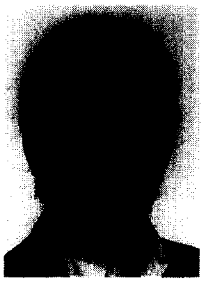

Benjamin Kuipers received the B.A. degree in mathematics from Swarthmore College in 1970, did his graduate work at the MIT Artificial Intelligence Laboratory, and received a Ph.D. degree in mathematics from MIT in 1977. He is now Associate Professor of Computer Sciences at the University of Texas at Austin. He has worked at Harvard University, Tufts University, and the Massachusetts Institute of Technology. His research interests are in various aspects of commonsense knowledge: qualitative reasoning with incomplete knowledge of physical mechanisms; resource-limited inference; and spatial exploration, learning, and problem-solving. Dr. Kuipers is a member of the Association for Computing Machinery, the American Association for the Advancement of Science, the American Association for Artificial Intelligence, IEEE, The Society for Values in Higher Education, and Computer Professionals for Social Responsibility.

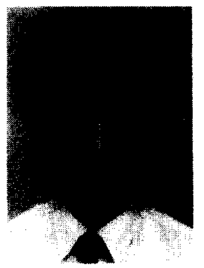

Yung-Tai Byung was born in Korea in 1954. He received the Bachelor of Science degree in Electrical Engineering from Seoul National University in 1977. He served in the R.O.K. Army from March 1977 to June 1979, and then worked for D.E.C. Korea. In August 1982, he entered Indiana University and obtained the Master of Science degree in Computer Sciences in August 1984. He then entered the doctoral program in the Department of Computer Sciences, University of Texas at Austin. He finished the program in 1989 and is currently an Assistant Professor at Hong-Ik University in Seoul, Korea. His research interests include building autonomous agents, spatial reasoning, qualitative reasoning, and neural networks.

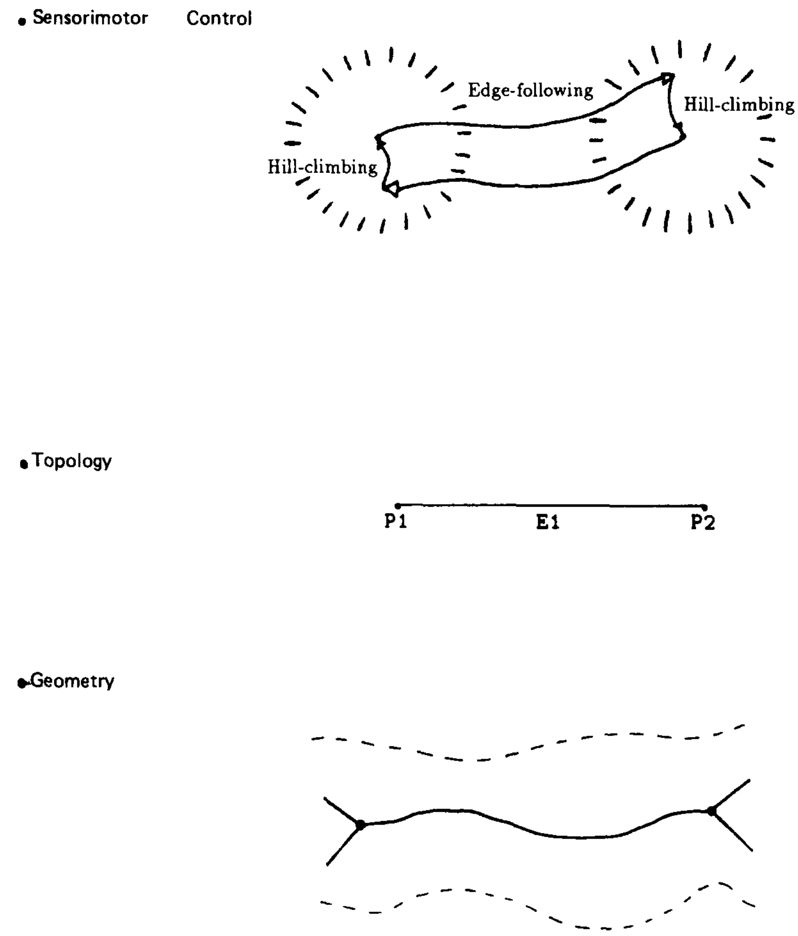

Fig. 1. The levels of spatial representation.

A layered structure isolates the different inference methods, and allows each level to establish the assumptions required by higher levels.

- (*Control*) When traveling between distinctive places, cumulative error is eliminated by alternating path-following with hill-climbing control strategies.
- (*Topology*) Elimination of cumulative error legitimizes the abstraction from a continuous physical world to a discrete topological network description.
- (*Geometry*) Geometric information is acquired about the places and paths, here in the form of a generalized cylinder description, including path length, shape, and envelope profile.

cedure to test hypotheses about the places' neighbors.

In the remainder of the paper, Section 2 reviews other approaches to spatial exploration and map-learning. Section 3 presents our hierarchical representation and its use in detail. Section 4 describes the specific instance of the hierarchical approach, the NX robot, that we have used in our research. Section 5 demonstrates the performance of NX as it explores a complex, large-scale environment, defines distinctive places and paths, links them into a topological network description, and accumulates an accurate geometrical description from metrical annotations on the topological map.

2. Background

Many researchers have studied spatial representation methods and exploration strategies: the *robot exploration and map-learning problem*. Since the goal of most approaches is a purely metrically accurate map, they are often brittle in the face of low mechanical accuracy and sensory errors [3,22,27]. However, humans perform well at spatial learning and spatial problem-solving in spite of sensory and processing limitations and frequently-incomplete knowledge [20,21]. We introduce the background of our qualitative method to the robot exploration and map-learning problem, and review the literature briefly.

2.1. Studies of the cognitive map

Many scientists [29,35,39] have observed that a cognitive map is organized into successive layers, and suggested that the central element of a useful and powerful description of the large-scale environment is a topological description. A layered model consists of the identification and recognition of landmarks and places from local sensory information; control knowledge of routes (procedures for getting from one place to another); a topological model of connectivity, order, and containment; and metrical descriptions of shape, distance, direction, orientation, and local and global coordinate systems. It appears that the layered structure of the cognitive map is responsible for humans' robust performance in large-scale space. Our approach attempts to apply these methods to the problem of robot exploration and map-learning.

The central description of environments in our qualitative approach is a topological model as in the TOUR model [19]. The model consists of a set of nodes and arcs, where nodes represent distinctively recognizable places in the environment, and arcs represent travel paths connecting them. The nodes and arcs are defined procedurally in terms of the sensorimotor control capabilities of the robot. Metrical information is added on top of the topological model.

2.2. Traditional approaches to robot exploration

Traditional spatial representation methods for known environments, and corresponding approaches to the robot exploration and map-learning problem in unknown environments, are based on the accumulation of accurate geometrical descriptions of the environment. These methods include Configuration Space [28], Generalized Cones [2], Voronoi Diagrams [15,30,31,38,44], the Grid Model [12,32,33], the Segment Model [7,42], the Vertex Model [17], the Convex Polygon Model [6,14,26], the Graph Model [15,34,36,42,44], and the Polygonal Region Model [30]. Some researchers (e.g. [16,18]) use several of these methods together.

Traditionally, sensor fusion methods such as those of [11,32,40] are used to integrate sensor input directly into one of the above geometrical representations. Within our framework, the same methods can be used, after the topological map has been created, to acquire accurate local metrical information about the places and paths of the topological structure. These can then be integrated into a relatively accurate global metrical map.

Because of low mechanical accuracy and sensory errors, it is often difficult to get accurate metrical information in large-scale space [3,6]. Some of the traditional methods perform reasonably well where environments are small enough to observe most important features from a single position. The problem is more difficult in large-scale space, as discussed by Brooks [3], Kuipers and Byun [23], and Levitt et al. [27]. A major goal of the qualitative approach reported in this paper is to overcome the fragility of purely metrical methods.

Several researchers use various types of graph model or topological model to represent the connectivity of the environment. Laumond [26] and Chatila and Laumond [6] build a topological model from the geometric model and then derive a semantic model, e.g. identifying 'rooms' and 'corridors', from the topological model. Their approach uses the topological description to represent map information at higher levels of abstraction. However, there is no use of the topological model to cope with metrical inaccuracy. Turchan and Wong [42] use an 'attributed graph' to represent the world, in which line segments and their attributes become vertices, and relations between adjacent line segments are represented by the arcs. This graph is completed by integrating local sensory information from several different locations. Their method is proposed as a way of finding a correct segment model for a large-scale environment with an error-free assumption, but it is vulnerable to errors. With the same error-free assumption, Oommen et al. [34] use a visibility graph, where vertices represent observable or actually visited meaningful points in the environment, and arcs show the connectivity of vertices for travel.

Our qualitative approach to exploration and mapping is quite consistent with the layered 'subsumption architecture' proposed by Brooks [4]. It is possible to view our procedural level as corresponding to Brooks' Level 2, 'Explore', and our topological and metrical levels as corresponding to Brooks' Level 3, 'Build maps'. We believe, however, that the structure of the exploration and mapping process is most clearly captured by the relationships between the three representations in the spatial semantic hierarchy.

Another qualitative method for place definition and navigation based on visual landmark recognition has been proposed by Levitt et al. [27]. They discuss the weakness of traditional navigation techniques and demonstrate successful exploration and navigation using a coordinate-free model of visual landmark memory, without an accurate map or metrical information. Their definition of place is based on regions, with virtual boundaries defined by line-segments connecting remote landmarks, whereas our definition of place is based on distinctiveness of a location within its neighborhood. Their methods are most appropriate in environments where remote point-like landmarks are easily observable.

Little of the literature discussing movement control strategy explicitly relates it to the topological model. Most researchers have used a goal-directed movement control strategy within a global Euclidean coordinate frame: repeat until reaching a goal (x, y), 'Try to go straight to the place, and if there is an obstacle in the way, move around until there is a possibly straight path to the goal'. Kadonoff et al. [16] use several local navigation strategies to avoid unexpected obstacles along a path without exact knowledge of the robot's position in a known world. Several different sensors and strategies are used to perform local navigation: Obstacle Avoider, Path Follower, Beacon Tracker, Wall Follower, Aisle Centerer, and Vector Summer. At any given time an arbiter using a production system dynamically selects the strategy to follow, a control structure quite similar to our procedural level. However, local navigation strategy information is neither used in describing the world nor saved for later use. Furthermore, unlike our emphasis on exploration of unknown environments, Kadonoff et al. [16] assume that a reasonably accurate world model already exists, and that beacons are available for periodically updating the robot's position.

In the machine learning community, Rivest and Schapire [37] have presented an approach for unsupervised learning in deterministic environments, a generalization of map learning. They use an extended version of our 'rehearsal procedure', which was initially developed in response to a problem posed by Rivest in 1984.

3. Building the hierarchical map

The central element of our hierarchical model is the topological network description, in which nodes correspond to distinctive places and arcs correspond to travel paths. We discuss in detail how to define distinctive places and travel paths, and their descriptions at the control and metrical levels. We also present a basic exploration strategy for building the topological model, and discuss certain implications of our approach.

A place corresponding to a node must be locally distinctive within its immediate neighborhood by some criterion definable in terms of sensory input. We introduce locally meaningful *distinctiveness* measures defined on a subset of the sensory features, by which some property can be maximized at a distinctive place. We define the *signature* of a distinctive place to be the subset of features, the distinctiveness measures, and the feature values, which are maximized at the place. A hill-climbing search is used to identify and recognize a distinctive place when a robot is in its neighborhood.

When returning to a known distinctive place, the robot is guided by the known signature. Travel paths corresponding to arcs are defined by local control strategies which describe how a robot can follow the link connecting two distinctive places. This local control strategy depends on the local environment. For example, in one environment, following the midline of a corridor may be reasonable; in another environment, maintaining a certain distance from a single boundary on one side is appropriate; in a third, moving toward a certain remote landmark is the best strategy.

The components in the topological model – the places and paths – are described at the control level in terms of control strategies and locally meaningful distinctiveness measures. Each component may also have local geometric information. This information and potentially derivable global metrical relations are the metrical level descriptions in the map.

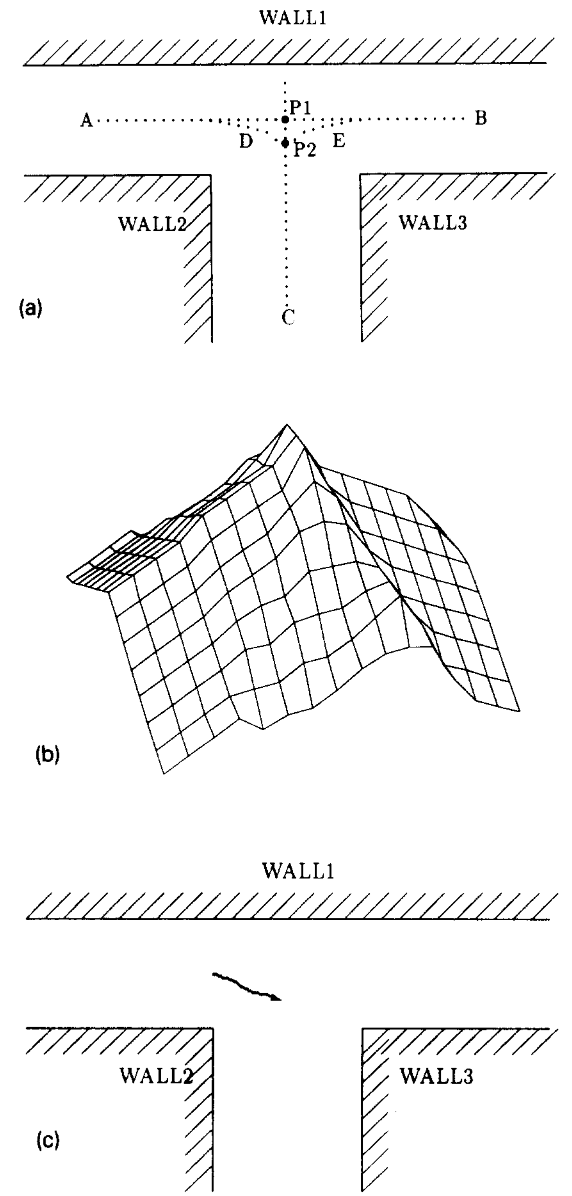

Fig. 2. A distinctive place. (a) A simple environment. (b) A distinctiveness measure (no sensory error). (c) A hill-climbing search.

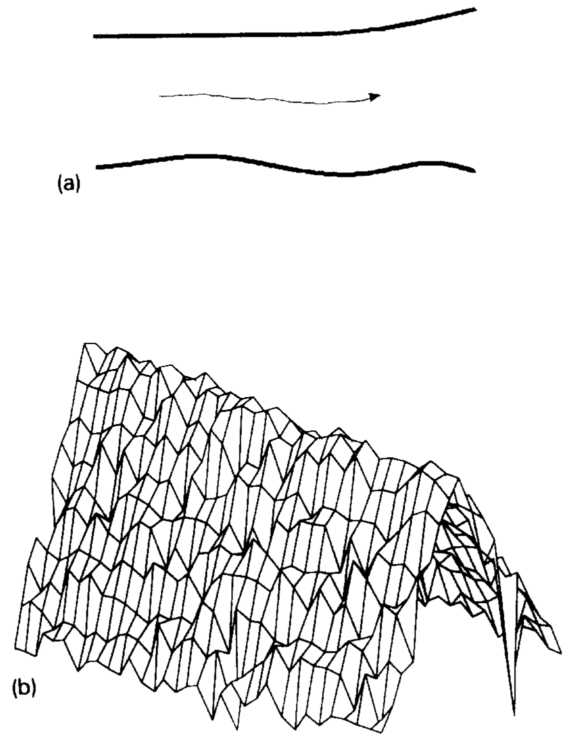

Fig. 3. A travel path. (a) Executing a local control strategy within a simple environment. (b) A distinctiveness measure (10 percent sensory error).

3.1. Distinctive places

The robot needs to identify distinctive places (DPs) in order to define the nodes of the network-structured topological model. Intuitively, if we consider the geometry of a simple 2-D local neighborhood in *Fig. 2a*, we can argue that the dashed lines represent points that are in some sense distinctive, and that the most distinctive points occur where the lines intersect, near the center.

However, the robot has only local, egocentric, sensorimotor access to the environment. [1] We need to determine which sensory characteristics provide the distinguishing features by which a place becomes locally distinctive, in order to formulate locally meaningful *distinctiveness* measures.

We hypothesize that any reasonably rich sensory system will support distinctiveness measures that can be defined in terms of low level

[1] I.e. the robot is a situated agent in the sense of Suchman [41] and Agre and Chapman [1].

sensory input. For a given sensorimotor system, we can specify the features to be maximized by corresponding distinctiveness measures. *Fig. 2b* shows the values, over a neighborhood, of a distinctiveness measure defined by a geometric feature, Equal-Distances-to-Near-Objects. The feature is specific to the sensory-motor system of this particular robot, but the notion of distinctive places is general.

Once the robot recognizes that it is in the neighborhood of a distinctive place, it applies a hill-climbing control strategy to move to the point where some distinctiveness measure has its local maximum value. *Fig. 2c* shows the result of the hill-climbing search with the same robot instance. Note here that it is not necessary for a place to be globally distinctive; it is only necessary to be distinguished from other points in its immediate neighborhood.

When connecting paths from or to a distinctive place are found, the place is described topologically in the model in terms of connecting paths and adjacent places. Metrical information from sensory devices is also used to describe a distinctive place, with information such as the distance and direction to nearby objects, the directions to true and false open space, the shape and apparent extent of nearby objects, etc. Metrical information is continuously accumulated during exploration and navigation, and averaged to minimize metrical error.

3.2. Distinctive travel paths

Just as a place is defined as a zero-dimensional local maximum, the paths followed during exploration are defined by some distinctiveness criterion that is sufficient to specify a one-dimensional set of points. *Fig. 3* shows the values of a distinctiveness measure in a corridor-like portion of the environment, and the result of performing the corresponding local control strategy, Follow the Midline.

With our current set of control strategies, the robot will follow the midline of a corridor, or walk along the edge of a large space, but will not venture into the interior of a large space, where the points have no qualitatively distinctive characteristics, at least to its limited-range sensory apparatus.

Travel paths connecting two distinctive places are defined in terms of local control strategies (LCS). Once a place has been identified, the robot selects an appropriate local control strategy for moving into an apparently open direction. While following a path with the chosen strategy, the robot continues to analyze its sensory input for evidence of new distinctive features. Once the next

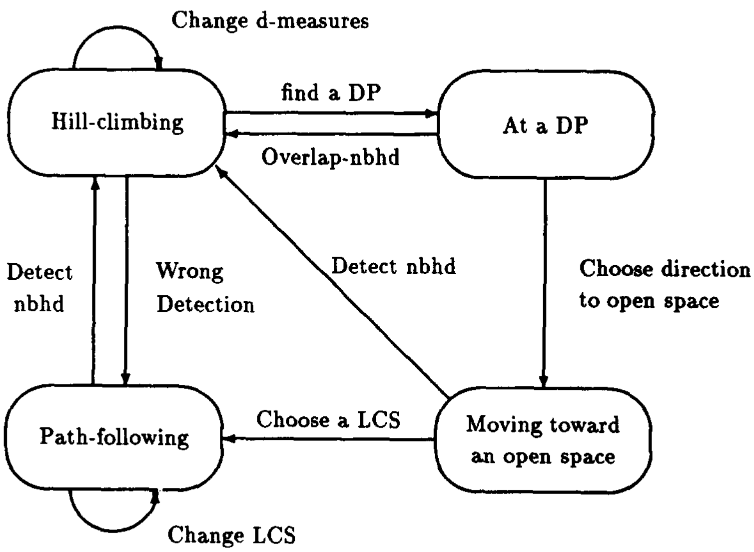

Fig. 4. A state-event diagram for the exploration strategy.

distinctive place has been identified and defined, the path connecting the two places is defined procedurally in terms of the LCS required to follow it. The knowledge used for selecting and performing the proper LCS is dependent on the robot's sensorimotor system.

Besides connectivity information, locally-observable metrical information is accumulated to describe the geometric features of a path, such as length, lateral width, curvature, net change in orientation, etc. Metrical information is continuously accumulated during exploration and navigation, and averaged to minimize metrical error.

3.3. The basic exploration strategy

We can summarize the exploration strategy by a simple state-event diagram (*Fig. 4*). The basic strategy cycles through the states in a clockwise sequence: (a) from a place, move into an open direction; (b) select a control strategy and follow a path; (c) detect a neighborhood, select a d-measure, and begin hill-climbing; (d) reach a local maximum that defines being at another distinctive place. The topological model is built as a side-effect of motion through this transition graph.

The other transitions in the graph handle exceptional cases, such as places that have overlapping neighborhoods so they are not separated by paths, and incorrect recognition of a neighborhood or choice of hill-climbing or path-following control strategy.

In pathological cases, if the robot leaves a distinctive place but cannot select an LCS, it can wander around until the available sensory information becomes sufficiently adequate to determine either a distinctive place or a local control strategy. This is similar to the relation between 'Explore' and 'Wander' processes in Brooks' [4] subsumption architecture. In terms of the state-event diagram in *Fig. 4*, control leaves this graph, and may return to either the hill-climbing or the path-following state.

3.4. Robust against errors

Although the robot is subject to sensory and movement errors, hill-climbing search based on continuous sensory feedback will bring it very near to a distinctive place. As long as the distinctiveness measures are defined and convex over an open region, it is not necessary for the robot to be located at the same (x, y) coordinates in an absolute coordinate frame.

The local control strategy taking the robot along a path from one place to another need not bring it precisely to the destination place. As long as it reliably brings the robot into the *neighborhood* of the place, hill-climbing will eliminate the error acquired during travel. It is not necessary for the robot to return to the same (x, y) position each time, as long as the behavior of the hill-climbing and path-following control strategies remains in correspondence with the topological map.

Thus, in a sufficiently well-behaved environment, by building a topological model based on an alternation of distinctive places and travel paths, our strategy effectively eliminates the problem of cumulative position error.

3.5. The position referencing problem

While a robot explores a given environment, it needs to know its current position in the map. This is the single most important task in the robot exploration and map-learning problem. In traditional approaches, the current position is represented by (x, y) in a global coordinate frame. As discussed in Section 2, it is not easy to maintain correct coordinates for the current position.

In our method, the current position is described at two levels: topological and metrical. At the topological level, the current position is described by either a distinctive place, or by a pair representing a path and a direction. At the metrical level, when the robot is at a distinctive place, the current local sensory information and its current orientation are given. When it is on a path, the robot's current position may be described in terms of the place it is coming from, the distance it has travelled, lateral distance information, and its current orientation.

3.6. The exploration agenda

During exploration, the robot uses an *exploration agenda* to keep information about *where* and *in which direction* it should explore further to complete the map. If (*Place*1, *Direction*1) is in the exploration agenda, it means that a robot has previously visited *Place*1 and left it in some direction(s) other than *Direction*1. Therefore, in order to delete (*Place*1, *Direction*1) from the exploration agenda, the robot should either (a) visit *Place*1 and leave in the direction *Direction*1, or (b) return to *Place*1 from the direction opposite to *Direction*1.

In general, directions are defined with respect to local coordinate frames at each place. Matching directions between visits to a place may require inference involving the sensory characteristics of the place and the estimated change of heading during travel. The particular robot instance we use in our experiments has an absolute compass, which simplifies this matching step.

When the robot reaches a place during exploration, the exploration agenda can be either empty or non-empty. If the exploration agenda is empty, it means that there is no known place with directions which require further exploration. Therefore the current place must be new, unless a robot has intentionally returned to a previously known place through a known path. If the exploration agenda is not empty, the current place could be one of the places saved in the exploration agenda. This is only possible when the current place's metrical description is similar to that of a place saved in the exploration agenda, and the difference between the current orientation and the direction saved on the agenda is approximately 180 degrees.

The ordering on the exploration agenda controls the overall behavior of the robot, but is largely independent of our navigation and mapping approach. It is easy to define priority schemes which tend to minimize the number of 'loose ends' on the exploration agenda, for example by giving priority to the sharpest turn from the current place that leads to an unexplored direction.

Alternatively, exploration and mapping can be treated as a background process, in which an unrelated goal-oriented process in the foreground controls the overt behavior of the robot by manipulating the order on the exploration agenda.

3.7. *The rehearsal procedure*

When a robot reaches a place during the exploration, the identification of the place is the most important task. If the place has been visited before and the robot comes back to that place, the robot should recognize it. A new place must be recognized as new, even if it is very similar to one of the previously visited places. Place matching is done using global topological constraints as well as local metrical comparison.

The current and stored place descriptions are first compared metrically, allowing a certain amount of looseness of match to provide robustness in the face of small variations in sensory input. If there is any possibility of a *false positive* match, the topological matching process is initiated. The *rehearsal procedure* is activated, and uses the topological model and control knowledge of paths and nearby distinctive places to test the hypothesis that the current place is equal to a previously known place.

The robot constructs routes between the known place and adjacent DPs. It then tries to follow the routes and return to the current place. If the

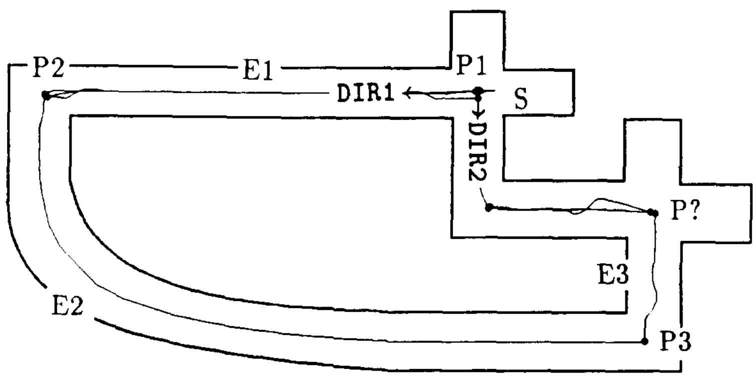

Fig. 5. An environment requiring the rehearsal procedure.

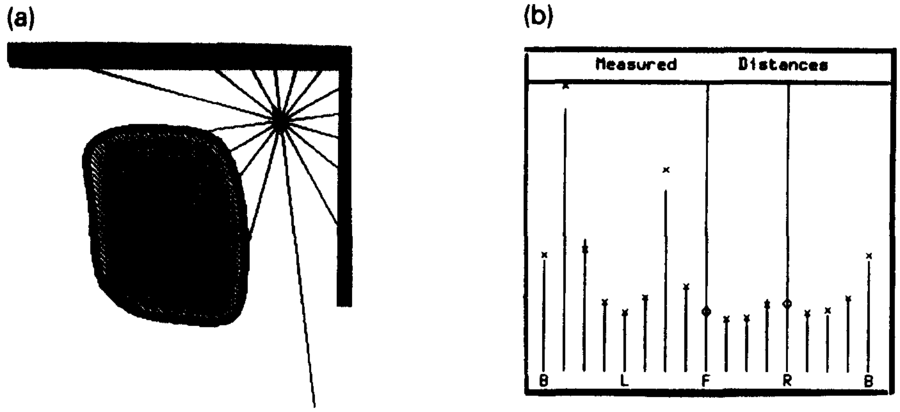

Fig. 6. (a) NX in its environment; (b) NX's sensory image.

routes performed as predicted, then the current place matches the previously known one, and the current place has been identified. If not, then the current place must be a new place with the same sensory description as the old one.

Fig. 5 shows an environment in which topological matching is necessary, and demonstrates the rehearsal procedure.

(1) The robot starts at S.
(2) It finds a DP *P1* and follows a path *E1*.
(3) When it leaves *P1*, it puts (*P1 Dir2*) in the exploration agenda.
(4) It finds *P2*, follows *E2*, and finds *P3*.
(5) It follows *E3* and gets to a place *P?* where the local sensory information is very similar to that of *P1*. *P?* may be *P1* or a different DP.
(6) It sets up a hypothesis: If *P?* is *P1*, then following *E1* (i.e. the path hypothesised to be *E1*) will bring it to *P2*.
(7) Then it tests this hypothesis by traveling along the planned route.
(8) However, it reaches a place *P5* at which the local sensory information is quite different from that at *P2*.
(9) Therefore it concludes that the hypothesis was incorrect, and *P?* is a new place *P4*.

For any fixed search radius of this topological match, it is possible to construct an environment that will nonetheless yield a false positive match. In the current implementation, to guarantee termination even in pathological environments, the rehearsal procedure is not called recursively to test for a successful prediction; only local sensory characteristics are considered. However, if there is a reference place that is marked so as to be globally unique (e.g., 'home'), a version of the rehearsal procedure can be constructed to eliminate false positive matches. These and several more sophisticated properties of the rehearsal procedure have been proved by Dudek et al. [10].

4. A robot instance: NX

We believe that our exploration and mapping approach is supported by any robot with sufficiently rich sensory input, which takes sufficiently small steps through its environment. We demonstrate our method as applied to a specific instance of such a robot.

4.1. The NX robot simulator

The robot NX exists in a two-dimensional simulated environment. The simulator is written in Common Lisp on the Symbolics 3600. Although we use this specific robot to test our qualitative method, our approach does not depend critically on the choice of sensors and movement actuators.

NX has sixteen sonar-type distance sensors covering 360 degrees with equal angle difference between adjacent sensors, two tractor-type chains for movement, and an absolute compass for global orientation. Thus the input to NX is a vector of time-varying, real-valued functions

$$[S_1(t), S_2(t), \ldots, S_{16}(t), Compass(t)]$$

represented by $(\mathbf{S}, \boldsymbol{\theta})$.

Fig. 6a shows NX's range-sensors, when it is near place P11 in *Fig. 8*. Figure 6(b) shows the 16 range-sensor readings as observed by NX at that instant. The middle line represents the direction straight ahead of the robot. The length of each line represents the perceived distance in each direction. For the aid of the researchers, an 'x' or 'o' indicates the *true* distance. The 'x' indicates that the perceived distance reflects a random error, and the 'o' indicates that the perceived distance reflects a systematic error due to specular reflection. This error simulation is based on Walter [43], Flynn [13], and Drumheller [9].

4.2. Coping with sensory errors

The first step in handling errors is a spatial smoothing operation. Basically, NX attempts to fit sensory information to a hyperbolic shape (e.g., one made by six sensors on the left side of *Fig. 6b*). This operation smoothes out random errors, and can also ignore the false open space reading that appears in the middle sensor in *Fig. 6b*. However, the second false open space reading, in the middle of the right side of *Fig. 6b*, still remains. NX considers this to represent a free space to explore between two objects on the right side.

The second step is a temporal smoothing operation, applied to sensory information accumulated over several small steps. In some cases, the second false open space in the figure can be eliminated by this operation.

The third step tests hypotheses about where objects are and where open spaces are. NX tries to check the hypothesis by moving near each open space, and determining whether its sensory image behaves as expected. A false open space will disappear or move when NX approaches it. By this method, the second false open space is completely eliminated from the description of the current surroundings. Readers can see the trace of the hypothesis-testing operation in *Fig. 8*.

In addition to random and systematic sensory errors, we simulate a five percent random error of movement control in *Fig. 8*. This can result in incorrect metrical information being accumulated about paths. Since all metrical information is local until it is propagated into a global metrical map, this does not affect the first two levels of description of the model. The effects of such errors are eliminated by the accumulation of information from several traversals. Systematic motor control errors should also be correctable by this method.

4.3. Distinctiveness measures for NX

To specify the domain-specific aspects of a navigation and mapping strategy for NX, within the framework we have described, we need to specify the distinctiveness measures, the local control strategies, the criterion for the event *Detect-Neighborhood* (*Fig. 4*), and the rules for selecting a distinctiveness measure or a local control strategy given the current sensory surroundings.

A set of production rules is used to decide whether NX is in the neighborhood of a DP and what distinctive features can be maximized in that neighborhood. Each rule checks a set of assumptions and suggests a distinctiveness measure.

The individual distinctiveness measures are an open-ended, environment- and sensor-specific set of measures. For our current robot, the measures we can define include the following.

- Extent of distance differences to near objects.
- Extent and quality of symmetry across the center of the robot or a line.
- Temporal discontinuity in one or more sensors, experienced over a small step.
- Number of directions of reasonable motion into open spaces around the robot.
- Temporal change in number of directions of motion provided by the distinct open spaces, experienced over a small step.
- The point along a path that minimizes or maximizes lateral distance readings.

The current local control strategies for paths are:

- Follow-the-Midline.
- Move-along-Object-on-Right.
- Move-along-Object-on-Left.
- Blind-Step.

As with the distinctiveness measures, a set of production rules selects a proper LCS depending on the current sensory information. Detailed descriptions of the measures, rules, and place and path structures are provided in [23,24].

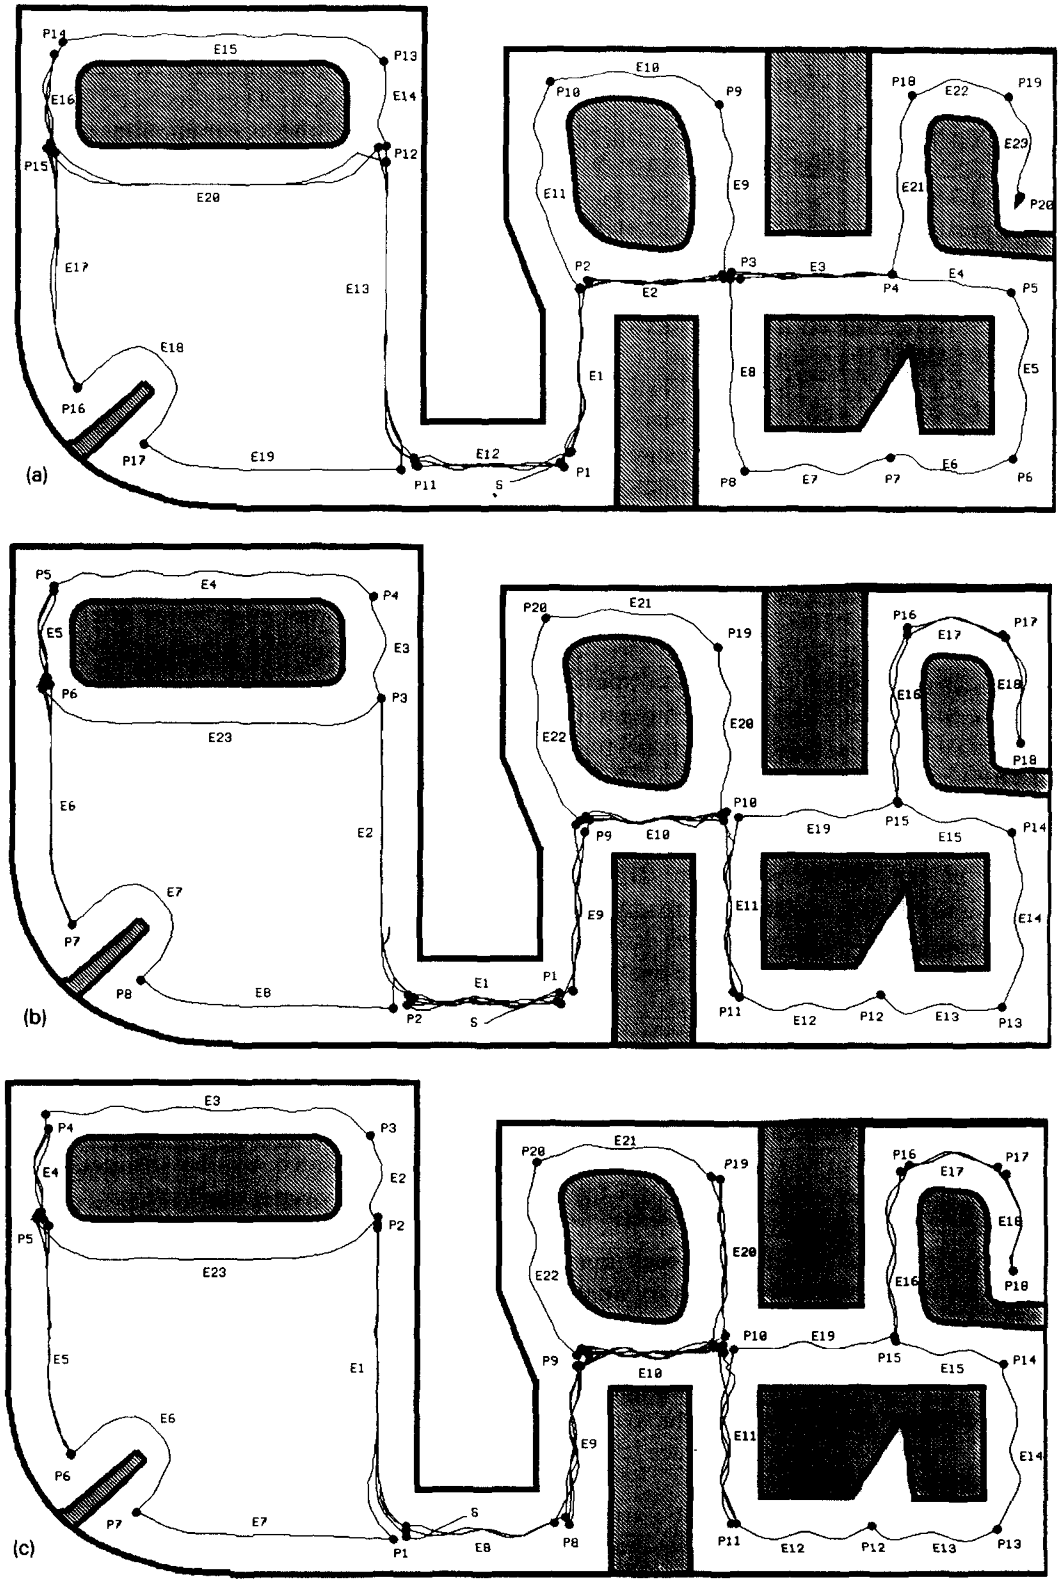

Fig. 7. The effect of random error on exploration. (a) Zero percent random error. (b) Five percent random error. (c) Ten percent random error.

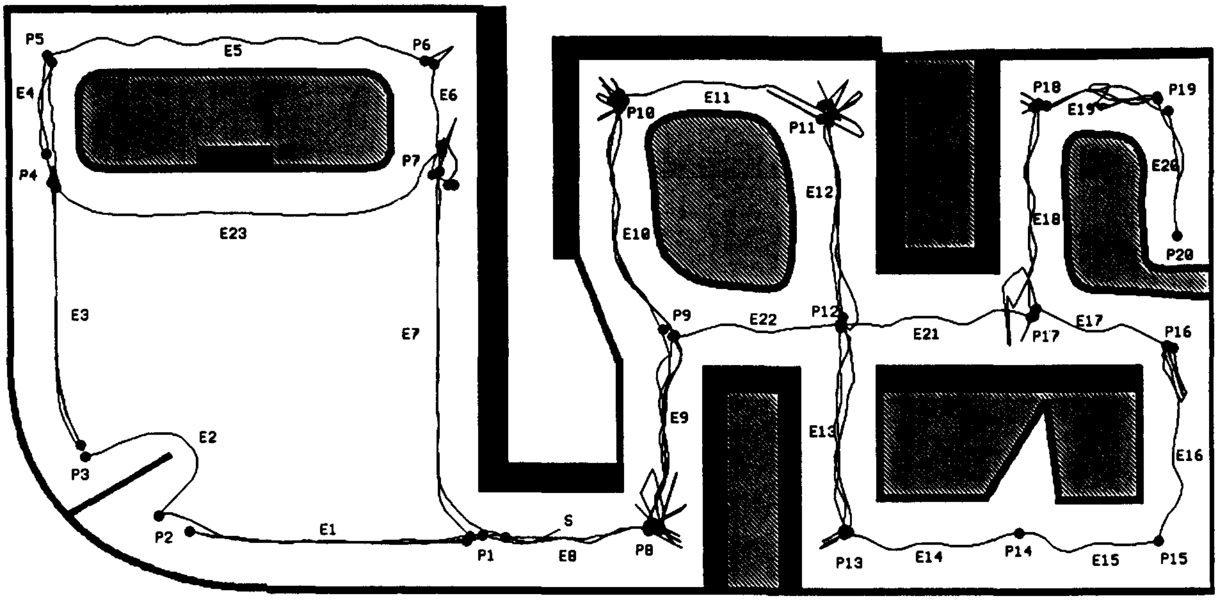

Fig. 8. Exploration results with systematic and 10% random error.

5. Exploration results

We present a detailed example showing how NX explores and builds a map, using the environment shown in *Fig 8*. (Thick black rectangles along the walls are considered surfaces which can cause systematic errors by specular reflection [9].)

To demonstrate the effect of sensory errors, we also show the exploration results for three different random error rates: the error-free case (*Fig. 7a*), five percent error (*Fig. 7b*), and ten percent error (*Fig. 7c*). NX constructs the correct map successfully in all cases, but careful examination of *Figs. 7* and *8* will reveal subtle differences.

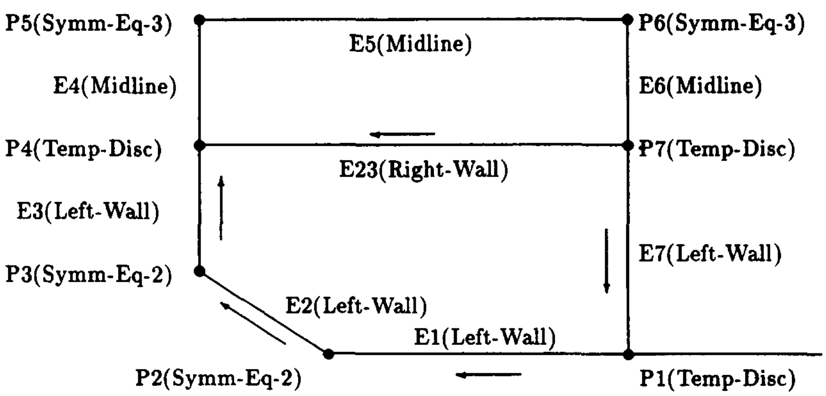

Fig. 9. Control and topological level information.

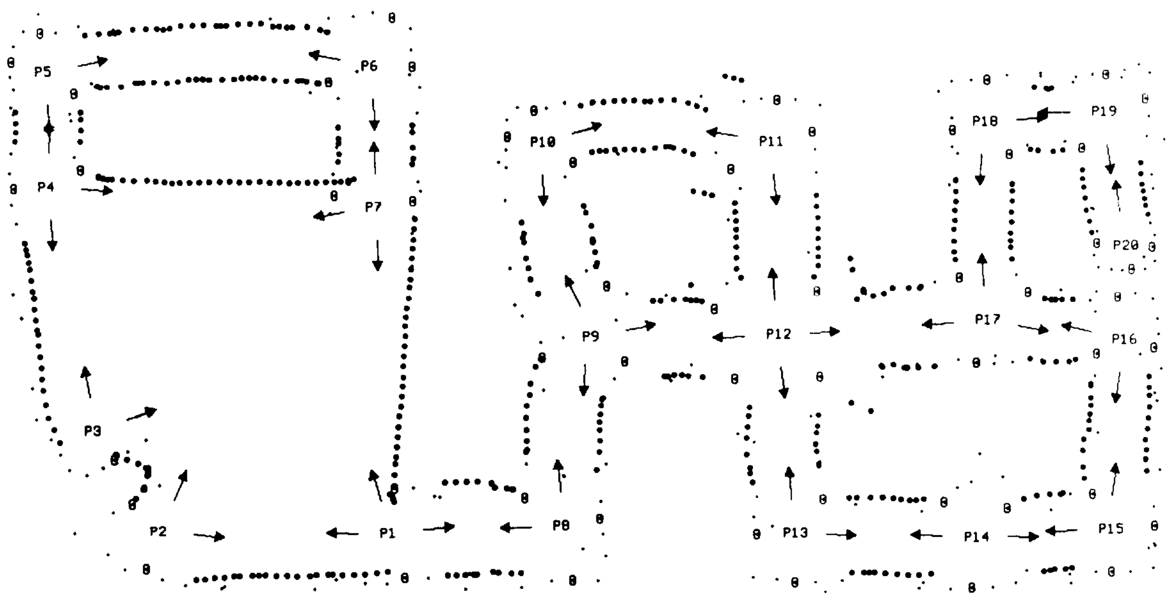

Fig. 10. Metrical level information.

Fig. 9 shows (part of) the topological model with control annotations on the paths, and *Fig. 10* shows the metrical map of the environment in *Fig. 8*.

(1) NX starts its exploration from point S in figure 8, between places P1 and P8 and directed toward P1. It determines that it is not in the neighborhood of a place, chooses the Follow-the-midline LCS, and moves to P1.

In the neighborhood of P1, it recognizes that there is a wide open-space in front, and that the angle between the directions to two near objects begins to change, after being roughly constant over a period. While executing the Follow-the-midline LCS, the angle to nearby objects was 180 degrees. (In terms of number of sensors, the angle is $N/2$ where N is the number of sensors.) As NX moves into the wide open-space near to P1, the angle becomes less than $N/2$ after being constant over a period.

This criterion is a more robust, continuous, implementation of the distinctiveness measure we originally defined as *Temporal discontinuity*. There is a large change of one sensor reading when NX moves a small amount near P1. No connectivity information is stored for P1 at this time.

The metrical information extracted from the sensory image is also recorded, and is shown graphically in figure 10. Two 'O's with small dots inside around P1 indicate the distances and directions to the nearest objects, and small dots show the rough shapes of nearby objects.

(2) There are three directions to choose from P1. If there is no particular reason to choose any particular direction, NX chooses the direction which requires the least rotation. When NX finds P1, the rotation angle to the direction toward P2 happens to be less than that toward P7 or P8. Therefore it rotates to the direction toward P2 and saves two other directions from P1 on the exploration agenda.

(3) When NX leaves P1, it chooses Move-along-object-on-left, since it has selected a direction for travel and there is a wide open-space on the right side.

(4) While it moves, NX continuously checks for the possibility of reaching the neighborhood of a DP. It finds P2 where it locally maximizes the value of distinctiveness measure Symmetry-Equal-Distances to near objects. Control and metrical features of P2 are recorded in its description, just as they were for P1.

(5) The control information about E1 (see *Fig. 9*) indicates the control strategy used for the path. The topological description of E1 says that E1 connects P1 to P2. Once E1 exists, the topological descriptions of P1 and P2 need to be changed to reflect the connectivity of the map.

Metrical information about E1 is also saved and is shown graphically in *Fig. 10*. NX records a sufficient amount of local metrical information, including leaving orientation, arriving orientation, delta orientation, travel history, distance, and lateral readings, so that a generalized cone description of each travel path is derivable. This information becomes more accurate when more traversals are made for the path.

(6) NX then follows E2 and finds P3. NX creates control, topological, and metrical descriptions for E2 and P3 as before.

(7) NX then finds and describes E3, P4, E4, P5, E5, P6, E6 and P7.

Notice that a place does not always need to be found at exactly the same physical location in the environment. We also see the trace of the hypothesis test of open-space around P6 and P7.

(8) From P7, NX explores downward to P1 in *Fig. 8* and finds a place which could be P1. The local sensory information at the place is very similar to that recorded at P1. In addition, the new place is being approached from a direction opposite to a direction saved in the exploration agenda when P1 was first seen. Therefore there is a good possibility that the current position is P1, which NX previously visited.

NX performs the rehearsal procedure as follows.

- If the new place is really P1, then NX knows from the topological map that it can reach P2 by following E1.
- It actually follows E1 and reaches P2 (or at least a place that appears identical to P2).
- It concludes that the new place actually is P1.

(9) The information saved in the exploration agenda for the direction from P1 toward P7 is deleted from the exploration agenda.

At this point, the exploration agenda now has three elements, (at P4, direction toward P7), (at P7, direction toward P4), and (at P1, direction toward P8). NX selects the third element of the exploration agenda and follows E8 to discover P8.

(10) The exploration process continues in much the same fashion. NX explores all areas of the environment, and finishes its exploration by traversing E23. At several points, the rehearsal procedure was invoked to determine whether a newly found place was the same as a previously seen place.

We see that NX had a more difficult time in *Fig. 8* than *Fig. 7*. Difficulties occurred when NX traversed between P18 and P19 and when NX performed the hill-climbing search for P17. In all four figures, NX shows a slightly different exploration order. Since there is random sensor error, the order of exploration is nondeterministic. NX continues its exploration until there is nothing in the exploration agenda and no more unexplored directions from the current place.

Once NX finishes its exploration completely, it repeatedly selects a place randomly and navigates to that place. Its control and topological level descriptions of the environment are complete, so route- finding and navigation are straightforward. However, while NX is navigating, it accumulates more metrical information to increase the metrical accuracy of its description of paths and places.

NX can also demonstrate the quality of its map by successfully orienting itself after being dropped at an unknown location within an already-explored environment, using the rehearsal procedure.

6. Summary

It is very difficult to build a metrically accurate map within a global coordinate frame, through exploration in an unknown unstructured environment. Instead, we use a hierarchical description of the spatial environment, in which a topological network description mediates between a control and a metrical level. Distinctive places and paths are defined by their properties at the control level, and serve as the nodes and arcs of the topological model. Each place and path can then accumulate local metrical information. Successful performance relying on the control and topological levels of the map is not vulnerable to errors at the metrical level, but can be improved as reliable metrical information becomes available. In suitable environments, therefore, our approach eliminates the cumulative metrical error problem of traditional approaches.

Robust performance in the face of sensory and motor errors is the result of a number of factors: the separation of semantic levels in the hierarchy, the robustness of control strategies for hill-climbing and path-following, the metrical matching with tolerance for place matching, and the rehearsal procedure for topological place-matching.

However, we can construct pathological environments where the current NX fails. The rehearsal procedure does topological matching out to a fixed radius (currently only one path), and can be deceived by a sufficiently uniform environment. The topological description may be ambiguous, due to sensory errors, or due to multiple topological models being nearly equally appropriate for a particular environment (a phenomenon we call *bifurcation*). Extensions to the straight-forward graph model of the topological description may be required to extend NX's capabilities from room-and-corridor environments like figures 7 and 8 to terrain-and-landmark environments like those studied by Levitt et al. [27].

Dynamic environments pose additional problems of three distinct types. First, the robot is currently considered low-speed or *friction-dominated*: without explicit action, no motion takes place. A more realistic model would be high-speed or *momentum-dominated*, requiring a more sophisticated set of control strategies to move through the environment. Second, there may be other moving agents ('pedestrians') moving either faster or slower than the robot within the same environment, requiring improved obstacle-avoidance capabilities. Third, the environment itself may change, as doors are opened or closed, or parked cars move around, requiring diagnosis to discriminate between fixed and changeable aspects of the environment. Research on these and other questions is under way.

References

[1] P.E. Agre and D. Chapman, Pengi: An implementation of a theory of activity, *Proc. National Conference on Artificial Intelligence (AAAI-87)* (Morgan Kaufman, Los Altos, CA, 1987).

[2] R.A. Brooks, Solving the find-path problem by good representation of free space, *Proc. AAAI-82* (1982) 381–387.

[3] R.A. Brooks, Visual map making for a mobile robot, *IEEE Proc. International Conference on Robotics and Automation* (1985) 824–829.

[4] R.A. Brooks, A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation* RA-2 (1) (1986) 14–23.

[5] Y.T. Byun, Spatial learning mobile robots with a spatial semantic hierarchical model. Ph.D. dissertation, Comp. Sci. Dept., Univ. of Texas at Austin, May 1990. (Also available as TR AI90-121, Computer Science Department, University of Texas, Austin, Texas, USA, 1990.)

[6] R. Chatila and J. Laumond, Position referencing and consistent world modeling for mobile robots, *IEEE Proc. International Conference on Robotics and Automation* (1985) 138–170.

[7] J.L. Crowley, Navigation for an intelligent mobile robot, *IEEE Journal of Robotics and Automation* RA-1 (1) (1985) 31–41.

[8] J.J. D'Azzo and C.H. Houpis, *Linear Control System Analysis and Design: Conventional and Modern*, 3rd ed. (McGraw-Hill, New York, 1988).

[9] M. Drumheller, Mobile robot localization using sonar, MIT AI-Lab TR AI-M-826 (1985).

[10] G. Dudek, M. Jenkin, E. Milios and D. Wilkes, Robotic exploration as graph construction, University of Toronto, Computer Science Technical Report RBCV-TR-88-23 (1988).

[11] H.F. Durrant-Whyte, *Integration, Coordination and Control of Multi-Sensor Robot Systems* (Kluwer, Boston, MA, 1988).

[12] A. Elfes, A sonar-based mapping and navigation system, *IEEE Proc. International Conference of Robotics and Automation* (1986) 1151–1156.

[13] A.M. Flynn, Redundant sensors for mobile robot navigation, TR 859, MIT A.I. Lab. (1985).

[14] G. Giralt, *Mobile Robots, Robotics and Artificial Intelligence*, in: M. Brady L.A. Gerhardt and H.F. Dividson, eds., NATO ASI series (1983) 375–393.

[15] S.S. Iyengar et al., Learned navigation edges for a robot in unexplored terrain, *IEEE Proc. the 2nd Conference on AI Application* (1985) 148–155.

[16] M.B. Kadonoff et al., Arbitration of multiple control strategies for mobile robots, *Proc. SPIE, Mobile Robots*, Cambridge, MA (1986).

[17] E. Koch et al., Simulation of path planning for a system with vision and map updating, *IEEE Proc. International Conference on Robotics and Automation* (1985) 146–160.

[18] D.T. Kuan et al., Natural decomposition of free space for edge planning, *IEEE Proc. International Conference on Robotics and Automation* (1985) 178–183.

[19] B. Kuipers, Modeling spatial knowledge, *Cognitive Science*, 2 (1978) 129–153.

[20] B. Kuipers, Commonsense knowledge of space: Learning from experience, Proc. IJCAI-79, Stanford, CA (1979) 499–501.

[21] B. Kuipers, Modeling human knowledge of routes: Partial knowledge and individual variation, *Proc. AAAI-83*, Washington, DC (1983).

[22] B. Kuipers and Y.T. Byun, A qualitative approach to robot exploration and map-learning, *Proc. Spatial Reasoning and Multi-Sensor Fusion Workshop*, Chicago, IL (1987).

[23] B. Kuipers and Y.T. Byun, A robust, qualitative method for robot spatial learning, *Proc. AAAI-88*, St.Paul/Minneapolis, MN (1988) 774–779.

[24] B. Kuipers and Y.T. Byun, A robust, qualitative approach to a spatial learning mobile robot, *Proc. SPIE Symposium on Advances in Intelligent Robotics System*, Cambridge (1988).

[25] B. Kuipers and T. Levitt, Navigation and mapping in large-scale space, *AI Magazine* 9 (2) (1988) 25–43.

[26] J. Laumond, Model structuring and concept recognition: two aspects of learning for a mobile robot, *Proc. IJCAI-83* (1983) 839–841.

[27] T.S. Levitt et al., Qualitative navigation, *Proc. DARPA Image Understanding Workshop* (Morgan Kaufmann, Los Altos, CA, 1987).

[28] T. Lozano-Perez, Automatic planning of manipulator transfer movements, *IEEE Trans. on Systems Man and Cybernetics* (1981) 781–798.

[29] K. Lynch, *The Image of the City* (MIT Press, Cambridge, MA, 1960).

[30] D. Miller, A spatial representation system for mobile robots, *IEEE Proc. International Conference on Robotics and Automation* (1985) 122–128.

[31] A.C.C. Meng, Free space modeling and geometric motion planing under location uncertainty, *Proc. Spatial Reasoning and Multi-Sensor Fusion Workshop*, Chicago, IL (1987).

[32] H.P. Moravec, Sensor fusion in certainty grids for mobile robots, *AI Magazine* 9 (2) (1988) 61–74.

[33] H. Moravec and A. Elfes, High resolution maps from wide angle sonar, *IEEE Robotics and Automation* (1985) 116–121.

[34] B.J. Oommen et al., Robot navigation in unknown terrains using learned visibility graphs, Part 1: The disjoint convex obstacle case, *IEEE Journal of Robotics and Automation RA-3* (6) (1987) 672–681.

[35] J. Piaget and B. Inhelder, *The Child's Conception of Space* (Norton, New York, 1967).

[36] N.S.V. Rao et al., Robot navigation in an unexplored terrain, *Journal of Robotic Systems* 3 (1) (1986).

[37] R.L. Rivest and R.E. Schapire, A new approach to unsupervised learning in deterministic environments, *Proc. of the Fourth International Workshop on Machine Learning* (1987).

[38] R.S. Rosenberg and P.F. Rowat, Spatial problems for a simulated robot, *Proc. IJCAI-81* (1981) 758–765.

[39] A.W. Siegel and S.H. White, The development of spatial representations of large-scale environments, in: H.W. Reese, ed., *Advances in Child Development and Behavior* (Academic Press, 1975).

[40] R. Smith and P. Cheeseman, On the representation and estimation of spatial uncertainty, *The International Journal of Robotics Research* 5 (4) (1986).

[41] L. Suchman, *Plans and Situated Actions* (Cambridge University Press, New York, 1987).

[42] M.P. Turchan and A.K.C. Wong, Low level learning for a mobile robot: Environment model aquisition, *Proc. 2nd International Conference on AI Application*, Miami Beach, FL (1985) 156–161.

[43] S. Walter, The sonar ring: Obstacle detection for a mobile robot, *IEEE Proc. International Conference on Robotics and Automation* (1987) 1574–1579.

[44] C.R. Weisbin, Intelligent-machine research at CESAR, *AI Magazine* 8 (1) (1987).