

A Robust, Qualitative Method for Robot Spatial Learning*

Benjamin J. Kuipers and Yung-Tai Byun
Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712

Abstract

We present a qualitative method for a mobile robot to explore an unknown environment and learn a map, which can be robust in the face of various possible errors in the real world. *Procedural knowledge* for the movement, *topological model* for the structure of the environment, and *metrical information* for geometrical accuracy are separately represented in our method, whereas traditional methods describe the environment mainly by metrical information. The topological model consists of *distinctive places* and *local travel edges* linking nearby distinctive places. A distinctive place is defined as the local maximum of some measure of distinctiveness appropriate to its immediate neighborhood, and is found by a *hill-climbing search*. Local travel edges are defined in terms of *local control strategies* required for travel. How to find distinctive places and follow edges is the procedural knowledge which the robot learns dynamically during exploration stage and guides the robot in the navigation stage. An accurate topological model is created by linking places and edges, and allows metrical information to be accumulated with reduced vulnerability to metrical errors. We describe a working simulation in which a robot, NX, with range sensors explores a variety of 2-D environments and we give its successful results under varying levels of random sensor error.

1. Introduction

Traditional approaches to the robot exploration, navigation and map-learning, based on the accumulation of accurate metrical descriptions of the environment, are highly vulnerable to metrical inaccuracy in sensory devices and movement actuators [Brooks, 1985; Chatila and Laumond, 1985; Koch et al., 1985; Moravec and Elfes, 1985; Rao et al., 1986; Turchan and Wong, 1985; Kadonoff et al., 1986]. Recent work taking a more qualitative approach [Kuipers and Byun, 1987; Levitt et al., 1987] shows great promise of overcoming the fragility of purely metrical methods.

Humans perform well at spatial learning in spite of sensory and processing limitations [Kuipers, 1979] and partial knowledge [Kuipers, 1983]. Many cognitive scientists [Lynch, 1960; Piaget and Inhelder, 1967; Siegel and White, 1975] observe that a cognitive map is organized into succes-

sive layers. These results suggest that the basic element of a useful and powerful description of the environment is a topological description. The layered model consists of the identification and recognition of *landmarks* and *places*, procedural knowledge of *routes*, a *topological model* of connectivity, order, and containment, and *metrical information* of shapes, distance, direction, orientation, and local and global coordinate systems. Our approach attempts to apply the method to the problem of robot exploration and map-learning.

The central description of the spatial environment in our qualitative approach is a topological model as in the TOUR model [Kuipers, 1978]. The model consists of a set of nodes and arcs, where nodes represent distinctively recognizable places in the environment, and arcs represent travel edges connecting them. The nodes and arcs are defined procedurally in terms of the sensorimotor capabilities of the robot. Metrical information is added on top of the topological model.

A place in the environment corresponding to a node in the topological model must be *locally distinctive* within its immediate neighborhood by one geometric criterion or another. We introduce locally meaningful "distinctiveness" measures defined on a subset of the sensory features, by which some distinctive features 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 the robot is in its neighborhood. When exploring, both the signature and the local maximum must be found. When returning to a known place, a robot is guided by the known signature.

Travel edges corresponding to arcs are defined by local control strategies which describe how the robot can follow the link connecting two distinctive places. This local control strategy depends on the local environment and there may be several possible strategies. 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.

We have implemented and tested successfully our approach with a working simulator. We will discuss our method in detail, simulation results, and further extension.

*Support for this research is provided by NASA, under grant number NAG9-200.

2. A Topological Model with Procedural and Metrical Information

The basic structure of a map, in our approach, is the topological model of which nodes are distinctive places and arcs are travel edges. We discuss how to define distinctive places and travel edges, and their procedural and metrical descriptions with a robot instance, NX.

2.1 A Robot Instance NX

We hypothesize that our approach is supported by any sensorimotor system that provides sufficiently rich sensory input, and takes sufficiently small steps through the environment. For simplicity and concreteness, we currently define a specific instance of a robot NX which 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 $[S_1(t), S_2(t), \dots, S_{16}(t), \text{Compass}(t)]$. Although we use NX to test our qualitative method, our approach does not depend critically on the choice of sensors and movement actuators.

2.2 Distinctive Places




Figure 1. Distinctive points in a neighborhood

In order to have the nodes of the network-structured topological model we need to look for distinctive places (DPs). If we consider the geometry of a simple 2-D local neighborhood in Figure 1, we can argue that the dotted lines define a set of places that are qualitatively distinctive for one reason or another. There is clearly a place which is the most distinctive compared to its surroundings. Our approach attempts to find a suitable criterion for defining a maximally distinctive place in any given neighborhood. In environments dominated by obstacles and extended landmarks, we believe that a map based on DPs and connecting edges provides a more robust topological representation than, for example, regions related by adjacency. In an environment dominated by remote, point-like landmarks, the reverse may be true [Levitt et al., 1987].

In order to formulate locally meaningful "distinctiveness" measures, we need to determine which sensory characteristics provide the distinguishing features by which a place becomes locally distinctive. We hypothesize that any reasonably rich sensory system will have distinctiveness measures that can be defined in terms of low level sensory input. Note 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.

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 consists of assumptions and a decision for the distinctive features. Here is an example:

```
(defrule DP-R10
  (if (>= (number-of-objects) 3)
      (not (all-objects-far-away))
      (not (there-is-wide-open-space)))
  (then (am-I-in-neighborhood-DP is
        'DP-symm-equal)))
```

Once NX knows what distinctive features can be maximized locally in the neighborhood of a DP, NX performs a hill-climbing search around the neighborhood looking for the point of maximum distinctiveness (e.g., minimizing differences of distances to near objects, if DP-R10 is true). When a DP is identified, it is added to the topological model with its distinctiveness measures, connectivity to edges, and metrical information.

The individual distinctiveness measures are an open-ended, domain- 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, given 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, with a small step.
- The point along a path that minimizes or maximizes lateral distance readings.

We summarize the levels of description of DPs: (An example is given in Section 3.)

- **Procedural knowledge** for a DP: Ability to recognize the neighborhood, knowledge of what features can be maximized in the neighborhood, and ability to perform the hill-climbing search to get to the DP. Learned in the exploration stage and used in the navigation stage.
- **Topological descriptions** of a DP: A node in the topological model, connected to edges and other DPs. Added to the topological model when it is found and possibly updated during the process of constructing the model.
- **Metrical information** about a DP: Local geometry like directions to OPEN-SPACE, shape of near objects, distances and directions to objects, etc. Continuously accumulated in the exploration and navigation stage and averaged to minimize metrical error.

2.3 Travel Edges

Travel edges are defined in terms of local control strategies (LCS). Once a DP has been identified, the robot moves to another place by choosing an appropriate control strategy.

While following an edge with a chosen strategy, the robot continues to analyze its sensory input for evidence of new distinctive features. Once the next place has been identified and defined, the arc connecting the two DPs is defined procedurally in terms of the LCS required to follow it.

The edges followed during exploration are defined by some distinctiveness criterion that is sufficient to specify a one-dimensional set of points. Therefore, following our 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.

As shown in Figure 2, when the robot is following a known edge from one node to another, it starts by using the hill-climbing algorithm to locate itself at the DP corresponding to the first node. It then follows the LCS associated with the arc and ends up somewhere in the neighborhood of the second place. Then the hill-climbing algorithm brings it to the DP corresponding to the second node. This method uses continuous sensory feedback to eliminate cumulative error.




Figure 2. Movement with Error

A set of production rules to decide a proper LCS depending on the current sensory information is given to NX. An example of a rule is given below.

```
(defrule LCS-rule10 ()
  (if (>= (number-of-objects) 2)
    (TwoWalls-near-to-each-other)
    (TwoWalls-almost-opposite-direction))
  (then (proper-LCS is
        'pass-on-the-midline)))
```

The current local control strategies are:

- Follow-Midline
- Walk-along-Object-Right
- Walk-along-Object-Left
- Blind-Step

In summary for edges: (An example is given in Section 3.)

- **Procedural knowledge:** Ability to choose and perform a proper LCS and knowledge of which control strategy defines the edge. Learned in the exploration stage and used in the navigation stage.
- **In the Topological model:** An edge with direction, connected to two end-places. Added to the topological model when the second end-place is found.

- **Metrical information:** Curvature, distance, change of orientation, lateral width while traveling, etc. Continuously accumulated in the exploration and navigation stage and averaged to minimize metrical error.

2.4 Position Referencing Problem

While NX explores the given environment, it needs to know the current position. The current position is described topologically rather than metrically. When NX is at a DP, the current position is described by the current place name, the current orientation in degrees, and a travel edge through which NX has come to the current place from the previous place. When NX is on an edge, the current position is described by the previous place name, the current orientation, and an indication "ON-EDGE".

2.5 Matching Process to Determine the Current Position

When NX reaches a place during its exploration, the identification of the place is the most important task. If a place has been visited before and NX comes back to that place, NX should recognize it. A new place must be recognized as new, even if it is very similar to one of the previously visited places. Our matching process is done topologically as well as metrically.

While NX explores, it uses an *exploration agenda* to keep the information about *where* and *in which direction* it should explore further to complete its exploration. If (Place1 Direction1) is in the exploration agenda, it means that Direction1 is a reasonable direction for travel (e.g., points to OPEN-SPACE) from Place1 and has not been explored. Therefore, in order to delete (Place1 Direction1) from the exploration agenda, NX should either visit Place1 later and leave in the direction Direction1, or return to Place1 from the opposite direction.

When NX gets to a place in the exploration stage, the exploration agenda can be either empty or not 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 NX has intentionally returned to a previously known place through a known edge. 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 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 current and stored place descriptions are compared metrically, allowing a certain amount of looseness of match to provide robustness in the face of small variations in sensory input. But mismatching is possible. If there is any possibility, the topological matching process is initiated. From the topological model and procedural knowledge of edges and nearby DPs, the *rehearsal procedure* [Kuipers 1985] is activated to test the hypothesis that the current place is equal to a previously known place. NX constructs routes between the known place and adjacent DPs. It then tries to follow the routes and return to the current place. If the routes performed as predicted, then the current place matches the previously known one, and NX has identified the current place. If not, then the current place must be a new place with the same sensory description as the old one (e.g., two intersections in

the first environment in Figure 5).

For any fixed search radius of this topological match, it is possible to construct an environment that will yield a false positive match. However, if there is a reference place that is somehow marked so as to be globally unique (e.g., "home"), false positives can be eliminated.

3. Simulator and Results

We have developed a simulation system NX-SIM. Figure 3 is a copy of the simulation window. NX is represented as a triangle at P4 in Figure 3. The metrical lines in the "Measured Distances" box in the upper right corner show the 16 sensor readings at the current instant. The length of the line represents the sensor reading perceived by the robot. In this example, the sensor readings are subject to a 10% random error, so the true distance is indicated by an "x" (perceived only by the researchers).

At the top left corner, the result of analysis of each distinctiveness measure considered currently is displayed. NX was located near Place1 initially. The first peak on the second row shows the symmetric and equal distance analysis while it tried to find Place1. The second, the third, and the fourth peaks correspond to Place2, Place3, and Place4, respectively.

We show the graphic exploration results of three different error rates: 0% error in Figure 4a, 5% error in Figure 4b, and 10% error in Figure 4c. NX starts near P1 in each case, marked S. P_i means Place- i and E_i means Edge- i . We will trace NX's movement with Figure 4c very briefly. It constructs the correct map successfully in all three cases, but careful examination of figures 4a-c reveals subtle differences.

Starting from S in Figure 4c, NX chooses Pass-on-the-midline and moves downward. Because of sensory error, it does not initially recognize that it is in a neighborhood. But

while continuing to perform Pass-on-the-midline, it recognizes a qualitative change, and so it performs a hill-climbing search to minimize the difference of distances to near objects. This search turns it around, converges on a local maximum, and defines the place P1. If we look at Figure 4a and 4b, we do not see this kind of backtracking around P1. NX recognizes the neighborhood sooner than in Figure 4c.

Once NX finds P1, it records P1's information in the map as follows.

PLACE :

Name = P1

Procedural: Symm-Equal (i.e., Symmetry & Equal distance)

Topological: Nil

Metrical:

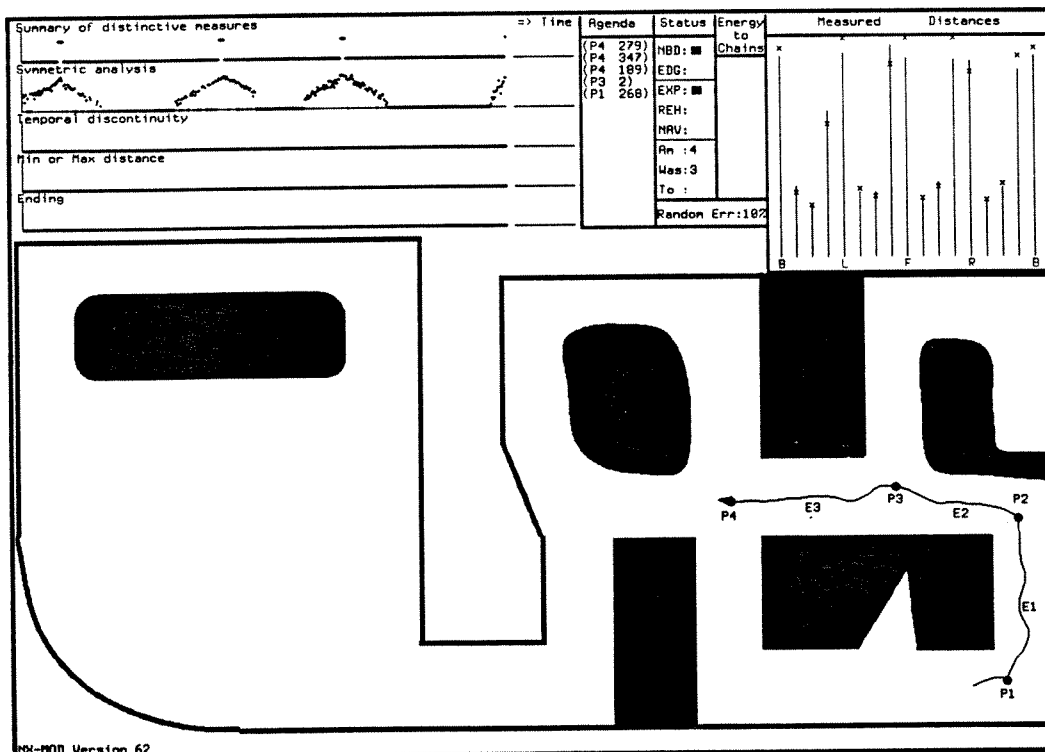
Direction-requiring-more-exploration: 345 & 282 deg.

Angle and Distance to Objects: (70 deg. 46 units)

(317 38)

(160 51)

There is no topological information for P1 at this time. There are two directions in which NX can go from P1. If there is no particular reason to choose an indicated direction, it chooses the direction which requires the *least* rotation. It rotates to the direction toward P2 and keeps the other direction on the agenda. This selection rule, of course, would cause NX to lose badly in an infinite environment. An alternate rule, selecting the direction requiring the *most* rotation, would cause the explored region to grow roughly concentrically. While NX is moving ahead from P1, it chooses Pass-on-the-midline and gathers metrical information about the edge such as distance, shape, width of the edge, change of the width, and so on. Then NX finds the second DP, P2, which is characterized by Temporal-discontinuity.



Status indicators;

NBD: a neighborhood of a DP

EDG: on an edge

EXP: Exploration

REH: Rehearsal

NAV: Navigation

Am : a current place

Was : a previous place

To : a destination place

(Values of several distinctive measures are shown on the left top, and measured distances with error are shown on the right top.)

Figure 3. NX-SIM Window




Figure 4-a. Exploration result with 0% sensor error




Figure 4-b. Exploration result with 5% sensor error




Figure 4-c. Exploration result with 10% sensor error

EDGE :

Name = E1

Procedural: Pass-on-the-Midline

Topological: from P1 to P2

Metrical:

Travel-history : ((DIR+ (8 10) (6 6) (11 9) (18 18)))

Distance : (43)

Lateral-width : ((DIR+ (81 ALMOST-STD 43)))

Minimum-width: 80

D-Orientation : ((DIR+ -8))

Once P2 has been defined, the above is recorded in the map for E1. Travel-history is a record of the number of rotations of each chain. DIR+ specifies the topological direction from P1 to P2. (81 ALMOST-STD 43) means that the distance between the two walls is approximately 81 units and almost steady while it moves approximately 43 steps. The minimum distance between the two walls along E2 is 80 units. D-Orientation gives the net change of orientation in degrees

along edges. It also updates the topological information of P1 at this moment, since E1 is connected to P1.

While NX leaves P2, NX thinks that Pass-on-the-midline is the appropriate LCS. You can see a line stretching to the direction between E2 and E6. But it soon realizes that Move-along-object-on-left or Move-along-object-on-right are more appropriate. Because it prefers smaller rotation angles, it chooses Move-along-object-on-left. We can see a significant difference between this and what happens in Figure 4a and 4b, as the result of the different amount of errors. However, the exploration process recovers from temporary errors, and is successful in all three cases.

Then NX finds P3, E3, P4, E4, P5, E5, and P6. It moves along E6 and finds a place which looks similar to P2. The rehearsal procedure is activated for topological matching. Notice here that NX does not make the same trace stretching to the middle direction between E2 and E6 as before, because it already knows that if the current place is P2, Move-along-object-on-left-side is the proper LCS. We need to emphasize that a place visited several times does not need to be exactly the same location in the environment. Accumulated metrical information and the rest of exploration are discussed in detail in [Kuipers and Byun, 1988]. We present more results with various environments in Figure 5.

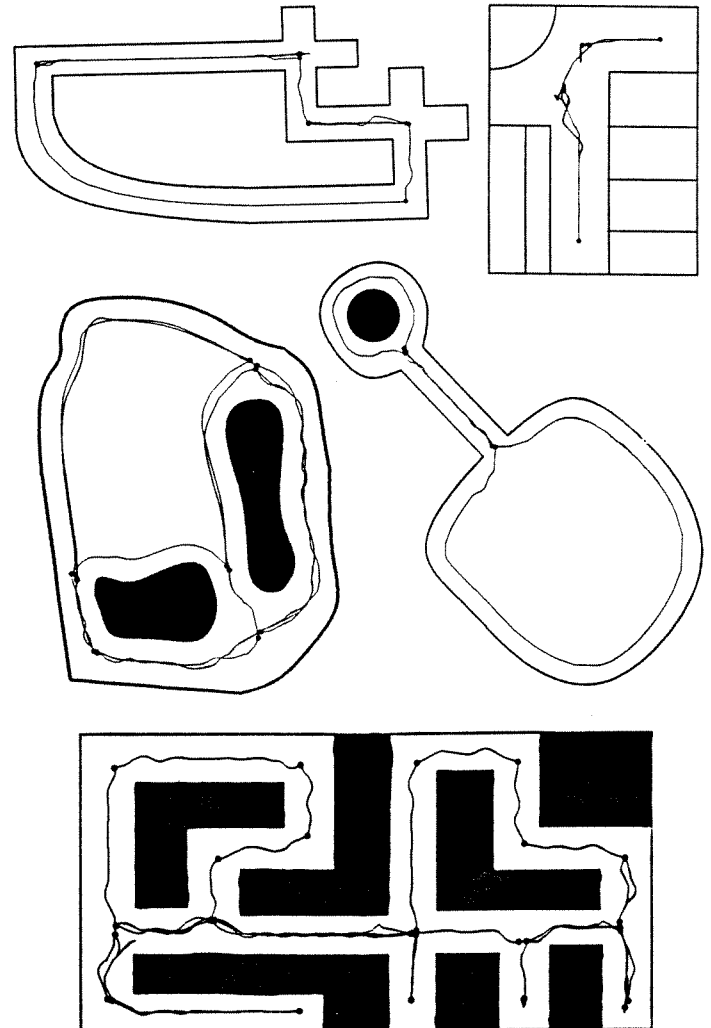


Figure 5. More exploration results

4 . Summary and Future Work

We have demonstrated a successful, robust, qualitative robot exploration and mapping method. The results show that our method can solve several of the problems of traditional approaches. The major achievement of our approach is the elimination of cumulative metrical error. Key development tasks developed currently or in the near future are listed below.

- Handling of systematic error (e.g., the acoustic peculiarities of sonar).
- Use of metrical information for optimizing routes, edge-following procedures, and correction of topological errors.
- Dynamic world (e.g., doors opening and closing; moving pedestrians).
- Removal of dependence on global compass, and use of local orientation frames and their connections.
- Hierarchical representation of complex maps.

References

- [Brooks, 1985] R.A.Brooks. Visual Map Making for a Mobile Robot, *IEEE Proc. International Conference on Robotics and Automation*, pp 824-829
- [Chatila and Laumond, 1985] R.Chatila and J.Laumond. Position Referencing and Consistent World Modeling for Mobile Robots, *IEEE Proc. International Conference on Robotics and Automation*, pp 138-170
- [Kadonoff *et al.*, 1986] M.B.Kadonoff, F.Benayard-Cheriff, and A. Franklin. Arbitration of Multiple Control Strategies for Mobile Robots, *Proc. SPIE, Mobile Robots*. Cambridge MA.
- [Koch *et al.*, 1985] E.Koch, C.Yeh, G.Hillel, A.Meystel and C.Isik. Simulation of Edge Planning for a system with Vision and Map Updating, *IEEE Proc. International Conference on Robotics and Automation*, pp 147-170
- [Kuipers, 1978] B.Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:pp 129-153
- [Kuipers, 1979] B.Kuipers. On representing commonsense knowledge, In N.V.Findler (Ed.), *Associative networks: The representation and use of knowledge by computers*. New York: Academic Press.
- [Kuipers, 1983] B.Kuipers. Modeling human knowledge of routes: Partial knowledge and individual variation. *Proc. AAAI-83*, Washington, D.C.
- [Kuipers, 1985] B.Kuipers. The Map-Learning Critter. UT AI Lab TR-18
- [Kuipers and Byun, 1987] B.Kuipers and Y.T.Byun. A Qualitative Approach to Robot Exploration and Map-Learning, *Proc. Spatial Reasoning and Multi-Sensor Fusion Workshop*, Chicago.
- [Kuipers and Byun, 1988] B.Kuipers and Y.T.Byun. A Robust Qualitative Method for Robot Exploration and Map-Learning, UT AI Lab TR-73
- [Levitt *et al.*, 1987] T.S.Levitt, D.T.Lawton, D.M.Chelberg and P.C.Nelson. Qualitative Navigation, *Proc. DARPA Image Understanding Workshop*, Los Altos.
- [Lynch, 1960] K.Lynch. *The Image of the City*, Cambridge, MA: MIT Press.
- [Moravec and Elfes, 1985] H.Moravec and A.Elfes. High Resolution Maps from Wide Angle Sonar, *IEEE Robotics and Automation*. pp 116-121
- [Rao *et al.*, 1986] N.S.V.Rao, S.S.Iyengar, C.C.Jorgensen and C.R.Weisbin. Robot Navigation in an Unexplored Terrain, *Journal of Robotic Systems*, Vol.3 No.4.
- [Piaget and Inhelder, 1967] J.Piaget and B.Inhelder. *The Child's Conception of Space*, New York: Norton.
- [Siegel and White, 1975] A.W.Siegel and S.H.White. The development of spatial representations of large-scale environments, *Advances in child development and behavior*, H.W.Reese Ed, New York: Academic Press.
- [Turchan and Wong, 1985] M.P.Turchan and A.K.C.Wong. Low Level Learning for a Mobile Robot: Environment Model Acquisition, *Proc. The 2nd International Conference on AI Application*, Miami Beach, FL. pp 156-161