

Bio-inspired Model of Robot Spatial Cognition: Topological Place Recognition and Target Learning

Alejandra Barrera and Alfredo Weitzenfeld

Abstract—In this paper we present a model designed on the basis of the rat’s brain neurophysiology to provide a robot with spatial cognition and goal-oriented navigation capabilities. We describe place representation and recognition processes in rats as the basis for topological map building and exploitation by robots. We experiment with the model by training a robot to find the goal in a maze starting from a fixed location, and by testing it to reach the same target from new different starting locations.

I. INTRODUCTION

SIMULTANEOUS localization and mapping (SLAM) addresses the problem of a mobile robot acquiring a map of its environment while simultaneously localizing itself within this map [1]. The past decade has seen extensive work in SLAM related problems. Different approaches to map building have been proposed, such as topological [2], metric [3], and hybrid maps combining these two approaches [4]-[6]. Additionally, many different issues have arisen as critical to practical and robust SLAM implementations. Among these, data association [7] and perceptual ambiguity [8] are crucial issues in recognizing paths and locations having been already traveled.

Data association relates to whether or not two features observed at different points in time correspond to one and the same object or place in the physical world [1]. On the other hand, perceptual ambiguity arises when trying to distinguish between places in the environment that may provide equivalent visual patterns [5]. Either, incorrect associations between new sensory data and previously mapped places or incorrect distinctions between perceptually similar places, can lead to a catastrophic failure in SLAM algorithms. Such failures become evident when a robot returns to a previously mapped region after a long excursion, the so-called loop-closure problem [8], [9].

Appearance-based methods and multi-hypothesis techniques address the data association problem. Appearance signatures and image similarity metrics related to color, shape and texture have been developed for recognizing places in topological mapping [10], and recently, appearance measures

have been applied also to detect loops in metric SLAM approaches [11]. Multi-hypothesis techniques, on the other hand, generate multiple data association hypotheses when new features are observed although adding computational overhead. Later, a single hypothesis becomes selected as more sensor data arrive. Examples of these techniques include multi-hypothesis Kalman filter (MHT) [12] and particle filter-based algorithms like FastSLAM [13].

Similar to SLAM algorithms having been developed for robots, animals such as rats and primates also rely on correct data association to solve spatial tasks in goal-oriented behaviors. Place recognition in rats is based on information stored in internal space representations often referred to as cognitive maps [14] that are located in an area of the brain known as hippocampus [15]. According to [16], a cognitive map is built gradually by means of an extensive exploration of the environment attaching topological and metric information based on the animal’s orientation and its estimation of distances to recognized objects.

Neurophysiological studies suggest that spatial learning in rodents is achieved by neurons in the hippocampus, known as place cells, codifying information about physical locations in the environment [15]. Experimental work has shown that the representation encoded by place cells integrates visual cues with kinesthetic feedback information in order to recognize places already visited thus distinguishing among perceptually similar places [17].

These biological studies offer the attractive prospect of taking inspiration from animals such as rats to incorporate in mobile robots adaptive navigation behavioral models. Taking inspiration from the rat’s hippocampus, several robotic navigation models have been proposed such as [18] that builds space representations from visual cues; [19] that combines kinesthetic with visual information based on hippocampal representation; and [20] that also integrates both kinds of information. In general, most of these models were tested in virtual simulations with very limited experimental work on real robots or any sort of linkage to topological or metric maps built during robot exploration.

Over the past months we have developed a navigation model based on the neurophysiology of the rat hippocampus that allows an actual robot to learn reward locations in different mazes, while building a topological map of the environment as well as returning home autonomously by means of Hebbian [21] and reinforcement learning [22]. This model relies on purely kinesthetic information to identify places in the environment and has been tested in different

This research was partially supported by collaboration projects UC MEXUS CONACYT (ITAM – UCSC), LAFMI CONACYT (ITAM – ISC), NSF CONACYT (ITAM – UCI) under grant #42440 and “Asociación Mexicana de Cultura, S. A.”

Alejandra Barrera and Alfredo Weitzenfeld are with the Computer Engineering Department – Robotics and CANNES Laboratories at the Instituto Tecnológico Autónomo de México. Río Hondo #1, Tizapán San Ángel, CP 01000, México DF, México (e-mail: abarrera@itam.mx, alfredo@itam.mx).

learning and mapping experiments, with partial results reported in [23] and [24]. In this paper we present our latest work in the integration of kinesthetic and visual information derived from artificial landmarks placed in the environment to improve the data association process involved in the robot spatial cognition. In this way, Section II of the paper presents place recognition and target learning processes in rats, Section III presents the place recognition and target learning model description, Section IV describes the robot experimentation and results, and we conclude in Section V.

II. PLACE RECOGNITION AND TARGET LEARNING IN RATS

Place recognition and target learning in rats have been clearly demonstrated through what is considered a “classical” neurophysiological experiment devised by Morris in 1981 [25]. Under this experiment, normal rats and rats with hippocampal lesions were independently placed in a circular tank filled with an opaque mixture of milk and water. They were required to swim until they located the platform, upon which they could stand and escape from the cold water. During the complete experiment, the platform remained at the same location in the tank. During the training phase, rats started from the same position in the periphery of the tank, whereas in the testing phase, rats started from different positions. In the case where the platform was visible, all rats were able to find it and escape during test trials. In a different set of experiments, the platform was submerged inside the tank, and visual cues were placed outside the tank around the arena. In this case, rats with hippocampal lesions were unable to find the platform, whereas normal rats found it during all test trials, taking direct routes to the platform although they had never swum from those starting positions.

An important contribution of the Morris experiment is the distinction between reference memory and spatial memory. When the platform is visible, rats just needed to swim towards the platform by employing their reference memory. However, when the platform was hidden, rats needed to relate the location of the platform to the location of external landmarks, thus using their spatial memory. The dependence of this capability on the existence of an unlesioned hippocampus provided a confirmation of O’Keefe and Nadel hypothesis [15] of a cognitive map located in the mammalian hippocampus. In SLAM terminology, we could argue that a data association process is employed by rats with unlesioned hippocampus to recognize the location of the hidden platform in the Morris experiment.

In a previous version of our model [23], [24], when the robot was located at a certain decision point in a maze from where the goal could be perceived, the robot was able to follow a direct route to the goal. Technically in this case, the robot took advantage of its reference memory to reach the goal. Later on, we restricted goal perception to nearby robot locations; yet, the robot was able to learn the route to the goal when starting from a unique departure position during all training and testing trials. In this case, the robot recognized

places previously visited, learned routes to these places, and distinguished between perceptually similar places due to path integration and reinforcement learning capabilities.

When the robot was tested to find the goal starting from different departure positions in the maze, path integration became an unreliable method to recognize previously visited places. Thus, we took inspiration from Morris experiment to extend our previous model and robot capabilities with spatial memory. We developed a landmarks processing module to manage visual cues placed around the maze while updating the place representation module accordingly. In this way, the robot can relate the location of landmarks to the location of the goal and the rest of the maze where landmarks can be perceived. The basic spatial cognition model and the newly added extensions are described in the next section.

III. A MODEL OF SPATIAL COGNITION

The rat-inspired model of spatial cognition comprises a number of modules: affordances, motivation, path integration, landmarks, learning, place representation, and action selection as shown in Fig. 1. A detailed description of the affordances, motivation, path integration, learning and action selection model components were presented in [23], [24]. In this section we quickly overview these modules and provide a detailed description of our newly developed landmark processing module and the improved place representation module.

A. Model Architecture

Sensory inputs to the model include affordances, motivation, kinesthesia, and landmarks. Affordances represent possible turns the rat can execute at any given time from -180° to $+180^\circ$ in 45° intervals and are coded by a linear array of cells called the affordances perceptual schema (AF). Determination of affordances is based on a local coordinate system that is relative to the rat’s head (see Fig. 2(a)).

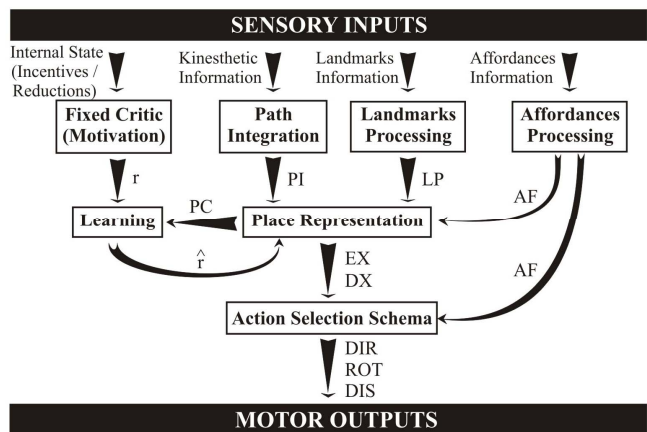


Fig. 1. The modules of the spatial cognition model and their interaction. r = immediate reward; PI = kinesthetic information pattern; LP = landmarks information pattern; AF = affordances perceptual schema; PC = place information pattern; \hat{r} = effective reinforcement; EX = expectations of maximum reward over a sequence of nodes and their corresponding directions (DX); and resulting DIR = rat direction, ROT = rat rotation, DIS = rat displacement.




Fig. 2. (a) Local coordinate system used in the model to determine affordances. (b) Global coordinate system used in the world graph layer of the model to build the topological map of the environment. This system is relative to the departure location (the base of the T in the figure).

The animal's motivation is related to its need to eat, i.e. the hunger drive. The Fixed Critic (FC) module of the model computes the hunger value and immediate reward (r) the animal gets by the presence of food.

Kinesthetic information relates to magnitude and direction of rat movement, and is used by the path integration module to update the position of the rat's point of departure each time it moves relative to its current location.

Goal oriented behavior is achieved by using reinforcement learning by means of an Actor-Critic architecture [22] composed of an Adaptive Critic (AC) and Actor units. AC estimates the future reward value of any particular place or location (PC) and computes the effective reinforcement signal (\hat{r}) that supports the learning process.

At a given location, the choice to turn in a specific direction is determined by the action selection schema (SS) by means of four signals corresponding to available affordances, a random rotation, rotations that have not been explored from the current rat's location (curiosity), and the global expectation of maximum reward. In this way, SS determines the motor outputs for the rat in terms of: direction of the rat's head (DIR) from 0° to 315° in 45° intervals, necessary rotation to point in the desired direction (ROT), and moving displacement (DIS). While returning to the departure point after having finished a training trial in an experiment, SS implements a backwards reinforcement over the nodes in the path followed by the rat. The return process and the backwards reinforcement are documented in detail in [26].

B. Landmarks Processing

Spatial information used by rats to navigate in the environment includes location of goal related to the landmarks. Experiments with rodents employing colored cylinders as landmarks [17] suggested that the location of food is coded relative to each landmark independently, computing the distance and direction of food from each landmark. Considering these findings, we also used colored cylinders as landmarks in the environment. Colors of cylinders are used to estimate the distance and relative orientation of each visible landmark to the rat. The dynamics of the landmarks processing module is shown in Fig. 3. Distance and orientation of each visible landmark from the rat are represented by two linear arrays of cells in a landmark perceptual schema (LPS). Each array PS represents distance or orientation in terms of a Gaussian function using (1):




Fig. 3. Landmarks processing module of the spatial cognition model. FDL stands for Feature Detector Layer; w = connection weights; L1= landmark 1 information pattern; Ln= landmark n information pattern; LP= landmarks pattern representing an egocentric view from the rat.

$$PS_i = h e^{-\frac{(i-a)^2}{2f^2}}, \quad (1)$$

where f is the width of the Gaussian, h is its height, and a corresponds to its medium position. The medium position of the orientation Gaussian corresponds to the relative rotation of the rat to the landmark from -180° to $+180^\circ$ in 45° intervals. The medium position of the distance Gaussian is determined by the proportion between the estimations of the current distance and the maximum distance to the landmark.

Every neuron in LPS is randomly connected to 50% of the neurons in a particular landmark feature detector layer (LFDL). Connection weights between layers are randomly initialized and normalized between 0 and 1. The activation level A_j of neuron j in LFDL is computed by adding the products between each input value I_i coming from neuron i in LPS and the corresponding connection weight w_{ij} as described in (2):

$$A_j = \sum_{i=1}^n I_i w_{ij}. \quad (2)$$

We apply Hebbian learning to update connection weights between layers in order to associate the activation of groups of neurons in LPS with the activation of groups of neurons in LFDL. In this way, we have incorporated into the model groups of neurons in LFDL that respond to specific landmark information patterns represented by LPS. The application of the Hebb rule is shown in (3):

$$\Delta w_{ij} = \alpha I_i w_{ij} L_j, \quad (3)$$

where α is the learning rate, and L_j is a new activation value of neuron j between 0 and 1 that depends on the magnitude of the original activation level A_j within the hierarchy of activation levels in LFDL. Updated weights are then normalized between 0 and 1.

One LPS layer and one LFDL layer is added to the model for each landmark in the environment. Then, all LFDL layers are combined into a single landmarks layer (LL) following the same connectivity pattern used to define the connections between a LPS and a LFDL, i.e. each LFDL is randomly connected to 50% of the neurons in LL. The activation level of cells in LL is computed as in any LFDL using (2), and the

connection weights between layers are updated also as in LFDL using (3). The activity pattern (LP) in LL represents in this way an egocentric view from the animal.

C. Place Representation and Recognition

Two layers in the model called Place Cell Layer (PCL) and World Graph Layer (WGL) carry out place representation and recognition processes. The dynamics of these layers is shown in Fig. 4. To represent a place, neurons in PCL behave like place cells in the rat hippocampus, i.e. their activation level are determined by the combination of kinesthetic information and landmarks information. We apply Hebbian connections between the path integration module and PCL and also between the landmarks processing module and PCL.

The topological map is implemented by WGL. Nodes represent distinctive places and arcs between nodes are associated with the direction of the rat's head when the animal moves from one node to the next one. Determination of the direction of the rat's head is based on a global coordinate system (see Fig. 2(b)). The model assumes that the animal can orient itself to eight directions at any given place. The eight activation patterns generated by PCL are stored in Actor units. Thus, every node in the map can be connected to eight Actor units, one for each direction. Every connection is associated with a weight that represents the expectation of finding a reward when orienting to the specific direction at the current location.

To determine whether or not the rat recognizes a place, WGL searches the current activation pattern produced by PCL (PC) within the Actor units belonging to all nodes in the map. This search involves the computation of the similarity degree SD between PC and every k stored pattern pat^k as described by (4).

$$SD^k = \sum_{i=1}^n \min(pat_i^k, PC_i) / \sum_{i=1}^n PC_i, \quad (4)$$

where n is the total amount of neurons in PCL, i is the index used in any activation pattern, and \min is a function that computes the minimum value between its two arguments. The stored pattern having the biggest SD and exceeding a threshold value close to 1 is considered the winner pattern together with its corresponding Actor unit. If there is no winner, PC is stored in a new Actor unit associated to the current rat's direction.




Fig. 4. Place representation module of the model. PI= kinesthetic information pattern; LP= landmarks information pattern; w= connection weights; PC= place information pattern; AF= affordances perceptual schema; \hat{r} = effective reinforcement signal; EX= expectations of maximum reward over a sequence of nodes and their corresponding directions (DX).

Either a winner pattern is detected or a new one is created, WGL activates or creates a node in the map depending on the current available affordances. If they are different from those perceived at time $t-1$ and a new Actor unit was created, then, a new node is created and linked to the active node through an arc associated with the current rat's direction. The new node becomes the active one, and the Actor unit just created is connected to it. Fig. 5(a) illustrates this case. If affordances did not change from $t-1$ to t and a new Actor unit was created, the model provides a representation of all similar places in an entire corridor of the environment. Thus, WGL computes the average between the activation pattern of the new Actor unit and the pattern currently stored in the Actor unit of the active node that is associated to the current rat's direction (see Fig. 5(b)). On the other hand, if there was a winner Actor unit and there exists an arc from the active node pointing to the node connected to that Actor unit, this node becomes the new active one (Fig. 5(c)). In case there is not such arc, a new arc is created to link the active node to the node associated to the winner Actor unit, and this is the new active node (Fig. 5(d)).

IV. ROBOT EXPERIMENTATION AND RESULTS

The rat cognitive model was designed and implemented using the NSL system [27]. The system can interact with a virtual or real robotic environment through an external visual processing module that takes as input the image perceived by the robot, and a motor control module that executes rotations and/or translations on the robot. We tested the model with a Sony AIBO ERS-210 4-legged robot. As sensory capabilities, we only used the 2D vision system of the robot, whose view field covers about 50° in the horizontal plane and 40° in the vertical plane. Using its local camera, the robot takes at each step three non-overlapping snapshots ($0^\circ, +90^\circ, -90^\circ$).

The experimental environment consists of the maze shown in Fig. 6(a). Three colored cylinders representing landmarks are placed outside the maze. Colored papers pasted over the walls inside the maze are used just to compute affordances, since we are using only the robot head camera to detect obstacles in our experiment.




Fig. 5. The activation/creation of nodes in the map. Dotted lines illustrate new components, and the crossed node is the active one. nA= new Actor unit; nN= new node; cD= current rat's direction; A= existing Actor unit; wA= winner Actor unit. (a) The creation of a new node. (b) The average between the activation patterns of two Actor units. (c) The activation of an existing node. (d) The connection between two existing nodes.

The objective of the overall experiment is considered as two-fold: (i) to test the place recognition process carried out by the robot, and (ii) to test the learning of the correct route to the goal. As we mentioned in Section II, the spatial task we implemented is inspired on Morris experiment [25]. We decided not to implement open-field navigation in order to exploit the affordances module of the model.

In every training trial of our experiment the robot starts from the fixed position TD shown in Fig. 6(a), and explores the maze until it finds the goal or the end of a corridor, then it returns to the departure location. The robot is programmed to recognize the goal just one step away from it. While the robot returns, the backwards reinforcement process takes place. According to the model, the robot needs to find the goal in 5 training trials to learn the route that leads towards it from TD (i.e., to increase the expectations of reward associated to the directions of the arcs linking the correct route). The map built by the robot during the training phase is shown in Fig. 6(b). It is composed of 15 nodes, and some of them represent more than one location in the maze.

After training, we placed the robot at different departure positions (D1, D2 and D3) during testing trials. Fig. 7(a) shows D1, D2, D3 and three examples of routes followed by the robot in locating the goal. We should point out that for all three starting positions the robot found the goal successfully during a single test. The updated map of the environment after completing the testing phase is shown in Fig. 7(b). As can be seen, the robot is able to recognize all nodes belonging to the correct route to the goal (black nodes), and to follow it as soon as it finds a place associated to any of these nodes. In the testing trials starting at D1 and D2, the robot followed the correct route when it found node 7 in the map, and in the trial starting at D3, it did so when it found node 3.

As shown in Fig. 7(b), the robot added four new nodes to the map (16, 17, 18 and 19). Considering for example the trial starting at D2, in its way to node 7, the robot did not recognize nodes 14, 12 and 10 corresponding to places already visited although not belonging to the learnt route. Why did the robot recognize node 7 and not node 10? The robot visited during training the location associated to node 7 oriented to 90° , 0° and 180° . Thus, there are three Actor units connected to node 7, one for each direction, storing the combination of the kinesthetic pattern and the landmarks pattern produced by PCL at that moment. In the test, the robot visited again this location with similar orientation of 180° . The activation pattern produced by PCL at this time is very similar to the one stored in Actor unit 180° of node 7 since the robot perceived the same configuration of landmarks. Note that kinesthetic information could vary due to different departure points in this testing trial.

On the other hand, the robot visited the location associated to node 10 oriented to 0° during training; thus, there is one Actor unit connected to node 10, corresponding to direction 0° and storing the activation pattern produced by PCL at that moment. In the test, the robot visited this location oriented to

180° , and the activation pattern produced by PCL at this time is different from the one stored in the Actor unit 0° of node 10 because the robot perceived a distinct configuration of landmarks. Thus, a new Actor unit is created corresponding to direction 180° and connected to a new node 17, which is created and linked to existing node 13 (see Fig. 7(b)), since affordances changed from time $t-1$ to t and there was not an arc from the active node 13 in direction 180° . Node 17 is the new active node in the map. The creation of nodes 16, 18 and 19 is due to similar reasons. The robot's behavior is shown in a "shortened" video found in [28].

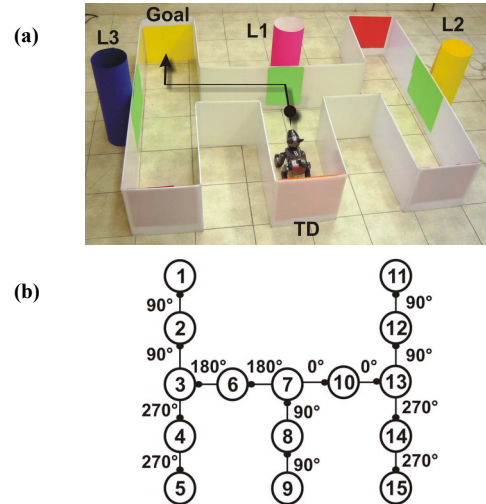


Fig. 6 (a) The physical maze used in the spatial cognition experiment. AIBO robot is located at the training starting position (TD), and the path to the goal is marked with an arrow. Landmarks are labeled as L1, L2, L3. (b) The map built by the robot during the training phase of the experiment. Nodes are identified by numbers and arcs are associated to the robot's direction when it went from one node to the next one.

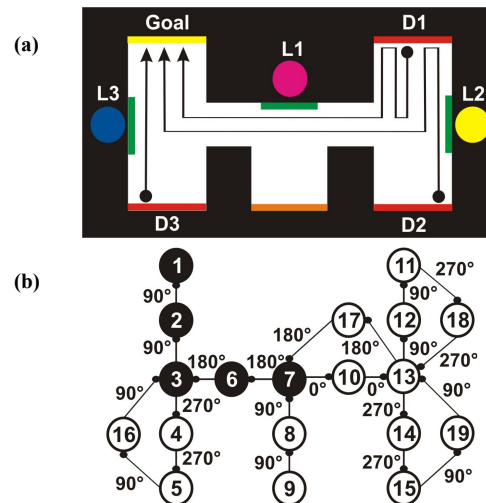


Fig. 7. (a) Routes followed by the robot in three testing trials starting from locations D1, D2 and D3. (b) The map updated by the robot during the testing phase. Black nodes belong to the learnt route and were followed by the robot to reach the goal during tests.

V. DISCUSSION AND CONCLUSIONS

In this paper we have presented a robotic navigation model based on the neurophysiology of the rat's brain. We have shown that the robot, after having explored the maze, having built a map of the environment, and having learnt the route to reach the goal starting from a fixed location, is able to recognize places already visited, distinguishing between similar places, and using the acquired knowledge to find the goal when starting from different locations in the maze.

The spatial cognition capability assigned to the robot is due to the combination of kinesthetic and landmarks information used to define every place in the explored environment. Based on its kinesthetic information, the robot can distinguish between visually similar places; e.g., places where none of the landmarks can be perceived. On the other hand, landmarks information allows the robot to recognize places already visited in a specific direction when it starts a test trial from a new departure position.

When the robot finds a node belonging to a learnt route to the goal, it is able to follow the rest of that route until it reaches the target by using local information stored in the topological map, i.e. the expectations of reward computed during training and associated to the direction of the arcs linking the nodes of the route. As shown in Fig. 7(a), when the robot started from D1 and D2, it explored the environment randomly until it found node 7 in the map (see Fig. 7(b)). Once the robot found this node, the rest is a direct path to the goal; however, the previous random exploration could last several minutes (approximately 1 to 4 in the simple maze presented) because of the intensive neuronal process the model uses.

Considering the robot's qualitative behavior, we can say that our results are similar to those obtained by Morris. Although we did not implement open field navigation, in both experiments, the rat and the robot demonstrated a place learning strategy enabling them to find the goal independently of the starting location during tests. In our case, the goal is not hidden as in Morris experiment; however, we simulated this fact by not allowing the robot to perceive the goal unless it is one step away.

At this point our experimental environment has been simplified to test the basic bio-inspired model in providing a robot with spatial cognition and goal-oriented navigation capabilities. With the work presented in this paper we are beginning to address some of the SLAM challenges previously mentioned in Section I: data association and perceptual ambiguity. A future goal is to test the model using more complex and dynamic mazes, i.e. environments having physical configuration that can change during the exploration process.

REFERENCES

- [1] D. Hähnel, W. Burgard, B. Wegbreit, and S. Thrun. "Towards lazy data association in SLAM." In Proceedings of the 11th International Symposium of Robotics Research (ISRR), Sienna, Italy, 2003.
- [2] M. O. Franz, B. Schölkopf, H. A. Mallot, H. Bülthoff. "Learning view graphs for robot navigation." *Autonomous Robots* 5: 111-125, 1998.
- [3] H. P. Movarec and A. Elfes. "High resolution maps from wide angle sonar." Proceedings of IEEE International Conference on Robotics and Automation (ICRA), pp. 116-121, 1985.
- [4] J. Guivant, E. Nebot, J. Nieto, F. Masson. "Navigation and mapping in large unstructured environments." *The International Journal of Robotics Research*, vol. 23, no. 4, pp. 449-472, 2004.
- [5] B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, F. Savelli. "Local metrical and global topological maps in the Hybrid Spatial Semantic Hierarchy." Proceedings of ICRA, New Orleans, USA, 2004.
- [6] Z. Zivkovic, B. Bakker, B. Kröse. "Hierarchical map building using visual landmarks and geometric constraints." Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Edmonton, Canada, pp. 7-12, 2005.
- [7] J. Folkesson, H. Christensen, "Graphical SLAM - A self-correcting map." Proceedings of ICRA, New Orleans, USA, 2004.
- [8] F. Savelli, B. Kuipers. "Loop-closing and planarity in topological map-building." Proceedings of IROS, pp. 1511-1517, 2004.
- [9] K. R. Beavers, W. H. Huang. "Loop closing in topological maps." Proceedings of ICRA, pp. 4378-4383, Barcelona, Spain, 2005.
- [10] I. Ulrich, I. Nourbakhsh. "Appearance-based place recognition for topological localization." Proceedings of ICRA, San Francisco, CA, pp. 1023-1029, 2000.
- [11] P. Newman, D. Cole, K. Ho. "Outdoor SLAM using visual appearance and laser ranging." Proceedings of ICRA, 2006.
- [12] Y. Bar-Shalom, T. E. Fortmann. "Tracking and Data Association." Academic Press, 1988.
- [13] M. Montemerlo, S. Thrun. "Simultaneous localization and mapping with unknown data association using FastSLAM." Proceedings of ICRA, pp. 1985-1991, 2003.
- [14] E. Tolman. "Cognitive maps in rats and men." *Psychological Review* 55, pp. 189-208, 1948.
- [15] J. O'Keefe and L. Nadel. "The hippocampus as a cognitive map." Oxford University Press, 1978.
- [16] B. Poucet. "Spatial cognitive maps in animals: new hypotheses on their structure and neural mechanisms." *Psychological Review*, 100 (2): 163-182, 1993.
- [17] T. S. Collett, B. A. Cartwright, B. A. Smith. "Landmark learning and visuo-spatial memories in gerbils." *Journal of Comparative Physiology A*, 158, 835-851, 1986.
- [18] N. Burgess, M. Recce and J. O'Keefe. "A model of hippocampal function." *Neural Networks*, Vol. 7, Nos. 6/7, pp. 1065-1081, 1994.
- [19] A. Guazzelli, F. J. Corbacho, M. Bota, and M. A. Arbib. "Affordances, motivation, and the world graph theory." *Adaptive Behavior*, vol. 6 (3/4), pp. 435-471, 1998.
- [20] A. Arleo, F. Smeraldi, W. Gerstner. "Cognitive navigation based on nonuniform Gabor space sampling, unsupervised growing networks, and reinforcement learning." *IEEE Transactions on Neural Networks*, Vol. 15, No. 3, pp. 639-652, May 2004.
- [21] D. O. Hebb. "The organization of behavior: a neuropsychological theory." Wiley-Interscience, New York, 1949.
- [22] R. S. Sutton and A. G. Barto. "Reinforcement learning: an introduction." MIT Press, Cambridge, MA, 1998.
- [23] A. Barrera, A. Weitzenfeld. "Bio-inspired model of robot adaptive learning and mapping." Proceedings of IROS, Beijing, China, October 2006.
- [24] A. Barrera, A. Weitzenfeld. "Biologically Inspired Neural Controller for Robot Learning and Mapping." Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN), Vancouver, Canada, July 2006.
- [25] R. G. M. Morris. "Spatial localization does not require the presence of local cues." *Learning and Motivation*, 12: 239 - 260, 1981.
- [26] A. Barrera and A. Weitzenfeld. "Return of the rat: biologically-inspired robotic exploration and navigation." Proceedings of the 1st IEEE / RAS-EMBS International Conference on Biomedical Robotics and Biomechanics (BioRob). Pisa, Italy. February 20 - 22, 2006.
- [27] A. Weitzenfeld, M. Arbib and A. Alexander. "The Neural Simulation Language." MIT Press, 2002.
- [28] A. Barrera and A. Weitzenfeld. "Robot experimentation results." <http://ftp.itam.mx/pub/alfredo/ABarrera/VideoCIRA07/>