

Rat-inspired Model of Robot Target Learning and Place Recognition

Alejandra Barrera¹ and Alfredo Weitzenfeld^{1,2}

¹Instituto Tecnológico Autónomo de México/Computer Engineering – Biorobotics Laboratory, Mexico City, Mexico

²University of South Florida/Computer Science Engineering – Unmanned Systems Laboratory, Tampa, USA

Abstract— 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 target learning and place recognition processes in rats as basis for topological map building and exploitation by robots. We experiment with the model in different maze configurations by training a robot to find the goal 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 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, such as data association [7] and perceptual ambiguity [8].

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 [9]. 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 catastrophic failures in SLAM algorithms, which become evident when a robot returns to a mapped region after a long excursion, the so-called loop-closure problem [6], [8].

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, have been applied also to detect loops in metric SLAM approaches [11]. Multi-hypothesis techniques 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].

Similarly 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 called cognitive maps [14] that are located in the brain area known as hippocampus [15].

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 [16]. These biological studies offer the attractive prospect of taking inspiration from animals to incorporate adaptive navigation behavioral models in mobile robots.

Taking inspiration from the rat’s hippocampus, several robotic navigation models have been proposed such as [17] that builds space representations from visual cues, and [18], [19] that combines kinesthetic with visual information based on hippocampal representation. Most of these models were tested in virtual simulations with very limited experimental work on real robots or linkage to topological or metric maps built during robot exploration.

Over the past months we have developed a rat-inspired navigation model 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. This model relies on purely kinesthetic information to identify places in the environment and has been tested in different learning and mapping experiments, with partial results reported in [20] and [21]. In this paper we present our latest work in the integration of kinesthetic and visual information derived from artificial landmarks to improve the data association process involved in the robot spatial cognition. The rest of the paper is organized as follows. Section II introduces target learning and place recognition processes in rats, Section III presents target learning and place recognition model description, Section IV describes robot experimentation results, and we conclude in Section V.

II. TARGET LEARNING AND PLACE RECOGNITION IN RATS

Rats’ capabilities to learn a route to a given target and to recognize places have been demonstrated through what is considered a “classical” neurophysiological experiment devised by Morris in 1981 [22]. Under the Morris experiment, normal rats and rats with hippocampal lesions were independently placed in a circular tank filled with an

This research is partially supported by ARO (USF), SPAWAR (USF), UC MEXUS CONACYT (ITAM – UCSC), LAFMI CONACYT (ITAM – ISC), NSF CONACYT (ITAM – UCI) and “Asociación Mexicana de Cultura, S. A.”

opaque mixture of milk and water. Rats were required to swim until they located the platform, upon which they could stand and escape from the cold water. During the experiment, the platform remained at the same location in the tank. During training, rats started from the same position in the periphery of the tank, whereas during tests, rats started from different positions. When the platform was visible, all rats were able to immediately swim towards the platform. 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 need to swim towards the platform by employing their reference memory. However, when the platform is hidden, rats need to relate the position of the platform to the location of external landmarks, thus using their spatial memory.

Recently, Hollup et al. [23] implemented variations to the Morris task by using a water maze with a circular corridor placed at the center of the tank and an equivalent land-based maze to compare place cells activity during rats' spatial learning. They concluded that hippocampal place fields in the open water maze and in the water maze restricted by a corridor are largely controlled by the same factors as on the dry land-based maze restricted by the corridor, in spite of differences in kinesthetic input.

III. A RAT-INSPIRED MODEL OF SPATIAL COGNITION

The modules included in the model are shown in Fig. 1. A detailed description of the affordances, motivation, path integration, learning and action selection model components were presented in [20], [21]. 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 information related to affordances, motivation, kinesthesia, and landmarks. Affordances are relative turns the rat can execute at any time from -180° to $+180^\circ$ in 45° intervals coded in an array of cells called affordances perceptual schema (AF).

The animal's motivation is related to its need to eat. The Fixed Critic (FC) module computes the hunger value and the immediate reward (r) the animal gets by the presence of food at any given time.

Kinesthetic information refers to the magnitude and direction of rat's movement that are 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. The output of this module is an array of cells PI that stores a kinesthetic information pattern.

Goal oriented behavior is achieved by using reinforcement learning by means of an Actor-Critic architecture, whose components, Adaptive Critic (AC) and Actor units, have been related to the circuits of the basal ganglia and associated brain structures [24]. 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.

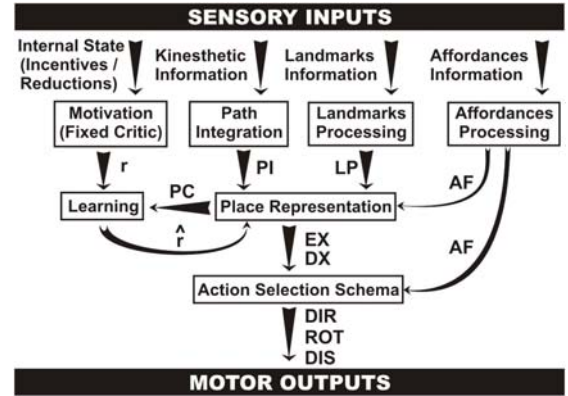


Figure 1. The spatial cognition model. r = immediate reward; PI = kinesthetic information pattern; LP = landmarks information pattern; AF = affordances perceptual schema; PC = place information pattern; \hat{r} = effective reinforcement; EX = expectations of maximum reward and their corresponding directions (DX); DIR = next rat direction; ROT = rat rotation; DIS = next rat moving displacement.

At a given location, the choice of the rat to turn to a specific direction at any time 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, and global expectation of maximum reward. SS determines the motor outputs in terms of rat direction (DIR) from 0° to 315° in 45° intervals, necessary rotation to point to 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 [25].

B. Landmarks Processing

Experiments with rodents employing colored cylinders as landmarks [26] suggested that food location is coded relative to each landmark independently, computing the distance and direction of food to each landmark. Fig. 2 shows the dynamics of the landmarks processing module of the model. Colors of cylinders are used to estimate the distance and relative orientation of each visible landmark from the rat, which are represented by two arrays of cells in a landmark perceptual schema (LPS).

In order to reduce the number of connections during processing, every neuron in LPS is randomly connected only to 50% of the neurons in a corresponding 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 coming from neuron i in LPS and the corresponding connection weight w_{ij} :

$$A_j = \sum_i LPS_i w_{ij} . \quad (1)$$

We apply Hebbian learning [27] to update connection weights between layers. To do this, we divided the population of neurons in LFDL in a number of neighborhoods including an equal amount m of cells. The

twenty most active neurons within each neighborhood are logically organized according to their activation at time t , $A_j(t)$, and identified by their place k in the hierarchy such as $OrderedSet(A_{jk}) | A_{jk}(t) \geq A_{j,k+1}(t)$ with $1 \leq k \leq 20$.

These jk neurons are associated with new activation values $G_{jk} = (20 - k + 1) / 20$ between 1 and 0.05 in 0.05 decrements. The rest $m - 20$ neurons j within the neighborhood are associated to $G_j = 0$. Activation values G_j of all n neurons j in LFDL are stored in an array L :

$$L_j = G_j | 1 \leq j \leq n. \quad (2)$$

Then, L is employed in the application of the Hebb rule:

$$\Delta w_{ij} = \alpha L P S_i w_{ij} L_j, \quad (3)$$

where α is the learning rate.

One LPS layer connected to one LFDL layer are added to the model for each landmark in the environment. All LFDL layers are combined into a single landmarks layer (LL) following the same connectivity pattern used to define the connections between any couple LPS – LFDL. We apply Hebbian learning to update connection weights between layers LFDL and LL by using an array of cells LP that stores an egocentric view from the animal. This learning rule allows the model to produce groups of neurons in LL that respond to specific information patterns derived from the integration of all landmarks.

C. Place Representation and Recognition

The dynamics of this module is shown in Fig. 3(a). To represent a place, neurons in Place Cell Layer (PCL) behave like place cells in the rat hippocampus, i.e. their activation are determined by the combination of kinesthetic and landmarks information patterns. Every neuron in PI and in LP is randomly connected to 50% of neurons in PCL. Connection weights between layers are randomly initialized and normalized between 0 and 1. The activation level A_j of a PCL unit j is computed as follows:

$$A_j = \sum_i P I_i w_{ij} + \sum_q L P_q w_{qj}, \quad (4)$$

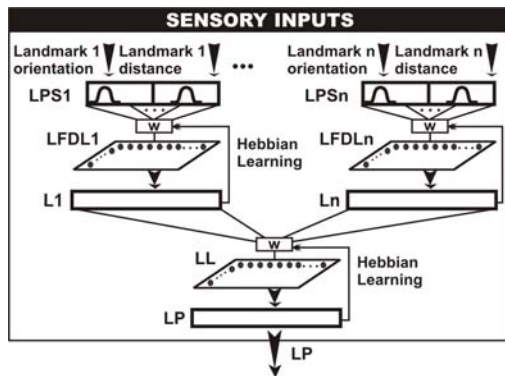


Figure 2. The landmarks processing module. LPS= landmark perceptual schema; LFDL= landmark feature detector layer; LL= landmarks layer; w = connection weights; $L1, Ln$ = landmark information patterns; LP = global landmarks pattern.

where $P I_i$ is the activation value coming from neuron i in the kinesthetic information pattern PI , $L P_q$ is the activation value coming from neuron q in the landmarks information pattern LP , and w_{ij}, w_{qj} are connection weights.

To update connection weights, activation values G_j of neurons j in PCL are stored in an array called PC , which is employed in the application of the Hebb rule in a similar way as described by (3). The activation pattern stored in PC encodes kinesthetic and visual information sensed by the rat being oriented in a given direction.

The topological map is implemented by World Graph Layer (WGL). Nodes represent places; arcs are associated with rat's direction when the animal moved from one node to the next one. The model assumes that the animal can orient itself to eight directions at any given place according to an allocentric coordinate system that is relative to the departure location in the exploration process (Fig. 3(b)). The eight activation patterns generated by PCL are stored in Actor units. Thus, every node in the map (a place) can be connected to eight Actor units (eight views).

To determine whether or not the rat recognizes a place, WGL searches the current pattern PC within the Actor units of all nodes in the map by computing the similarity degree SD between PC and every stored pattern. The 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 with the current rat's direction. Then, WGL activates or creates a node in the map depending on available affordances. In general, only if affordances at time t are different from those at time $t-1$, a new node is created and the new Actor unit is connected to it. Otherwise, WGL activates the node associated with the winner Actor unit.

Every connection between Actor units and map nodes is associated with a weight that represents the expectation of the rat to find reward when orienting itself to the direction of the Actor unit at the location corresponding to the node. To anticipate the next direction the rat should orient to get the greatest reward, WGL analyzes the weights of three nodes in sequence from the active node in the map to obtain the biggest weights EX and their corresponding directions DX . These values DX and EX are used by the action selection schema of the model to compute the global expectation of maximum reward.




Figure 3. (a) The place representation module. PCL= place cell layer; WGL= world graph layer; PI= kinesthetic pattern; LP= landmarks pattern; w = connection weights; PC= place information pattern; AF= affordances perceptual schema; \hat{r} = effective reinforcement; EX= expectations of reward and their directions (DX). (b) Allocentric coordinate system relative to the departure location.

IV. EXPERIMENTATION RESULTS

The spatial cognition model was designed and implemented using the NSL system [28]. The model 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 local 2D vision system of the robot to take at each step three non-overlapping snapshots (0° , $+90^\circ$, -90°) obtaining 45° affordance intervals.

To test the model, we implemented a spatial task inspired on Morris experiment [23] in a maze surrounded by landmarks. On the basis of the results presented by Hollup et al. [24], we supported our decision of using a land-based maze, and not implementing open-field navigation to exploit the affordances module.

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. The following subsections describe three scenarios of the experiment, keeping the same landmarks configuration in the maze during the whole task, interchanging landmarks during tests, and removing one of them. The purpose of these last two cases was to verify that the robot truly employs the position and color of landmarks to recognize different places of the maze.

A. Keeping same landmarks configuration

The experiment includes two phases, training and testing, carried out in the environment presented in Fig. 4(a), where three colored cylinders placed outside the maze represent landmarks. Colored papers pasted over the walls inside the maze were used just to compute affordances, since we have exploited only the robot head camera to detect obstacles.

Training Phase. In every trial of this phase, the robot starts from the fixed position B4 (Fig. 4(a)), and explores the maze until it finds the goal or the end of a corridor, then it returns to the departure location.

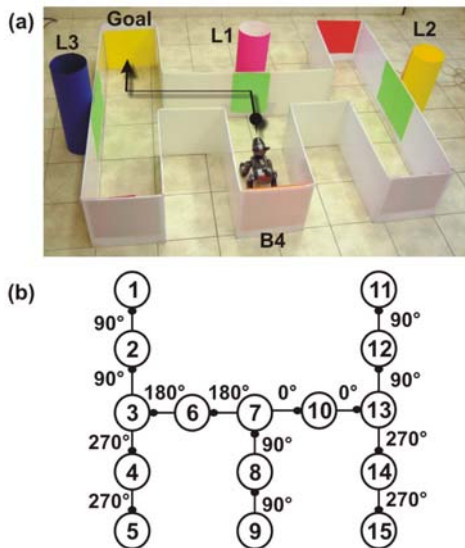


Figure 4. (a) The physical maze used in the spatial cognition experiment. AIBO robot is located at the training starting position (B4), and the path to the goal is marked with an arrow. Landmarks are labeled as L1, L2, L3. (b) Example of a map built by the robot during training.

The robot is programmed to recognize the goal just one step away from it. While the robot returns by reading arc directions in the map, the backwards reinforcement process takes place. The robot needs just one training trial finding the goal to learn the route that leads towards it from B4. After this event, the training phase is finished and the testing phase begins. An example of a map built by the robot during training is shown in Fig. 4(b). Although robot actions are performed randomly during exploration, we programmed the robot to avoid rotate 180° when there exist other possible rotations (0° , $+90^\circ$, -90°), thus accelerating the exploration process to find the goal. For these reason, arc directions between nodes in the map are one way.

Testing Phase. During tests, we placed the robot at different departure positions (B1, B2 and B3). Fig. 5(a) shows examples of routes followed by the robot from those positions in locating the goal. We should point out that for all starting positions tested, the robot found the goal successfully. The updated map of the environment after completing tests is shown in Fig. 5(b). Black nodes were recognized by the robot as belonging to the correct route to the goal. However, the robot did not recognize some places already visited although not belonging to the learnt route. Considering, for example, the robot trajectory during test starting at B1 (Fig. 5(a)), it did not recognize nodes 12 and 10 in directions 270° and 180° respectively, thus nodes 18 and 17 were created and added to the map (Fig. 5(b)). In general, the robot did not recognize some places when it was oriented in specific directions because it previously visited them in different directions, thus the landmarks configuration perceived at test time was not similar to the one stored in the corresponding nodes.




Figure 5. (a) Examples of robot trajectories in tests B1, B2 and B3. (b) The map updated by the robot during tests. Black nodes belong to the learnt route and were followed by the robot to reach the goal.

B. Interchanging landmarks

The training phase proceeds as in the original experiment just described. After the robot finds the goal, we used the same training configuration of landmarks in tests B1 and B2, but we changed it in trials B3 and B4.

As in the scenario just describe, during tests starting at B1 and B2, experimental results showed that the robot followed a direct path to the goal once it recognized a node belonging to the learnt route. Fig. 6(a) shows examples of robot trajectories during tests B3 and B4, where landmarks L2 and L3 were interchanged. The whole map resultant from this experiment is shown in Fig. 6(b), where black nodes were created during training.

In both tests B3 and B4, the robot recognized places from where landmarks L2 and L3 were not perceptible during training, such as nodes 1, 2, 7, 6, 10 and 11. On the other hand, places from where landmarks L2 and L3 were visible during training were not recognized by the robot during tests, thus adding new nodes in the map to represent them. For example, the relevant decision locations corresponding to nodes 3 and 9 were not recognized, and nodes 19 and 22, created instead during tests, were not associated with expectations of future reward in any direction, thus the robot performed random actions, reaching the goal location by chance.

C. Removing one landmark

We carried out this experiment keeping the usual training phase. During tests starting from B1 and B2, we kept the same training configuration of landmarks, whereas during tests starting at B3 and B4, we removed one of the landmarks. Fig. 7(a) shows examples of robot trajectories during tests B3 and B4, where landmark L3 was removed. The whole map resultant from this experiment is shown in Fig. 7(b).

Test starting at B3. Node 5 was not recognized due to the absence of L3; therefore, node 17 was added and connected to node 6 in direction 90° . At this time, nodes 14, 3 and 8 linked in direction 180° represented the horizontal corridor of the maze. However, since L3 was not visible anymore in this direction from any of those places, the robot did not recognize them, nodes 18, 19 and 20 were created and connected in direction 180° . The robot did not recognize node 9 in direction 180° until it chose by random to turn right orienting itself to 90° . However, it chose then to turn right again motivated by random and curiosity more than by the expectation of reward stored in the node. At this moment, connection 20-9 was defined in direction 180° , Actor 180° of node 8 was reassigned to node 20, Actor 180° of node 3 was reassigned to node 19, and the existent connection 8-9 was removed, as well as node 8 and connection 3-8 since it is not possible to reach node 9 in direction 180° from two different places. Then, the robot moved forward in its current direction 0° , and node 21 was added and connected to node 9. The robot did not recognize node 19 until it decided then to turn right twice orienting itself to 180° , thus connecting 21-19 in direction 0° . From node 9, the robot reached the goal location.

Test starting at B4. The robot did not recognize node 3 in direction 90° , but when it chose to turn left by random, it did recognize Actor 180° of node 19 and moved forward in this direction motivated by the expectation of reward inherited from node 3. At this point, connection 2-19 was defined in direction 90° , connection 2-3 was removed, Actor 0° of node 3 was reassigned to node 19, connection 3-4 was removed, connection 19-4 was added in direction 0° , node 3 and connection 14-3 were removed. From node 19, the robot finally followed the correct route to the goal.

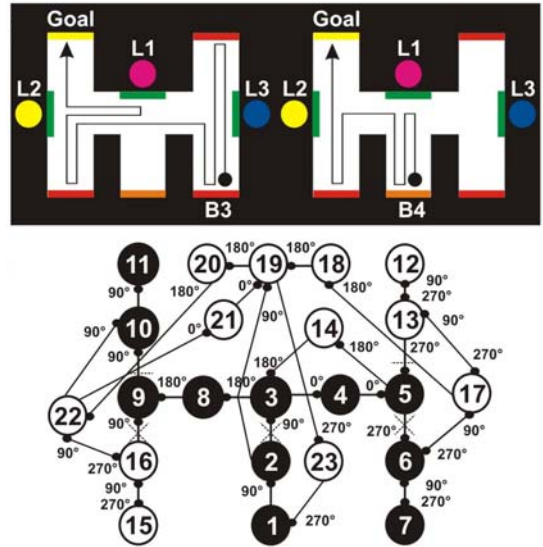


Figure 6. (a) Examples of robot trajectories in tests starting at B3 and B4, where landmarks L2 and L3 were interchanged. (b) The map built by the robot during the whole experiment. Black nodes were created during training; arcs marked as "X" were eliminated during tests.

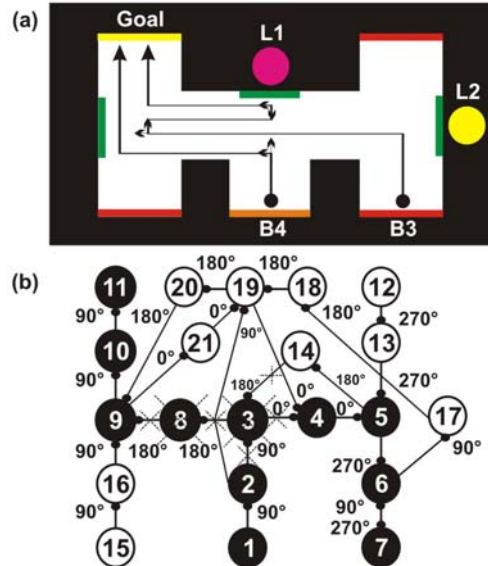


Figure 7. (a) Examples of robot trajectories during tests B3 and B4, where landmark L3 was removed. (b) The map built by the robot during the whole experiment. Black nodes were created during training, whereas nodes and arcs marked as "X" were eliminated during tests.

V. DISCUSSION AND CONCLUSIONS

We have presented a rat-inspired model of robot spatial cognition. We have shown that the robot, after having explored a maze randomly in one way during independent trials from a fixed training position, and having built a map of the environment, it is able to

- learn the correct route to reach the goal after a single training trial,
- recognize places already visited, distinguishing between places that look alike, and
- use the acquired knowledge to find the goal when starting from different locations in the maze.

The spatial cognition capability implemented on the robot is due to the combination of kinesthetic and

landmarks information used to define places in the explored environment. Based on its kinesthetic information, the robot can distinguish between places that look alike; e.g., places where none of the landmarks can be perceived, and recognize places already visited in a specific direction. Landmark information allows the robot to recognize places already visited in a specific direction when starting from a new departure point.

Target learning capability implemented on the robot results from the Actor-Critic architecture. By use of reinforcement learning, the model associates expectations of future reward to routes in the topological map, and the robot employs this local information to find the target. Specifically, when the robot starts from new different departure positions and finds a node belonging to a learnt route, it is able to follow the rest of that route until it reaches the target.

We took inspiration from Morris experiment to design the spatial task carried out by the robot during our experiments. We were particularly interested in demonstrating that the robot truly employs its “spatial memory” to recognize places through the use of position and color of external landmarks, as rats do in the Morris task. For these reason, we implemented variations in the landmarks configuration of the maze during tests by changing the position of some landmarks or removing them. According to the robot behavior, we can argue that the robot definitely relies on information of the external landmarks to recognize places and find the target. Therefore, depending on the kind of change implemented on the landmarks configuration after training, the robot can fail to find the goal location, can find it by chance building a complete different map, or find it by recognizing some places after random rotations and reassigning expectations of reward in the best scenario.

At this point, our experimental environments have 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 addressing some of the SLAM challenges in data association and perceptual ambiguity by taking a perspective inspired on rat neurophysiological studies that provides a different approach to learning and adaptability in robotics systems.

REFERENCES

- [1] Hähnel, D., Burgard, W., Wegbreit, B. and Thrun, S. “Towards lazy data association in SLAM.” Proceedings of 11th International Symposium of Robotics Research (ISRR), Sienna, Italy, 2003.
- [2] Franz, M. O., Schölkopf, B., Mallot, H. A., Bühlhoff, H. “Learning view graphs for robot navigation.” *Autonomous Robots* – 5: 111-125, 1998.
- [3] Movarec, H. P. and Elfes, A. “High resolution maps from wide angle sonar.” In Proceedings of IEEE International Conference on Robotics and Automation (ICRA), pp. 116-121, 1985.
- [4] Guivant, J., Nebot, E., Nieto, J., Masson, F. “Navigation and mapping in large unstructured environments.” *The International Journal of Robotics Research* – 23(4): 449-472, 2004.
- [5] Kuipers, B., Modayil, J., Beeson, P., MacMahon, M., Savelli, F. “Local metrical and global topological maps in the Hybrid Spatial Semantic Hierarchy.” Proceedings of International Conference on Robotics and Automation (ICRA), New Orleans, USA, 2004.
- [6] Bosse, M., Newman, P., Leonard, J., Teller, S. “SLAM in large-scale cyclic environments using the Atlas Framework.” *Journal on Robotics Research* – 23(12):1113–1139, 2004.
- [7] Folkesson, J., Christensen, H. “Graphical SLAM - A self-correcting map.” In Proceedings of IEEE International Conference on Robotics and Automation (ICRA), New Orleans, USA, 2004.
- [8] Savelli, F., Kuipers, B. “Loop-closing and planarity in topological map-building.” Proceedings of International Conference on Intelligent Robots and Systems (IROS), pp. 1511-1517, 2004.
- [9] Frese, U. “A discussion of Simultaneous Localization and Mapping.” *Autonomous Robots* – 20: 25-42, 2006.
- [10] Ulrich, I., Nourbakhsh, I. “Appearance-based place recognition for topological localization.” In Proceedings of IEEE International Conference on Robotics and Automation (ICRA), San Francisco, CA, pp. 1023-1029, 2000.
- [11] Newman, P., Cole, D., Ho, K. “Outdoor SLAM using visual appearance and laser ranging.” Proceedings of International Conference on Robotics and Automation (ICRA), Florida, 2006.
- [12] Bar-Shalom, Y., Fortmann, T. E. “Tracking and Data Association.” Academic Press, 1988.
- [13] Montemerlo, M., Thrun, S. “Simultaneous localization and mapping with unknown data association using FastSLAM.” In Proceedings of IEEE International Conference on Robotics and Automation (ICRA), pp. 1985-1991, 2003.
- [14] Tolman, E. “Cognitive maps in rats and men.” *Psychological Review* – 55: 189-208, 1948.
- [15] O’Keefe, J. and Nadel, L. “The hippocampus as a cognitive map.” Oxford University Press, 1978.
- [16] Jeffery, K. J., O’Keefe, J. M. “Learned interaction of visual and idiothetic cues in the control of place field orientation.” *Experimental Brain Research* – 127: 151-161, 1999.
- [17] Burgess, N., Recce, M. and O’Keefe, J. “A model of hippocampal function.” *Neural Networks* – 7(6/7): 1065-1081, 1994.
- [18] Guazzelli, A., Corbacho, F. J., Bota, M. and Arbib, M. A. “Affordances, motivation, and the world graph theory.” *Adaptive Behavior* – 6(3/4): 435-471, 1998.
- [19] Arleo, A., Smeraldi, F., Gerstner, W. “Cognitive navigation based on nonuniform Gabor space sampling, unsupervised growing networks, and reinforcement learning.” *IEEE Transactions on Neural Networks* – 15(3): 639-652, 2004.
- [20] Barrera, A. Weitzenfeld, A. “Bio-inspired model of robot adaptive learning and mapping.” Proceedings of IROS, Beijing, 2006.
- [21] Barrera, A. and Weitzenfeld, A. “Biologically Inspired Neural Controller for Robot Learning and Mapping.” Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN), Vancouver, 2006.
- [22] Morris, R.G. M. “Spatial localization does not require the presence of local cues.” *Learning and Motivation* – 12: 239 – 260, 1981.
- [23] Hollup, S. A., Molden, S., Donnett, J. G., Moser, M. and Moser, E. I. “Place fields of rat hippocampal pyramidal cells and spatial learning in the watermaze.” *European Journal of Neuroscience* – 13: 1197-1208, 2001.
- [24] Barto, A. G. “Adaptive critics and the basal ganglia,” in Models of information processing in the basal ganglia, edited by J. C. Houk, J. L. Davis and D. Beiser, MIT Press, Cambridge, MA, pp. 215-232, 1995.
- [25] Barrera, A. and Weitzenfeld, A. “Return of the rat: biologically-inspired robotic exploration and navigation.” In Proceedings of the IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechanics (BioRob), Pisa, Italy, 2006.
- [26] Collett, T. S., Cartwright, B. A., Smith, B. A. “Landmark learning and visuo-spatial memories in gerbils.” *Journal of Comparative Physiology A* – 158: 835-851, 1986.
- [27] Hebb, D. O. “The organization of behavior: a neuropsychological theory.” Wiley-Interscience, New York, 1949.
- [28] Weitzenfeld, A., Arbib, M. and Alexander, A. “The Neural Simulation Language.” MIT Press, 2002.