

Computational Modeling of Spatial Cognition in Rats and Robotic Experimentation: Goal-Oriented Navigation and Place Recognition in Multiple Directions

Alejandra Barrera, *IEEE Member*, and Alfredo Weitzenfeld, *IEEE Senior Member*

Abstract—A computational model of spatial cognition in rats is used to control an autonomous mobile robot while solving a spatial task within a cyclic maze. In this paper we evaluate the robot’s behavior in terms of place recognition in multiple directions and goal-oriented navigation against the results derived from experimenting with laboratory rats solving the same spatial task in a similar maze. We provide a general description of the bio-inspired model, and a comparative behavioral analysis between rats and robot.

I. INTRODUCTION

Two fundamental neuroscientific theories support the fact that rats are able to generate a spatial representation of any explored environment: the cognitive map by Tolman [1], and the hippocampus as a cognitive map by O’Keefe and Nadel [2].

In the hippocampus, neurons called place cells increase the frequency of action potential discharge when the animal is in a specific physical region of the environment, which defines the place field of the cell. The representation encoded by place cells integrates visual cues with kinesthetic feedback information in order to recognize places already visited [3].

In goal-oriented behaviors, rats are able to learn spatial tasks by associating rewards with locations in the environment, or rewards with procedures [2]. The place learning strategy depends on the cognitive mapping system provided by the hippocampus, and allows the animal to relate the goal location with the location of external visual cues, recognize the goal location within the cognitive map, and navigate towards it. The procedure learning strategy, on the other hand, depends on the reward system of the brain attributed to the striatum [4], and the rat manifests it by learning the procedure that leads to the target in terms of rotations and displacements relative to its body midline.

The purpose of the research that our group has been carrying out over the past months consists on identifying the mechanisms underlying learning and memory in order to computationally model rats’ abilities such as the generation

of the cognitive map, learning and unlearning of goal locations, and map exploitation during navigation to goal locations from any given point of departure.

Our computational model of spatial cognition and navigation in rats (i) relies on the integration of kinesthetic and visual information and on the use of Hebbian learning [5] to build a topological-metric spatial representation during the environment exploration, (ii) considers the internal motivational state of the animal to implement reinforcement learning through an Actor-Critic architecture [6] to enable learning and unlearning of goal locations in their neighborhood within the environment, and (iii) complements this learning architecture with a backward reinforcement method applied to the route followed by the animal during every trial in a given experiment. In this method, the adaptation of reward expectations associated to appropriate directions at different locations of the route enables learning and unlearning of routes within environments including several choice points, where the goal is perceptible only at its location. Once the rat has explored and learnt the environment, the exploitation of maximum expectations of reward stored in the nodes of the cognitive map allows the animal to reach the goal following a direct route from any given departure location. Additionally, the model implements an on-line adaptation process of the cognitive map to situations where changes in the environment physical configuration occur during exploration.

While experimenting with the model, we have tested the learning and unlearning of reward locations obtaining results derived from implementing with robots the classical spatial reversal task carried out by O’Keefe [7] with normal and hippocampal-lesioned rats in a T-maze and in an 8-arm radial maze. In [8] we show the behavioral similarity between robot and rat, and also present results from experimenting with the robot in a spatial task that involves a similar protocol used in a classical goal-oriented navigation experiment devised by Morris [9], employing in our case a multiple T-maze surrounded by landmarks. In [10] we discuss results from testing the robot in that multiple T-maze when changing the physical configuration of landmarks after having trained the robot to find the goal.

The current paper extends this work by evaluating the robot’s behavior while solving a spatial task within a cyclic maze in terms of place recognition in multiple directions and

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 faculty members in the Computer Engineering Department at Instituto Tecnológico Autónomo de México. Río Hondo #1, Tizapán San Ángel, CP 01000, Mexico City, Mexico. (e-mail: abarrera@itam.mx, alfredo@itam.mx).

goal-oriented navigation against the results derived from our experiments with laboratory rats solving the same spatial task in a similar maze.

Section II of the paper describes distinctive aspects of our approach against existing models, Section III introduces our bio-inspired spatial cognition model, Sections IV and V present results obtained from rats and robots solving the spatial task in the cyclic maze, Section VI analyzes these results comparatively, and we conclude in Section VII.

II. RELATED WORK

Taking inspiration from rat's spatial cognition system, several robotic navigation models have been proposed such as [11]-[15]. The main distinctive aspects of our approach include: (i) generation of a holistic topological-metric map by considering activity patterns derived from the complete population of place units to define distinctive places and their relationships instead of determining isolated place fields as a spatial representation, (ii) modeling of the rat's unlearning ability, (iii) modeling of the interaction between the hippocampus and the striatum to allow the animal to exploit expectations of future reward during reinforced spatial tasks, (iv) a map exploitation process to enable goal-directed navigation, (v) a map adaptation process that permits on-line representations of changes in the physical configuration of the environment, and (vi) validation of our robotic architecture not limited to test well-known spatial tasks performed with rats, but also by designing and implementing new experiments with rats. Refer to [8] for further detail on the comparative analysis between our model and the abovementioned approaches.

III. A BIO-INSPIRED MODEL OF SPATIAL COGNITION

The model comprises distinct functional modules shown in Fig. 1 that capture some properties of rat brain structures involved in learning and memory as follows:

Motivation. This module is related to the lateral hypothalamus controlling food seeking and food intake [16]. The module computes the value of the rat's hunger drive and produces the primary reward the animal gets by the presence of food (r).

Path Integration. The posterior parietal cortex (PPC) is part of a neural network mediating path integration [17], where the retrosplenial cortex (RC) is also involved [18]. We attribute to PPC the representation of the updated position of the rat's point of departure each time the animal moves in relation to its current position through a dynamic remapping perceptual schema (DR), and to RC, the generation of kinesthetic information patterns (PI) carried out by a path integration feature detector layer in our model.

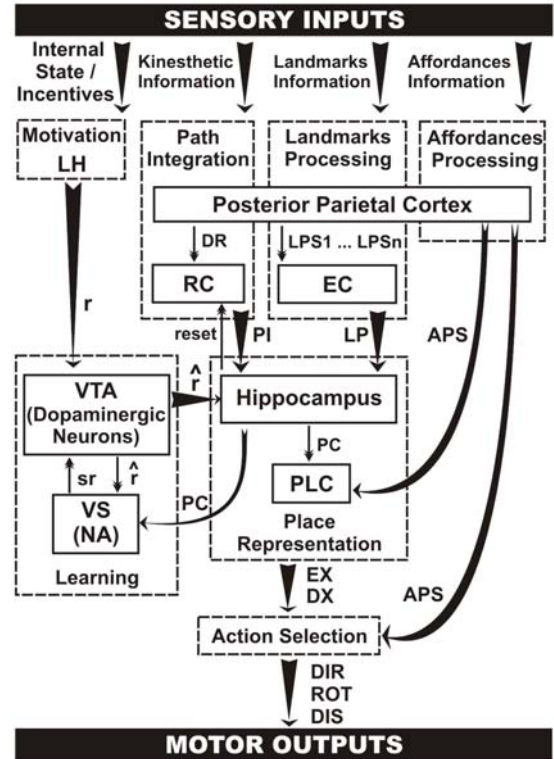


Fig. 1. The computational spatial cognition model and its underlying biological framework. Glossary: LH – Lateral Hypothalamus; RC – Retrosplenial Cortex; EC – Entorhinal Cortex; VTA – Ventral Tegmental Area; VS – Ventral Striatum; NA – Nucleus Accumbens; PLC – Prelimbic Cortex. Inputs/Outputs: r = primary reinforcement; sr = secondary reinforcement; \hat{r} = effective reinforcement; DR = dynamic remapping perceptual schema; LPS = landmark perceptual schema; APS = affordances perceptual schema; PI = kinesthetic information pattern; LP = landmarks information pattern; PC = place information pattern; EX = expectations of maximum reward and their corresponding directions (DX); DIR = next rat direction; ROT = rat rotation; DIS = next rat moving displacement.

Landmarks Processing. The entorhinal cortex (EC) is involved in landmarks processing. EC receives spatial information about landmarks from PPC [12], i.e. distance and relative orientation of each landmark encoded in a landmark perceptual schema (LPS), and then, landmarks information patterns are produced and integrated in a single pattern representing the egocentric view from the rat (LP).

Affordances Processing. It has been suggested that, preceding the rat's motion, nearly half of the cells in PPC exhibit movement-related activity discriminating among basic modes of locomotion: left turns, right turns, and forward motion [19]. We attribute to PPC the generation of the affordances perceptual schema (APS) encoding possible turns the rat can perform at any given time being at a specific location oriented to a certain direction.

Place Representation. This module comprises a place cell layer (PCL) and a world graph layer (WGL). The hippocampus receives kinesthetic and visual information from RC and EC respectively. The activity of place cells results from the integration of both information sources. Overlapping place fields in the collection of neurons in PCL are associated with a physical area in the rat's environment that is identified directionally by the ensemble activity

pattern (*PC*). These associations are represented by WGL through a holistic spatial map. WGL also performs place recognition, and we assume that its functionality could be corresponded to the prelimbic cortex, which is involved in the control of goal-directed behaviors [20]. The activation patterns generated by PCL when the animal is oriented to diverse directions are stored in Actor units. Thus, every node in the map (a place) can be connected to several Actor units (different views), one for each direction, and every connection is associated with a weight (representing the expectation of getting a reward when orienting to the particular Actor unit direction at the current location), and an eligibility trace (marking the connection as eligible to be reinforced later in time).

Learning. Reward information is processed in the basal ganglia by its dopaminergic neurons, which respond to primary and secondary rewards, and their responses can reflect “errors” in the prediction of rewards, thus constituting teaching signals for reinforcement learning. Neurons in the ventral striatum (nucleus accumbens) are activated when animals expect predicted rewards, and adapt expectation activity to new reward situations [4]. Houk et al. [21] proposed that the striatum implements an Actor-Critic architecture [6], in which an Adaptive Critic predicts reward values of any given place in the environment and produces the error signal (\hat{r}), and a number of Actor units are included representing possible actions to be performed by the rat. In our model, reward expectations associated to Actor units are adapted by means of \hat{r} . Recently, it has been suggested that rats with lesions of the hippocampal dentate gyrus (DG) are severely impaired in their ability to update spatial information essential to guide goal-oriented behaviors [22]. Thus, we suppose that Actor units could be located in DG, and suggest that the striatum should influence the hippocampus through DG by sending the expectations of future reward corresponding to the animal’s actions that are exploited by DG to allow the appropriate performance of the animal during reinforced spatial tasks.

Action Selection. The rat motor decision is determined by considering all possible rotations to execute from current location and orientation, the curiosity to execute rotations not yet explored, and the expectations of maximum reward. To enable goal-directed navigation, the module implements a backward reinforcement method where the eligibility traces of the Actor units are updated in the direction of the arcs connecting the nodes in the path followed by the rat. The strategy involved in this method is based on the existence of a factor referred to as goal gradient by Hull [23], according to which the reinforcement effect is the most at the goal location and diminishes progressively as the animal moves backward through the maze. Refer to [8] for further detail on the formal specification of the model.

IV. ANIMAL EXPERIMENTATION RESULTS

We employed a group composed of three male rats to

carry out the experiment in the maze shown in Fig. 2. This is a maze with 95 x 105 cm of area and six internal corridors. It is surrounded by a white curtain over which four different colored geometrical figures were disposed to represent spatial cues or landmarks.

In order to motivate the animals to learn the spatial task, the consumption of water was restricted during the previous week of the experiment, and absolutely prohibited during the prior 24 hours. The tip of a water dispenser was placed at the goal location indicated in Fig. 2 to provide sweet water to rats when reaching this location at the end of any experiment’s trial. The solution used as the reinforcement consisted on 5 gr. of sucrose dissolved in 50 ml. of water, and the water dispenser was filled with 400 ul of the solution before beginning every trial. Rats were independently trained and tested as described in the following subsections.




Fig. 2. Top view of the cyclic maze employed in the experiment with rats. Landmarks L1, L2, L3 and L4 are illustrated, as well as locations TD, D1, D2, D3, D4, and the goal location.

A. Training Procedure

At the beginning of every training trial, the rat is placed at location TD (Fig. 2) being oriented to north (90°), and the trial consists on exploring the maze freely until reaching the goal location and drinking the sweet water. Then, the rat is removed from the maze.

The training procedure concluded when the animals reached the learning criterion that consisted on stabilizing the latency (duration) of a trial. In this way, each rat was trained during 35 trials distributed in six sessions throughout three days. Fig. 3 shows the average latency of each trial considering the performance of three rats.

During the first four sessions, rats followed indirect routes from location TD to the goal (i.e. routes including cycles). From session 5, some trials consisted on direct although not optimal routes. Particularly, the rat that best learned the task followed a direct and optimal path in three trials of the fifth session and all trials of the sixth, whereas the other two animals even failed in two and three trials of the last session. The average latency of each trial was 122 sec, and the average duration of the entire training procedure was 72 min per rat.




Fig. 3. Average latency of each training trial considering the performance of three rats during the spatial task in the cyclic maze. The average latency at the end of the training procedure was 12 sec.

B. Testing Procedure

One day after having concluded the training procedure, rats were tested during 16 trials being liberated from four different locations in the maze: D1, D2, D3, and D4 (Fig. 2), pointing to four different orientations: north, south, east, and west. In every testing trial, the rat is placed at the corresponding location and orientation, and it is required to explore the maze freely until reaching the goal and drinking the sweet water. Then, the rat is removed from the maze.

During tests, latencies and routes followed by the animals in locating the water were registered. The performance of the rat that best learned the task is illustrated in Fig. 4, showing that the rat not always follows the shortest path to the goal. Specifically, when departing from location D1 being oriented to east, the rat followed an indirect route passing through the same places twice, and when departing from D2 being oriented to north, it executed a direct though not optimal route. Successful cases following optimal routes included the departure from D3 and D4 towards any direction. When the animal was liberated from D1 and from D2 (in some trials), it moved towards location TD and executed the same trajectory learnt during training. This behavior suggests that the rat employed a procedure learning strategy. However, the optimal routes followed from D3 and D4 suggest a place learning strategy and the exhibition of the goal gradient proposed by Hull [23].

V. ROBOT EXPERIMENTATION RESULTS

The rat cognitive model was designed and implemented using the NSL system [24]. The model can interact with a real robotic environment through an external visual processing module that receives three non-overlapping snapshots taken by the robot at each step using its local camera, and a motor control module that executes rotations and/or translations on the robot. Refer to [8] for further detail on the robotic implementation of the model.

The AIBO robot carried out the experiment in the maze shown in Fig. 5. This is a maze with 180 x 210 cm of area and six internal corridors. It is surrounded by four colored cylinders representing spatial cues or landmarks. During the experiment, the robot passed through three phases as described in the following subsections.

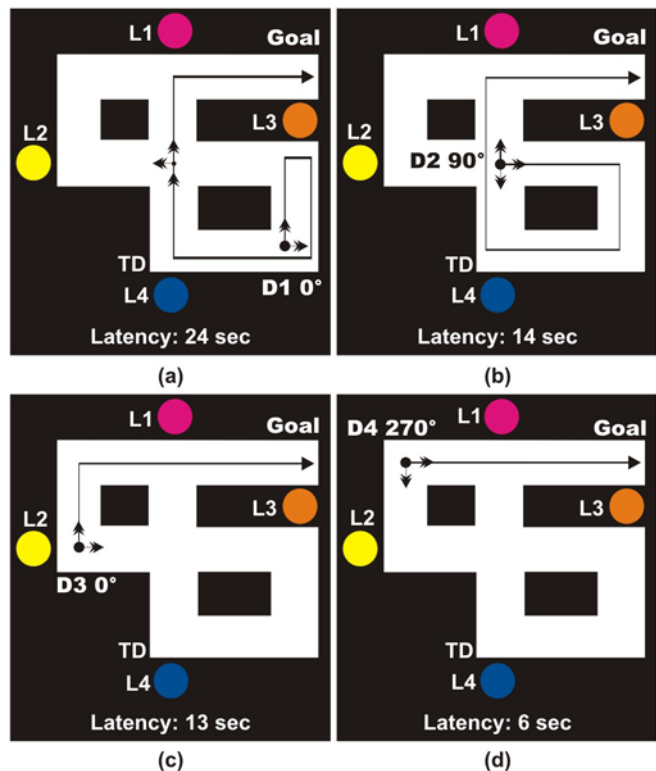


Fig. 4. Routes followed by the rat that best learned the task in locating the water dispenser departing from locations (a) D1 to east (0°), (b) D2 to north (90°), (c) D3 to east (0°), and (d) D4 to south (270°). Small arrows over the routes represent head or body turns executed by the animal.




Fig. 5. Top view of the cyclic maze employed in the experiment with an AIBO robot. Landmarks L1, L2, L3 and L4 are illustrated, as well as locations TD, D1, D2, D3, D4, and the goal location HE.

A. Habituation Procedure

The experiment with the robot begins with a habituation procedure where it departs from location TD (Fig. 5) and explores the maze freely until reaching location HE. No reinforcement is provided to the robot while exploring the maze, thus it does not carry out any learning process of goal locations, and its actions are determined just by curiosity and randomness. The duration of this procedure was variable considering three model executions. In average, the robot took 135 min to reach location HE from TD. As a result of habituation, the robot builds a cognitive map of the maze including nodes to represent the explored locations, and Actor units associated with the different directions in which they were explored. The map generated in the best of the model executions represents all maze locations in all possible directions as illustrated in Fig. 6.




Fig. 6. The cognitive map generated by the robot during the habituation procedure of the spatial task in the cyclic maze. Nodes are numbered in order of creation, and bidirectional arcs between nodes are labeled with the all-centric robot's direction when it moved from one node to the next one.

B. Training Procedure

After habituation, the reinforcement is introduced at location HE of the maze, the robot is placed at TD being oriented to north (90°), and the training procedure begins from the cognitive map previously generated.

The robot perceives the same visual and kinesthetic information sensed during habituation, thus the activity patterns registered by place cells in the model are similar to those stored in the map, which allows the robot to recognize all explored locations. In every trial, the robot explores the maze from TD until finding the goal location. During exploration, it marks the links between explored nodes in the corresponding direction as many times as it visits them. In the model, the map is analyzed at the end of the trial, and two situations may occur: (i) if the robot visited at least one node more than once, the route from TD to the goal is not reinforced since it represents an indirect path that includes cycles; otherwise, (ii) if the robot visited each node just once, the backward reinforcement process is applied to each map node belonging to the direct route.

During training, the curiosity is not anymore a factor that determines the actions of the robot. Nevertheless, the random factor is still present promoting that the robot follows indirect routes to the goal, even more in regions located faraway from the target, and while the robot's reward expectations are not big enough to exceed the randomness. This led to long training trials, thus we decided to adjust the model parameters related to the backward reinforcement process to allow the robot to learn a direct route from TD to the goal in just one trial executing it. In three model executions, eight trials were carried out lasting 9 min in average, and in the best case, the robot explored all possible direct routes from TD to the goal.

C. Testing Procedure

The robot was tested during several trials departing from different locations and orientations. Fig. 5 illustrates locations D1, D2, D3, and D4, from where the robot departs

being oriented to north (90°), north (90°), east (0°) and south (270°), respectively. For each position, Fig. 7 shows the best route followed by the robot in locating the goal.

Due to the exhaustive exploration performed during habituation, the robot was able to recognize all places previously visited in the same direction, thus it did not need to add any new node to the cognitive map. The recognition process allowed the robot to exploit information stored in Actor units to reach the goal location by executing rotations associated with maximum reward expectations.

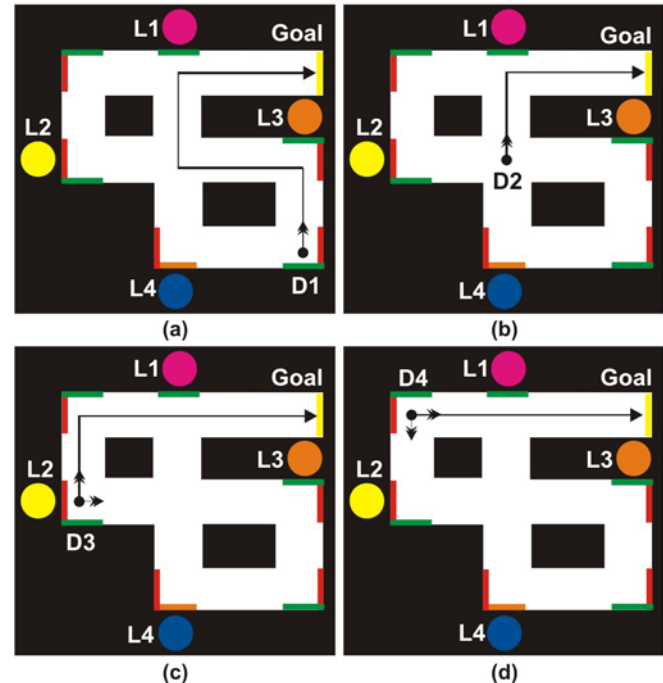


Fig. 7. Routes followed by the robot in locating the goal departing from locations (a) D1 to north (90°), (b) D2 to north (90°), (c) D3 to east (0°), and (d) D4 to south (270°). Small arrows over the routes represent body turns executed by the robot. Latencies of these testing trials were 150 sec from D1, 80 sec from D2, 130 sec from D3, and 90 sec from D4.

VI. COMPARATIVE BEHAVIORAL RESULTS

Although the experiment with rats did not include an explicit habituation phase as occurred with the robot, they habituated to the maze naturally during long trials registered in earliest training sessions. The habituation procedure allowed both, rats and robot, to improve notably their performance during training and tests enabling them to recognize maze locations in multiple directions.

The training average duration was the same for rats and robot (72 min) though different number of trials (8 robot trials vs. 35 rat trials), since one rat trial lasts 2 min vs. 9 min of one robot trial. As a result of training, the latency of rat trials stabilizes, and the robot learns different direct routes toward the goal from the designated fixed location.

During tests, the direct and optimal robot trajectories from D3 and D4 coincide with those executed by the rat that best learned the task, whereas considering the cases from D1 and D2, the rat followed a long way to the water dispenser, and the robot executed an optimal route to the goal. As expected,

the latency of the robot testing trials was bigger than that reported by the rat since the robot employs 10 sec just in taking snapshots of the environment at each step.

Results derived from experimentation with rats do not reveal an efficient performance in all tested situations. As abovementioned, the rat executed indirect routes in some cases where it visited places twice even after having learnt the task, thus suggesting the use of a procedure learning strategy. On the other hand, although the robot employed a place learning strategy during training adding information to the cognitive map related to reward expectations, the successful robot's behavior reported during tests does not guarantee an efficient performance in every situation either, since its performance depends on a cognitive map built and maintained precisely during the habituation and training procedures of the task. Specifically, situations where the robot does not explore the environment exhaustively during habituation, or does not visit all possible direct routes from the fixed departure location to the goal during training, will involve (i) the addition of new nodes to the map affecting the place recognition process during tests and hence the exploitation of the maximum reward expectations that lead to the goal directly, and (ii) the random navigation of the robot until recognizing a node in the map belonging to the route or routes learnt. Table I summarizes the average behavioral results derived from rat and robot.

TABLE I
AVERAGE BEHAVIORAL RESULTS FROM RAT AND ROBOT

Training	Rat		Robot	
Number of trials	35		8	
Trial latency (sec)	122		540	
Total latency (sec)	4270		4320	
Testing	Trial D1		Trial D2	
	Rat	Robot	Rat	Robot
Latency (sec)	24	150	14	80
Route quality	Indirect	Optimal	Direct	Optimal
Learning strategy	Procedure	Place	Procedure	Place
Testing	Trial D3		Trial D4	
	Rat	Robot	Rat	Robot
Latency (sec)	13	130	6	90
Route quality	Optimal	Optimal	Optimal	Optimal
Learning strategy	Place	Place	Place	Place

VII. CONCLUSIONS AND DISCUSSION

As described throughout the paper, we are concerned with studying the underlying mechanisms of rats' spatial cognition, incorporating relative physiological data in a robotic architecture, and evaluating it against biological experimentation results. In this way, we have initiated a contribution to experimental neuroscience by providing our system as a tool to test new hypotheses with robots.

According to the robot experimentation results presented in this paper, the exhaustive exploration during habituation and training required so much time, although we employed a small maze during the spatial task. However, the exhaustive exploration was the factor that enabled the successful robot performance, thus we will be working on reducing the exploration time by improving the visual processing the

robot carries out when getting a view of the environment at each step. With this improvement we expect to get a better robot training latency in relation to that reported by rats, considering that the robot needs a smaller amount of trials to learn any spatial task.

According to the rat experimentation results, on the other hand, we expected that rats manifested a place learning strategy during the task since we provided enough, big and stable landmarks around the maze. Nevertheless, some of the tests with rats suggested a procedure learning strategy, thus we will continue experimenting on the same spatial task with other group of rats in order to identify precisely the conditions that determine the learning strategy followed by the animals.

Our model of the hippocampal function, as other existing computational models, does not include the current perspective relative to functional differences between hippocampal substructures CA1, CA3 and DG [25]. We will do it as future work in order to produce a more realistic model that helps to verify those new neuroscientific concepts through robotic experimentation, and whose results might suggest new theoretical or experimental alternatives that promote a deeper development of our understanding about the hippocampal function.

REFERENCES

- [1] Tolman, E., "Cognitive maps in rats and men," *Psychological Review* 55, 189-208, 1948.
- [2] O'Keefe, J., and Nadel, L., *The hippocampus as a cognitive map*, Oxford University Press, 1978.
- [3] 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.
- [4] Schultz, W., Tremblay, L., Hollerman, J., "Reward prediction in primate basal ganglia and frontal cortex," *Neuropharmacology* 37(4-5): 421-429, 1998.
- [5] Hebb, D. O., *The organization of behavior: a neuropsychological theory*, Wiley-Interscience, New York, 1949.
- [6] 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.
- [7] O'Keefe, J. "Spatial memory within and without the hippocampal system," in *Neurobiology of the Hippocampus*, edited by W. Seifert, Academic Press, New York, pp. 375 - 403, 1983.
- [8] Barrera, A., Weitzenfeld, A., "Biologically-inspired Robot Spatial Cognition based on Rat Neurophysiological Studies," *Autonomous Robots*, Springer Netherlands, DOI 10.1007/s10514-007-9074-3, 2008.
- [9] Morris, R. G. M., "Spatial localization does not require the presence of local cues," *Learning and Motivation* - 12: 239 - 260, 1981.
- [10] Barrera, A., Weitzenfeld, A., "Rat-inspired Model of Robot Target Learning and Place Recognition," *Proceedings of the Mediterranean Conference on Control and Automation (MED)*, Athens, Greece, 2007.
- [11] Burgess, N., Recce, M., O'Keefe, J., "A model of hippocampal function," *Neural Networks* - 7(6/7): 1065-1081, 1994.
- [12] Redish, A., Touretzky, D., "Cognitive maps beyond the hippocampus," *Hippocampus* 7(1): 15-35, 1997.
- [13] Guazzelli, A., Corbacho, F. J., Bota, M., Arbib, M. A., "Affordances, motivation, and the world graph theory," *Adaptive Behavior* - 6(3/4): 435-471, 1998.
- [14] 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.
- [15] Milford, M., Wyeth, G., Prasser, D., "RatSLAM on the edge: revealing a coherent representation from an overloaded rat brain," *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Beijing, China, pp. 4060 – 4065, 2006.
- [16] Kelley, A., "Ventral striatal control of appetitive motivation: role in ingestive behavior and reward-related learning," *Neuroscience and Biobehavioral Reviews* 27(8): 765-776, 2004.
- [17] Parron, C., Save, E., "Evidence for entorhinal and parietal cortices involvement in path integration in the rat," *Experimental Brain Research* 159(3): 349-359, 2004.
- [18] Cho, J., Sharp, P., "Head direction, place, and movement correlates for cells in the rat retrosplenial cortex," *Behavioral Neuroscience* 115(1): 3-25, 2001.
- [19] McNaughton, B., Mizumori, S., Barnes, C., Leonard, B., Marquis, M., Green, E., "Cortical representation of motion during unrestrained spatial navigation in the rat," *Cerebral Cortex* 4: 27-39, 1994.
- [20] Granon, S., Poucet, B., "Involvement of the rat prefrontal cortex in cognitive functions: a central role for the prelimbic area," *Psychobiology* 28(2): 229-237, 2000.
- [21] Houk, J. C., Adams, J. L., Barto, A. G., "A model of how the basal ganglia generate and use neural signals that predict reinforcement," in *Models of information processing in the basal ganglia*, edited by J. C. Houk, J. L. Davis and D. G. Beiser, MIT Press, pp. 249-270, 1995.
- [22] Hernández-Rabaza, V., Barcia, J., Llorens-Martín, M., Trejo, J., Canales, J., "Spared place and object-place learning but limited spatial working memory capacity in rats with selective lesions of the dentate gyrus," *Brain Research Bulletin* 72(4-6): 315-323, 2007.
- [23] Hull, C. L., "The goal gradient hypothesis and maze learning," *Psychological Review* 39: 25-43, 1932.
- [24] Weitzenfeld, A., Arbib, M., Alexander, A., *The Neural Simulation Language*, MIT Press, 2002.
- [25] Leutgeb S, Leutgeb JK., "Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map," *Learning and Memory* 14(11):745-57, 2007.