

Spatial Memory and Learning: Towards a Set of Metrics to Evaluate Task Performance in Rats and Robots

Alejandra Barrera, *Member, IEEE*, Alfredo Weitzenfeld, *Senior Member, IEEE*, Alejandra Cáceres, and Victor Ramirez-Amaya

Abstract— The study of spatial memory and learning in rats has inspired the development of multiple computational models that have led to novel robotics architectures. Evaluation of computational models and resulting robotic architecture is usually carried out at the behavioral level by evaluating experimental tasks similar to those performed with rats. While multiple metrics are defined to evaluate behavioral performance in rats, metrics for robot task evaluation are very limited mostly to success/failure and time to complete task. In this paper we present a set of metrics taken from rat spatial memory and learning evaluation to further analyze performance in robots. The proposed set of metrics, learning latency and ability to navigate minimal distance to goal, should offer the robotics community additional tools to assess performance and validity of models in biologically-inspired robotic architectures at the task performance level. We also provide a comparative evaluation using these metrics between similar spatial tasks performed by rat and robot in comparable environments.

I. INTRODUCTION

As argued by Webb [15], in addition to have robots modeled *on* animals by using the biology as a source of inspiration, robots can also serve as models *of* animals to test biological hypotheses. And although nature is much more complex than robotics, robots can offer an experimental platform to test hypotheses about animal behavior without the need of extended training required in animals that may involve weeks of training and testing just for a single task evaluation.

In the past, studies of spatial memory and learning in rats, i.e. spatial cognition, have inspired various computational models and robotics systems, such as Redish and Touretzky [14], Guazzelli et al. [7], Arleo and Gerstner [1], Arleo et al. [2], Milford and Wyeth [10], and Barrera and Weitzenfeld [4]. In general, development of these models require: (a)

Alejandra Barrera is with the Computer Engineering Department at the Instituto Tecnológico Autónomo de México (ITAM), Mexico City, 01000, México (e-mail: abarrera@itam.mx).

Alfredo Weitzenfeld is with the Information Technology Department at University of South Florida Polytechnic, Lakeland, FL, 33180, USA (corresponding autor: tel 1-863-6677769; fax 1-863-6677752; e-mail: aweitzenfeld@poly.usf.edu).

Alejandra Cáceres is with the Plastic Neural Networks Laboratory, Neurobiology Institute, Universidad Nacional Autónoma de México (UNAM), Querétaro, 76230, México.

Victor Ramirez-Amaya is with the Plastic Neural Networks Laboratory, Neurobiology Institute, Universidad Nacional Autónoma de México (UNAM), Querétaro, 76230, México (email: ramirez@inb.unam.mx).

studying rat brain regions at different levels of granularity, e.g. behavioral and electrophysiological; (b) design of a computer-based model at the corresponding level of granularity; (c) testing of model initially by computer simulation and then by robot experimentation; and (d) evaluation of results obtained from robot experimentation against original rat studies. While various models of spatial cognition have been developed and tested in robots, most of these models do not “close the loop” in terms of a final evaluation against similar experiments in rats as described in (d). In this paper, we present metrics for such a comparative study between similar spatial memory and learning tasks performed in rats and robots. In prior papers, Barrera and Weitzenfeld [3][4] present a computational model and robotics architecture inspired on behavioral and electrophysiological studies of hippocampus and striatum brain regions in rats. Barrera and Weitzenfeld [3][4] also describe and contrast experimental results obtained from “normal” and “hippocampus-lesioned” models, and corresponding robotic experiments, inspired on classical rat spatial tasks including the T-maze and 8-arm radial maze as described by O’Keefe [12] and variations of the Water Maze by Morris [11]. In this paper we go beyond prior work by incorporating evaluation metrics taken from rat behavioral studies to further evaluate robot results under a new set of experimental tasks and evaluation environments. In general, these metrics go beyond “qualitative” measures to include “quantitative” aspects that include learning latency and ability to navigate minimal distances to the goal. For this purpose we use a specially designed cyclical maze having external landmarks and providing different possible paths between to reach the goal.

In the rest of the paper we present related work (Section II), the spatial cognition (memory and learning) model developed by our group for task experimentation (Section III), animal experimentation results (Section IV), robot experimentation results (Section V), finishing with conclusions and a discussion (Section VII).

II. RELATED WORK

Taking inspiration in the rat hippocampus, several robotic navigation models have been proposed, as previously mentioned, such as Redish and Touretzky [14], Guazzelli et al. [7], Arleo and Gerstner [1], Arleo et al. [2], and Milford and Wyeth [10][17]. In this section we contrast some salient

features between these models and ours (Barrera and Weitzenfeld) to be overviewed in the following section. Key aspects characterizing these models are: (i) spatial memory - how place locations are modeled, (ii) spatial learning - how place locations are learned, and (iii) scope of tasks experimented with the models. In the model by Redish and Touretzky [14], representation of places integrates both vestibular and visual information codified by local view cells. In our model, we consider spatial representations as holistic topological-metric maps. The study by Guazelli et al. [7] proposed the TAM-WG model providing both taxon and locale navigation systems, and spatial representation combining kinesthetic and visual information. We consider our model as an extension to Guazelli by including (a) a map exploitation process to enable goal-directed navigation in a mobile robot, (b) the ability to learn goal locations from multiple departure positions within mazes, (c) the ability to learn maps with two or more decision points, and (d) a map adaptation process that permits on-line representations of changes in the physical configuration of the environment perceived by the robot. The main components of the neural architecture proposed by Arleo and Gerstner [1] and Arleo et al. [2] are similar to those found in our model: (a) the integration of allothetic (visual) information and idiothetic (path integration) signals at the level of the hippocampal representation, (b) the use of Hebbian learning to correlate these inputs, (c) the mapping of place cell population activity into spatial locations, and (c) the application of reinforcement learning to support goal-oriented navigation. We include in our model the use of affordances information instead of population vector coding to map the ensemble dynamics of place cells into spatial locations. We also include an explicit construction of a topological map of places and their metric relations, while implementing an Actor-Critic reinforcement architecture that predicts, adapts and memorizes reward expectations during exploration, suggesting a mutual influence between the hippocampus and the striatum. The focus of our approach also differs from the one followed by Milford et al. [10] primarily concerned with the effectiveness of the hippocampus models in mobile robot applications that explore larger environments with natural cues. Instead, our interest has been in providing mobile robots with spatial cognition capabilities similar to those found in rodents in order to produce comparable behavioral results and eventually provide experimental neuroscience with valuable feedback. Nevertheless, our model coincides with Milford et al. in some aspects related to mapping and map adaptation. In contrast to our separation of place and orientation information, Milford incorporates the use of experiences in the topological map storing the activity within pose cells that codify physical localization and orientation, together with local view cells that encode visual information.

III. A MODEL OF SPATIAL COGNITION

The spatial memory and learning model developed as part

of this work considers the following aspects: (i) interaction of different brain structures in rats, in particular hippocampus and striatum (work by Ramirez-Amaya et al. in detecting regions using immediate early gene Arc [13], to demonstrate skills associated with global and relative positioning in space, (ii) integration of traveled path, (iii) use of kinesthetic and visual cues during orientation, (iv) generation of topological-metric spatial representation of the environment, (iv) adaptation using Hebbian learning [8], (v) representation of animal internal motivational states based on hunger and thirst drives, and (vi) management of rewards implemented by reinforcement learning using an Actor-Critic architecture [5].

The model developed consists of a number of modules: motivation, learning, kinesthetic processing, landmark processing, affordances processing, place representation, and action selection modules as shown in Fig 1.

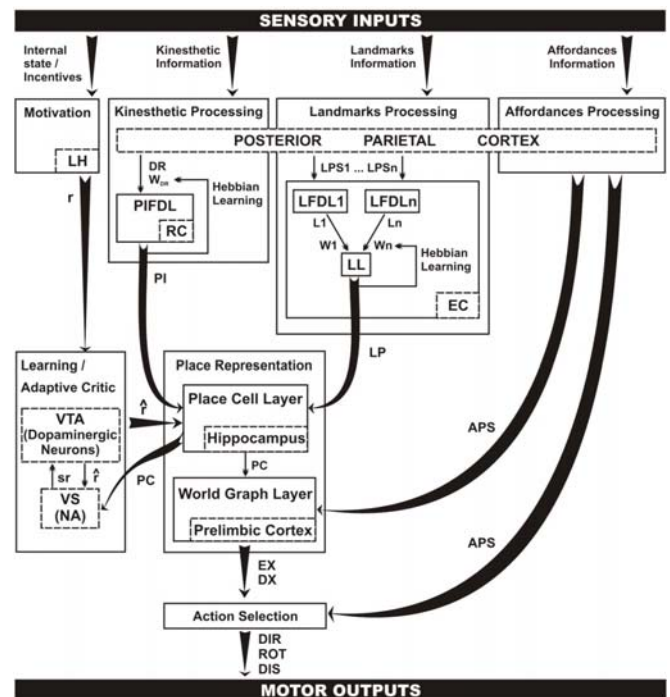


Fig. 1. The modules of the spatial cognition model and their interaction. Glossary: LH – Lateral Hypothalamus; RC – Retrosplenial Cortex; EC – Entorhinal Cortex; VTA – Ventral Tegmental Area; VS – Ventral Striatum; NA – Nucleus Accumbens; PIFDL – Path Integration Feature Detector Layer; LFDL – Landmark Feature Detector Layer; LL – Landmarks Layer. Inputs/Outputs: r = primary reinforcement; sr = secondary reinforcement; \hat{r} = effective reinforcement; DR= dynamic remapping perceptual schema; W_{DR} = connection weights between DR and PIFDL; LPS= landmark perceptual schema; APS= affordances perceptual schema; PI= kinesthetic information pattern; L= information pattern for one landmark; LP= landmarks information pattern; W_1, W_n = connection weights between L_1, L_n and LL; PC= place information pattern; EX= expectations of maximum reward and their corresponding directions (DX); DIR= next animat direction; ROT= animat rotation; DIS= next animat moving displacement.

These modules capture some of the properties of the rat's brain structures involved in spatial memory and learning. A detailed functional and mathematical depiction of each module is presented by Barrera and Weitzenfeld [4].

Motivation in the rat is related to its need to eat, i.e. its

hunger drive. The motivation module corresponds to the rat's lateral hypothalamus that controls food seeking and food intake computing a hunger value and immediate reward (r) the *animat* (animal or robot) gets by the presence of navigational goals.

Learning module represents dopaminergic neurons in the ventral tegmental area and ventral striatum processing reward information by use of an Actor-Critic architecture [5] where an Adaptive Critic (AC) estimates future reward values at any particular location (PC) in the environment. This module generates an effective reinforcement signal (\hat{r}) that supports the learning process. Actor-Critic architectures have been extensively used in modeling goal oriented behaviors in rats (e.g. Foster, Morris and Dayan [6]).

Kinesthetic processing module obtains information from two internal rat body systems: (i) vestibular organs in the semicircular canals of the inner ear, and (ii) feedback from muscles controlling movement. This module represents the updated animat position in relation to its point of departure (anchor) by integrating past rotations and translations through a dynamic remapping perceptual schema (DR), generating as output specific kinesthetic information patterns (PI).

Landmark processing module computes landmark-related spatial positioning, i.e. distance and relative orientation of each landmark to the animat. Spatial information about each landmark is encoded in a perceptual schema that produces and later responds to specific landmark patterns (LP) derived from the integration of all landmarks in the environment.

Affordances represent possible turns the animat can execute at any given time coded by the affordances perceptual schema (APS). McNaughton et al. [9] have suggested that preceding the rat's motion, nearly half of the cells in the Posterior Parietal Cortex (PPC) exhibit movement-related activity discriminating among basic modes of locomotion: left turns, right turns, and forward motion. Affordances determination is based on a local coordinate system that is relative to the animat's head.




Fig. 2. 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 and their corresponding directions (DX).

Place representation module, as shown in Fig. 2, receives input from kinesthetic processing (PI), landmark processing (LP) and affordances (AF) modules in addition to its interaction with the learning module. The place representation module comprises a place cell layer (PCL) and

a world graph layer (WGL). Place cell layer (PCL) represents the rat's hippocampus. Overlapping place fields in the collection of neurons in PCL are associated with a physical area in the environment that is identified directionally by the ensemble place cell activity pattern (PC). Synaptic efficacy between layers is maintained by Hebbian learning producing groups of neurons in PCL that respond to specific place cell information patterns (PC) derived from kinesthetic and egocentric visual information sensed by the animat while being at certain location and orientation. Associations between overlapping place fields and physical areas are represented and stored by a world graph layer (WGL) in a topological map enabling navigation between locations in the environment. Nodes in WGL represent associations between kinesthetic and visual information patterns and place cell population activity, whereas transitions between nodes are associated with metric information derived from animat locomotion such as moving direction and displacement. Every node in the map (a location) connects to a maximum of eight Actor units (eight possible orientations). Every Actor connection is associated with a weight (representing the expectation of reward when orienting to a particular direction from the current location), and an eligibility trace (marking the connection eligible to be reinforced later in time). In this way, Actor units compete to select the next moving direction (DX) from the current location or node that allows the animat to get the greatest reward (EX).

Action selection module computes motion direction (DIR), rotation (ROT) and displacement (DIS) by considering four signals corresponding to available affordances, a random rotation, rotations that have not been explored from the current location (curiosity), and the global expectation of maximum reward.




Fig. 3. Top view of the cyclic maze employed in the experiment with rats (left) and robots (right). Both mazes include landmarks L1, L2, L3, L4, locations TD, D1, D2, D3, D4, and goal (HE in the robot maze).

IV. ANIMAL EXPERIMENTATION RESULTS

We carried out animal experimentation in the Neurobiology Institute at UNAM in Mexico. We used a group of three male rats to carry out the behavioral experiments using a 95 x 105 cm² maze having six internal corridors as shown in Fig. 3 (left). The maze was placed in a room illuminated using a tenuous light, and surrounded by a white curtain with four different colored geometrical figures representing allocentric (external) 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 to provide sweet water to rats when reaching this location at the end of any training trial. The solution used as reward consisted of 5 gr of sucrose dissolved in 50 ml of water, whereas the water dispenser was filled with 400 μ l of the solution before beginning each trial. During the behavioral experiment, rats pass through three procedures: (A) habituation, (B) training, and (C) testing. These are described below in more detail.

A. Habituation Procedure

The task begins with a habituation session (pre-training), where the animal is placed in the maze and allowed to explore it freely from location TD oriented “north” (looking towards L1) until reaching the water dispenser. Water is not provided during the habituation session.

B. Training Procedure

During training, the water dispenser provides sweet water as reward to rats. At the beginning of each training trial the rat is placed at location TD oriented “north”. Each trial consists on freely exploring the maze until reaching the goal location and drinking the sweet water. The rat is then removed from the maze and placed in a cage for 60 sec after which a new trial is performed. Each rat daily training session included 9 trials, with each trial ending with the rat reaching the target for a total of 17 sessions. Latencies and routes followed by the animals in locating the water were registered by using a “SMART” system from San Diego Instruments. Fig. 4 shows routes followed by one of the rats during three training sessions, with sessions selected from beginning (Session 1), middle (Session 7), and end (Session 17) of the procedure. It is notorious how the learning latency of routes decreases as training evolves until the animal learns the shortest path to the goal.




Fig. 4. Routes followed by one rat in locating the water (Goal) during three training sessions (1, 7 and 17) from a fixed departure location (TD). Lines inside each maze illustrate 9 different routes. The average latency of arrival to the target is shown below each maze.

Fig. 5 shows the average length of path traveled by the rats during each training session normalized to the shortest path length possible in the maze from the fixed starting point TD to the goal location (i.e., the number of times that rats traveled the minimal distance possible to the goal). It is possible to observe how this number gets closer to 1.0 as rats learn the shortest path to the goal.

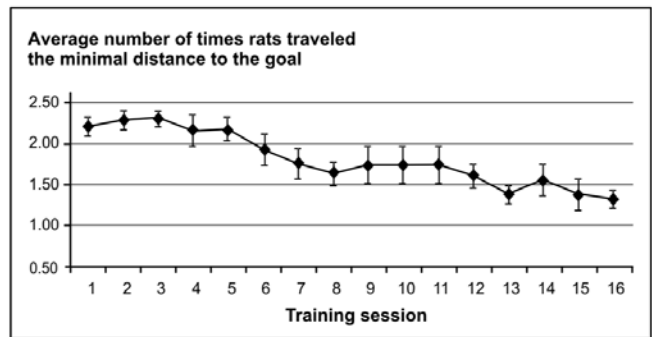


Fig. 5. Average performance during training in the task measured through the normalized route length to the goal location; i.e., the average number of times that rats traveled the minimal distance possible from the initial fixed location to the goal. Both graphs show the variability of the results.

A. Testing Procedure

In order to evaluate their ability to reach the target after having concluded the training procedure, rats were separately tested during 12 trials starting three times from four different locations in the maze, D1, D2, D3, and D4 (see Fig. 3 left). During each testing trial, the rat is placed at a corresponding initial location and orientation, and the trial consists on exploring the maze freely until reaching the goal location and drinking the sweet water. The rat is then removed from the maze. Fig 6 depicts latencies and routes followed by one of the animals in locating the water registered by using the SMART system, where it can be observed the rat not always follows an optimal performance following the shortest route.

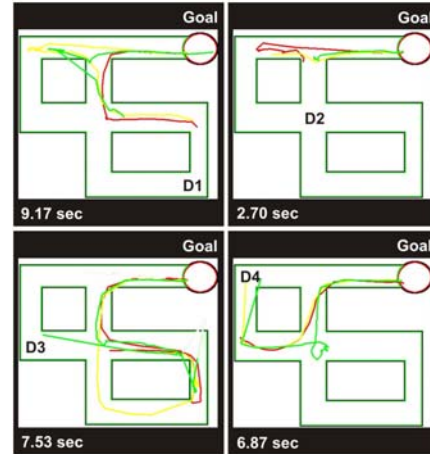


Fig. 6. Routes followed by one of the rats in locating the water (Goal) during tests departing from locations D1 to north, D2 to north, D3 to east, and D4 to south in the cyclic maze. Lines inside each maze illustrate three different routes. The average latency of arrival to the goal is shown below each maze.

V. ROBOT EXPERIMENTATION RESULTS

The rat cognitive model was designed and implemented using the Neural Simulation Language NSL system [16]. The computational model is used to control an AIBO robot (shown inside the maze in Fig. 3 (right)) in a 180 x 210 cm² maze having six internal corridors as shown in Fig. 3 (right). The maze is similar to that used with rats. The local AIBO camera takes as input three non-overlapping snapshots (left, forward, right) of the environment during each step. The

maze is surrounded by four colored cylinders representing allocentric cues or landmarks. During the experiment, similarly to rats, we used three robots passing through three phases: (A) habituation, (B) training, and (C) testing. These are described below in more detail.

A. Habituation Procedure

Similarly to rats, the experiment begins with an habituation procedure where the robot departs from location TD oriented “north” (looking towards L1). The robot explores the maze freely until reaching the goal location. After reaching the goal, the robot is removed from the maze. 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. As a result of the habituation procedure, the robot builds a cognitive map of the maze including nodes to represent the explored locations.

B. Training Procedure

After habituation, reinforcement is introduced into the model. The robot is placed at TD oriented north, and the training procedure begins from the cognitive map previously generated. The robot perceives the same visual and kinesthetic information sensed during habituation since the environmental configuration remains constant. The robot recognizes during training previously explored locations that were navigated in similar directions. When the robot visits unexplored locations or orientations, it adds new nodes and connections to the world graph map (see Fig 2). Additionally, when experimenting with different orientations at any given location, the robot merges possible redundant nodes by creating new ones and reconstructing links as necessary. In a typical training trial, the robot explores the maze from TD until finding the goal location. Exploration is interrupted if the robot spends more than 600 sec without reaching the goal. Depending on how exhaustively the robot explores the maze during habituation, its motor actions during training are more determined by the random factor and less determined by the curiosity factor since it recognizes all visited maze locations. Initially, while the robot’s reward expectations are not big enough to exceed the randomness, it will follow indirect routes to the goal (i.e., routes that involve visiting places more than once), or direct but not optimal routes. More direct routes are learned in time.

Training phase consisted of one session of 20 trials per robot. As with rats, latencies and routes followed by robots to reach the goal were recorded. Fig. 7 illustrates routes registered by one of the robots during three training trials selected from beginning (Trial 1), middle (Trial 8), and end (Trial 20) of the procedure. It can be observed how the learning latency decreases until reaching the goal location by following the optimal path.



Fig. 7. Routes followed by one of the robots in locating the goal during three training trials from a fixed departure location (TD) being oriented north in the cyclic maze. Illustrated trials were taken from beginning (Trial 1), middle (Trial 8), and end (Trial 20) of the learning process. The line inside each maze represents the route followed by the robot. The small arrow placed at location TD indicates the robot orientation. The average latency of arrival to the target in those three training trials is shown below each maze.

Fig. 8 shows the average number of times robots traveled the minimal distance possible from the initial fixed location TD to the goal during each training trial. As in the case of rats, it is possible to observe how this number gets closer to 1 as robots learn the shortest path to the goal.




Fig. 8. Average performance during training in the task measured through the normalized route length to the goal location; i.e., the average number of times that robots traveled the minimal distance possible from the initial fixed location to the goal.

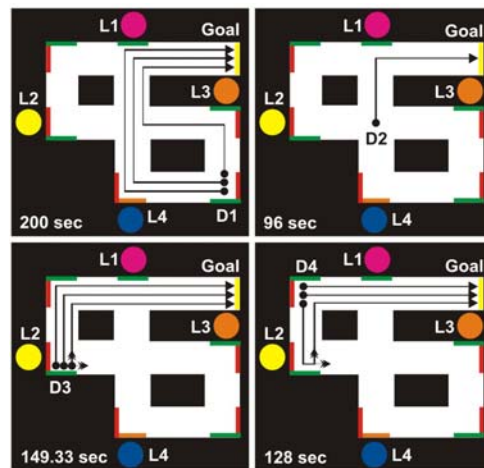


Fig. 9. Routes followed by one of the three robots while locating the goal during tests departing from locations D1 to north, D2 to north, D3 to east, and D4 to south, within the cyclic maze. Lines inside each maze illustrate routes recorded from three trials. When the robot followed the same route in those three trails, the corresponding maze shows only one line. The average latency of arrival to the goal location in three trials per departure point is shown below each maze.

C. Testing Procedure

In order to evaluate robots’ ability to reach the goal, we tested robot trajectories during 12 trials with robots liberated three times from the same initial locations and orientations as

used to test rats. During tests, latencies and routes followed by robots to reach the target were recorded, as shown by Fig. 9 for one of the robots. Each robot exploits the cognitive map built during habituation and adapted during training to reach the target successfully. During some of the testing trials, robots still modified their spatial maps by adding new nodes to represent unexplored locations and/or merging existent nodes.

Fig. 10 illustrates a graphical comparison between the average performance of rats and robots during tests departing from locations D1, D2, D3 and D4 in the maze, showing the number of times the rats/robots navigated the minimal distance possible to the goal location (i.e., the normalized length of the traveled path).




Fig. 10. Average performance of rats and robots during tests in terms of the number of times the subject navigated the minimal distance possible to the goal location.

VI. DISCUSSION AND CONCLUSIONS

The aim of the work presented in this paper is to introduce a set of metrics to evaluate task performance in spatial memory and learning that is common to both rat and robot experimentation. By having such metrics, not only roboticists will be able to better compare results against those obtained from animal experimentation, but also offer robotic architectures to researchers in neurobiology and neuroethology as an alternative platform to study, analyze and predict animal behavior. As part of this work we developed a computational model and robotic architecture of spatial memory and learning in rats that keeps a close relationship with hippocampus and striatum regions in the rat brain. By use of comparative experimental studies of spatial memory and learning in rats and robots, we demonstrated the behavioral performance match between robots and rats during the execution of a cyclical maze task by following a habituation, training and testing procedure where learning latency and normalized length of paths to reach the goal were contrasted. During these experiments, both rats and robots reached successfully the learning criterion during training, and were able during tests to find the target regardless of the departure location. Rats did not always show an optimal performance following the shortest route, and this behavior was also exhibited by robots during some tests. In both systems, efficient performance depends on how exhaustively they explore the environment during training visiting all

possible direct routes from the fixed departure location to the goal. Behavioral differences between rats and robots include: (i) robots learned faster than rats; (ii) robots reported less non-optimal trials during tests; and (iii) robots traveled less number of times the minimal distance to the goal during any testing trial than rats did. Exploring what contributes to this enhanced performance in robots may yield insights as to how learning can be improved in rats. Additional studies with rats will also enhance the robot models. In the future, we plan to extend this work with more complex tasks to evaluate that include adaptation to internal changes in the maze such as closing of existing corridors or opening of new ones.

REFERENCES

- [1] Arleo, A., Gerstner, W. (2000). "Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity", *Biological Cybernetics* 83: 287-299.
- [2] Arleo, A., Smeraldi, F., and Gerstner, W. (2004). "Cognitive navigation based on nonuniform Gabor space sampling, unsupervised growing networks, and reinforcement learning". *IEEE Transactions on Neural Networks* 15 (3): 639-652.
- [3] Barrera, A., and Weitzenfeld, A. (2007). "Rat-inspired Model of Robot Target Learning and Place Recognition", *Proceedings of the 15th Mediterranean Conference on Control and Automation – MED*, Athens, Greece.
- [4] Barrera, A., and Weitzenfeld, A. (2008). "Biologically-inspired Robot Spatial Cognition based on Rat Neurophysiological Studies", *Autonomous Robots*, Springer, Vol 25, No. 1-2, pp. 147-169.
- [5] Barto, A. G. (1995). "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.
- [6] Foster, D. J., Morris, R. G. M., Dayan, P. (2000). "A Model of Hippocampally Dependent Navigation, Using the Temporal Difference Learning Rule", *Hippocampus* 10: 1-16.
- [7] Guazzelli, A., Corbacho, F. J., Bota, M. and Arbib, M. A. (1998). "Affordances, motivation, and the world graph theory". *Adaptive Behavior* 6 (3-4): 435-471.
- [8] Hebb, D. O. (1949). *The organization of behavior: a neuropsychological theory*. Wiley-Interscience, New York.
- [9] McNaughton, B., Mizumori, S., Barnes, C., Leonard, B., Marquis, M., and Green, E. (1994). "Cortical representation of motion during unrestrained spatial navigation in the rat". *Cerebral Cortex* 4: 27-39.
- [10] Milford, M., Wyeth, G. (2007). "Spatial mapping and map exploitation: a bio-inspired engineering perspective". In Winter, S., Duckham, M., Kulik, L. and Kuipers, B. (Eds.), *Spatial Information Theory* (pp. 203-221), Heidelberg, Germany, Springer-Verlag.
- [11] Morris, R. G. M. (1981). "Spatial localization does not require the presence of local cues". *Learning and Motivation* 12: 239 – 260.
- [12] O'Keefe, J. (1983). "Spatial memory within and without the hippocampal system," in *Neurobiology of the Hippocampus*, edited by W. Seifert, Academic Press, New York, pp. 375-403.
- [13] Ramirez-Amaya, V., Vazdarjanova, A., Mikhael, D., Rosi, S., Worley, P. F., and Barnes, C. A. (2005). "Spatial exploration-induced arc mRNA and protein expression: Evidence for selective, network-specific reactivation". *Journal of Neuroscience* 25: 1761-1768.
- [14] Redish, A. and Touretzky, D. (1997). "Cognitive maps beyond the hippocampus". *Hippocampus* 7 (1): 15-35.
- [15] Webb, B. (2001). "Can robots make good models of biological behaviour?" *Behavioral and Brain Sciences* 24: 1033-1050.
- [16] Weitzenfeld, A., Arbib, M., Alexander, A. (2002). *The Neural Simulation Language*. MIT Press.
- [17] Wyeth, G. Milford, M., (2009). "Spatial cognition for robots", *IEEE Robotics and Automation Magazine*, Vol 16 No 3, pp 24-32, September.