

# Place Sequence Learning for Navigation

Olivier Trullier<sup>1,2</sup> and Jean-Arcady Meyer<sup>1</sup>

<sup>1</sup> AnimatLab, Département de Biologie, École Normale Supérieure,  
46, rue d'Ulm, 75230 Paris Cedex 05, France

<sup>2</sup> CNRS-Collège de France, LPPA, 11, place Marcelin Berthelot, 75005 Paris, France

**Abstract.** A model of the hippocampus as a “cognitive map”, inspired by the models of Burgess et al. (1994) and Jensen et al. (1996), is proposed. Simulations show that the resulting navigation behavior is as efficient as the behavior exhibited by previous models. However, the architecture of the proposed model and the mechanisms governing the temporal characteristics of the neurons in the model are more realistic. In particular, the proposed model assigns distinct and specific roles to the entorhinal cortex, the dentate gyrus and the hippocampal CA3 region. In essence, the dentate gyrus could act as a short-term memory that maintains the representation of a sequence of recently visited places. It could then force the corresponding CA3 place cells to fire and to learn the spatial relationships between places through a Hebbian rule. This “topological representation” could then serve as a basis for predicting places ahead of the animal and drive “goal cells”, i.e. cells that represent the direction to the goal, as proposed by Burgess et al.

## 1 Introduction

The discovery of *place cells* in areas CA3 and CA1 of the rat hippocampus [7] – cells that discharge selectively in restricted regions of the environment (place fields) – led to the idea that the hippocampus functions as a *cognitive map* of space [8]. However, cells with spatial selectivity have subsequently been found in the entorhinal cortex [10] and in the dentate gyrus [4]. Thus the different regions of the hippocampal formation successively process position information. Cells that fire as a function of the animal’s orientation in space, *head-direction cells*, have also been evidenced in various other regions of the rat brain (e.g. postsubiculum [11]). Representations of position and orientation, necessary for any navigation system, are thus present in the rat brain.

Among the existing computational models of animal navigation, i.e. *animats* [6,12], the model by Burgess et al. (hereafter called the BROm) [1] is one of the most effective in terms of the resulting navigational capabilities while maintaining some resemblance to the rat hippocampus. One problem however is that the BROm relies on an additional property of the place cell firing – phase coding – for which it provides no biologically plausible mechanism. Phase coding, discovered by O’Keefe and Recce [9], refers to the specific temporal relationship

---

In W. Gerstner, A. Germond, M. Hasler, and J.-D. Nicoud, editors, Proceedings of the 7th International Conference on Artificial Neural Networks (ICANN’97), pp. 757-762, Springer-Verlag.

between the discharge of place cells and the overall theta ( $\theta$ ) rhythm, a sinusoidal EEG oscillation observed in the rat hippocampus when the animal is engaged in locomotor behaviors. As the animal traverses a place field, the firing of the corresponding place cell precesses from late to progressively earlier phases of the  $\theta$  cycle. Burgess et al. call upon an ad hoc mechanism to define the phase at which the place cells should fire by looking where the animal is with respect to estimations of where place fields would develop during exploration.

Jensen et al. [2] recently presented a neural network model (hereafter called the JILm) that is capable of learning sequences of items. The JILm exploits the fact that 40Hz gamma oscillations superimposed on the  $\theta$  rhythm (6Hz) can divide each  $\theta$  cycle into “chunks” that each represent a memorized item. Thus, each  $\theta$  cycle contains a sequence of seven or less items and synaptic connections between cells firing during successive gamma cycles are enhanced. Jensen and Lisman [3] showed that the JILm could learn sequences of visited places and could, as a consequence, account for the phase precession phenomenon. However, the JILm can only learn sequences of seven or less items and the introduction of an eighth item dramatically disrupts the learned sequence. Furthermore, learned place sequences are one-dimensional.


This paper describes a model that remedies the weaknesses of the approaches presented by Burgess et al. and Jensen et al. by adapting and combining them. Temporal sequences of visited places are transformed into a topological representation of the environment, a graph where nodes correspond to places and edges indicate that places are neighbors [12]. By labeling each edge according to the direction of movement, the graph becomes two-dimensional. Phase coding results from the predictions of places lying ahead of the animal. This is then incorporated into the BROm. We then demonstrate that an animal that uses such a neural network can efficiently navigate in a simple environment.

## 2 The Model

The global architecture of our model (Fig. 1) closely emulates the architecture of the rat hippocampus. Each of the 4 layers has a specific and distinct function.

**Entorhinal Cortex:** The entorhinal cortex (EC) is the major entry of spatial information into the hippocampal formation. We postulate here that the EC neurons are selective for places [10] because they respond to local views. They then transmit this orientation-dependent position information to the short-term memory and to the topological graph learning layers.

**Dentate Gyrus:** This is the intermediate short-term memory store, inspired by the JILm. An intrinsic mechanism (see [2] for details) enables dentate granular (DG) cells to fire repeatedly at each  $\theta$  cycle, once they have been triggered to fire. The evidence for such a mechanism comes from *in vitro* slice studies on pyramidal cells. We postulate that similar mechanisms exist in DG cells. We suggest that the EC cells send convergent projections onto DG cells to give orientation-independent location selectivity. These cells would fire for the current place, but also for previously visited places, and they would repetitively fire in the order of the visited sequence. However, as the animal moves, the sequence of all the previously visited places grows longer and longer and cannot be stored in



**Fig. 1.** Architecture of the proposed model. Large triangles stand for pyramidal cells, unfilled circles for granule cells, and small shaded circles for inhibitory interneurons. An activity-dependent afterdepolarizing current (ADP) is intrinsic to each granule cell. 8 goal cells (4 shown) represent North, North-East, East and so on.

its entirety. Thus, we propose here that when the animat moves about, the DG stores a *sliding* sequence that has a maximum length, where representations of new places continuously enter the queue and representations of already visited places leave it. To this end, we propose a modification to the intrinsic mechanism used by Jensen et al. Originally, the cyclically repetitive firing of the DG cells representing a visited place was solely driven by the intrinsic afterdepolarizing current (ADP) – a current that builds up after a cell fires and eventually triggers another spike. In our model, it is driven by both the ADP and the continuous synaptic input coming from the EC that triggered the first spike. Neither ADP nor the synaptic input alone can maintain the repetitive firing. In other words, a representation is maintained in short-term memory as long as the corresponding “sensory” input is present. A consequence of this mechanism is that when the animat turns and the local views change (as does the whole input representation), an entirely new sequence is initiated. Thus, an active sequence in short-term memory corresponds to a straight path in a specific direction. The links between all the active place cells can thus be labeled with the same head-direction, i.e. the current one. This is done in CA3.

**CA3:** With its place cells and its dense recurrent connectivity, the CA3 region is well suited to function as a topological representation of the environment. During exploration, the DG cells that fire would force the CA3 place cells to fire in the same sequence through the strong mossy fiber connections. In contrast to the BROm where the CA3 recurrent connections were not taken into account, we postulate here that CA3 acts as a *heteroassociative* network. The synaptic weights are modified through Hebbian learning. Taking into account Mel’s suggestion [5] that pyramidal cells could behave like formal Sigma-Pi neurons, in our model, inputs from head-direction cells modulate the propagation of activity between place cells by gating each synapse. Consequently, each synapse potentially represents the fact that the place field of the post-synaptic cell is situated next to the place field of the pre-synaptic cell, in the direction defined by the gating head-direction cell. Individual sequences that correspond to linear segments in

the environment are merged into a two-dimensional spatial representation where any two neighboring places are connected by an orientation-modulated synapse.

**Goal Cells:** Similarly to the BROM, our model employs goal cells, downstream of the hippocampus. These cells code for where the animat is with respect to the goal. For instance, the “East” goal cell fires when the animat is to the East of the goal. To learn goal locations, the animat, upon receiving a reinforcement, looks about in all directions. The conjunctive activation of the place cell representing the goal location and of specific head-direction cells propagates through the orientation-modulated synapses, i.e. in the corresponding directions. It thus triggers *all* the place cells that represent locations in the respective region (e.g. East) with respect to the goal. The connections from these activated place cells to the respective goal cell are enhanced through a Hebbian learning rule.

### 3 Simulations

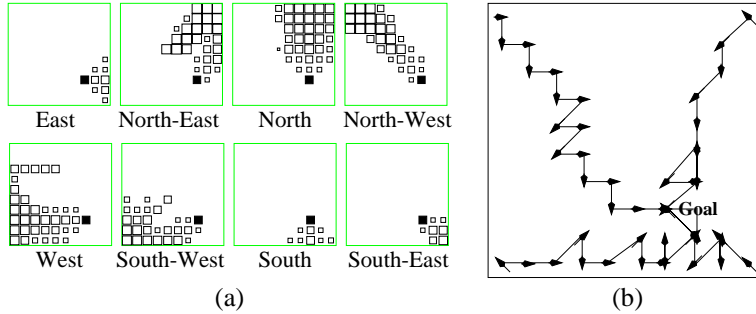
Each neuron is described as a single compartment model where the membrane potential obeys the following differential equation:

$$C \frac{dV}{dt}(t) = \sum_k g_k(t) \cdot (V - V_k)$$

where  $k$  stands for the respective ionic currents,  $g_k$  and  $V_k$  for the corresponding conductances and reversal potentials. The model includes the following currents: the leakage current, the fast afterhyperpolarization, the  $\theta$  modulation, the feedback inhibition, the synaptic input (AMPA or NMDA), and the afterdepolarization (the latter current appears only for the DG cells). Conductances are described as alpha functions or by differences of exponentials.

Simulations are separated into 3 phases: (i) the animat explores the environment and the network learns connections between predefined place cells; (ii) the animat stands at the goal, looks around and the network predicts which place cells will fire in various directions and associates them with goal cells; (iii) the network uses these learned connections to predict where the goal is.

Our current simulations consist of grid-like environments where the animat moves from place to place at every  $\theta$  cycle, in one of eight directions (labeled “North”, “North-East”, and so on). To reduce computation times the discharge profile, and thus the synaptic inputs to the DG cells, are analytical approximations based on observations from pilot simulations. Only the DG cells are simulated during exploration but, since they force the discharge in CA3 through the strong mossy fiber connections, the appropriate CA3 cells are assumed to fire and all the synaptic weights would undergo Hebbian learning. For the learning of the goal location and the search phase, only CA3 place cells and goal cells are simulated, with appropriate synaptic inputs. There is one DG cell and one CA3 place cell for each position in the environment. The mechanism by which EC cells respond to local views is not modeled explicitly here. EC cells start firing at one specific position and when the animat faces a specific direction and keeps firing if the animat moves in that direction, over a distance corresponding to  $5\theta$  cycles. The animat has a certain probability (0.15) of turning  $\pm 45^\circ$  at each step and randomly chooses a new direction when it encounters a wall.



**Fig. 2.** (a) Weight distributions of synapses from place cells to goal cells. Each labeled large square represents the 10 by 10 arena shown in (b). The size of each small square indicates the synaptic weight from the corresponding place cell to the labeled goal cell (e.g. the filled square represents the synaptic weight from place cell (9,8) to the East goal cell). (b) Trajectories to the memorized goal from different starting positions. The animat uses the goal representation shown in (a). The fact that trajectories tend to oscillate near the goal is probably due to the low level of discretization of the arena.

## 4 Performance

Once the topological representation is acquired and the goal location learned (Fig. 2a), the animat can return to the goal from any starting position (Fig. 2b). At the beginning of each  $\theta$  cycle, the place cell corresponding to the current location fires and triggers the firing of place cells with fields situated in the direction the animat is facing, by virtue of enhanced connections learned during exploration. These cells in turn drive the goal cells. The animat then moves in the direction opposite to that coded by the goal cell population. We have limited most of our simulations to a 10 by 10 square environment for improved computation times but simulations on bigger environments (not shown) revealed no qualitative differences in performance.

## 5 Conclusion and Perspectives

We have modified and combined two distinct models of the hippocampus, one expressed in terms of a cognitive map and the other as a sequence memory storage, and shown that the resulting animat can reach the goal in almost direct paths from different starting positions, within a very simple environment. The generated trajectories are not as direct as the ones generated by the BROM, but contrary to the BROM, we have encountered no limitation on the size of the environment, that is, our goal cells were active even at long distances from the goal. The overall performance of the present navigation system doesn't surpass the BROM. However, we would argue that our model is more biologically plausible and improves on both the BROM and the JILm: (1) its architecture is more closely related to the actual rat hippocampus and we tried to attribute *specific* and *distinct* functions to the respective regions (entorhinal cortex as relay between input and memory, dentate gyrus as sequence storage and CA3 as topological representation); (2) our model involves fewer ad hoc mechanisms. In

particular, it palliates the major weakness of the BROM, namely the computation of the firing phase, which is now intrinsic to the model; (3) an inherently one-dimensional sequence learning is transformed into a two-dimensional topological representation through the use of sigma-pi connections.

These are preliminary results. Presently, the environment is simple and the movements of the rat are restricted. We are currently relaxing these constraints by enabling the animat to move in a smooth and continuous manner.

*Acknowledgments* We would like to thank members of the “groupe modèle” at LPPA (S.Wiener, J.Droulez, N.Brunel & A.Berthoz) for helpful discussions. Work supported by the CNRS GIS “Sciences de la Cognition”. O.T. supported by a MESR fellowship.

## References

1. N. Burgess, M. Recce, and J. O’Keefe. A model of hippocampal function. *Neural Networks*, 7:1065–1081, 1994.
2. O. Jensen, M. A. P. Idiart, and J. E. Lisman. Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: Role of fast NMDA channels. *Learn. Mem.*, 3:243–256, 1996.
3. O. Jensen and J. E. Lisman. Hippocampal CA3 region predicts memory sequences: Accounting for the phase precession of place cells. *Learn. Mem.*, 3:279–287, 1996.
4. M. W. Jung and B. L. McNaughton. Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, 3(2):165–182, 1993.
5. B. W. Mel. Synaptic integration in an excitable dendritic tree. *J. Neurophys.*, 70:1086–1101, 1993.
6. J.-A. Meyer and A. Guillot. Simulation of adaptive behavior in animats: Review and prospect. In J.-A. Meyer and S. W. Wilson (eds.), *SAB90*. MIT Press, 1991.
7. J. O’Keefe and J. Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Res.*, 34:171–175, 1971.
8. J. O’Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford, 1978.
9. J. O’Keefe and M. L. Recce. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3:317–330, 1993.
10. G. J. Quirk, R. U. Muller, J. L. Kubie, and J. B. Ranck. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *J. Neurosci.*, 12:1945–1963, 1992.
11. J. S. Taube, R. U. Muller, and J. B. Ranck. Head-direction cells recorded from the postsubiculum in freely moving rats: Description and quantitative analysis. *J. Neurosci.*, 10:420–435, 1990.
12. O. Trullier, S. I. Wiener, A. Berthoz, and J.-A. Meyer. Biologically-based artificial navigation systems: Review and prospects. *Prog. Neurobiol.*, 51:483–544, 1997.