A MODEL OF SPATIAL MAP FORMATION IN THE HIPPOCAMPUS OF THE RAT

Kenneth I. Blum* and L.F. Abbott

Center for Complex Systems
Brandeis University
Waltham, MA 02254

Published in *Neural Comp.* **8**:85-93 (1996).

Abstract

Using experimental facts about long-term potentiation (LTP) and hippocampal place cells, we model how a spatial map of the environment can be created in the rat hippocampus. Sequential firing of place cells during exploration induces, in the model, a pattern of LTP between place cells that shifts the location coded by their ensemble activity away from the actual location of the animal. These shifts provide a navigational map that, in a simulation of the Morris maze, can guide the animal toward its goal. The model demonstrates how behaviorally generated modifications of synaptic strengths can be read out to affect subsequent behavior. Our results also suggest a way that navigational maps can be constructed from experimental recordings of hippocampal place cells.

*Current address: Dept. of Brain and Cognitive Sciences, MIT E25-236, 45 Carlton St., Cambridge, MA 02139.

Blockade of long term potentiation (LTP) and hippocampal lesions drastically impair the ability of rodents to navigate to a goal using distal cues (Morris *et al.* 1986; Morris *et al.* 1982; Jarrard 1993; O'Keefe and Nadel 1978). This has led to suggestions that the hippocampus plays a role in navigation (McNaughton *et al.* 1991; Worden 1992; Hetherington and Shapiro 1993; Burgess *et al.* 1994; Wan *et al.* 1994) by providing a cognitive map of the spatial environment (O'Keefe and Nadel, 1978). It has further been suggested that this cognitive map is stored by potentiated synaptic weights representing both spatial (Muller *et al.* 1991; Traub *et al.* 1992) and temporal (Levy 1989) correlations. In recent experiments the activity of an ensemble of place cells was decoded to reveal the location represented by their collective firing (Wilson and McNaughton, 1993). The ability to decode place cell ensemble output provides an opportunity to examine ideas about spatial maps quantitatively and to explore specific mechanism by which they can be created and stored. We examine the effect of LTP on the position encoded by place cell ensemble firing and show how it can produce a cognitive map useful for navigation.

What information is stored in the hippocampal cognitive map, how is it stored and how is it read out to guide navigation? We propose that this information resides in shifts of the position coded by hippocampal place cell activity that arise from synaptic potentiation. The prevalence of place cells in the hippocampus might lead to the assumption that this region primarily serves to represent the spatial location of the animal. However, knowledge of spatial location is not the function normally associated with a map. Rather, a map serves to suggest directions for future movement based on knowledge of present location. We suggest that the cognitive map in the hippocampus plays this role. We assume that the spatial location of the animal is determined from sensory input outside the hippocampus, that this information is available to the animal and is also transferred to the hippocampus. In our model, the role of the hippocampus is not merely to report this position, but to suggest directions of future motion on the basis of previous experience. We will first show how the location represent by place cell activity shifts in a direction that reflects the past experience of the animal in the environment and then discuss how this shifted location can be compared with the present location of the animal to provide a navigational cue.

Our model for the storage and read-out of a navigational map makes use of three key ingredients, all supported by experimental data. First, NMDA-dependent LTP in hippocampal slices occurs only if presynaptic activity precedes postsynaptic activity by less than approximately 200ms (Levy and Steward 1983; Gustafsson *et al.* 1987). Presynaptic activity following postsynaptic firing produces either no LTP or long-term depression (Debanne *et al.* 1994). Second, place cells – neurons broadly tuned to location – exist in the hippocampus and make synaptic connections with each other both within the CA3 region and between CA3 and CA1 (O'Keefe and Dostrovsky 1971; Amaral 1987). Third, a spatial location can be determined by appropriately averaging or fitting the activity of an ensemble of hippocampal place cells as has been done in other systems (Georgopoulos *et al.* 1986; Salinas and Abbott 1994). This coded position is near, but not necessarily identical to, the true location of the animal (Muller and Kubie, 1989; Wilson and McNaughton 1993).

These three observations imply that when an animal travels through its environment causing different sets of place cells to fire, information about both temporal and spatial aspects of its motion will be reflected in changes of the strengths of synapses between place cells. Since this LTP affects subsequent place cell firing, it can shift the spatial location collectively coded by place cell activity. We compute how this coded location is shifted relative to the true position of the animal and find: 1) If LTP occurs while an animal is sitting at a specific point, the coded location is shifted toward this point. 2) If LTP occurs while an animal traverses a specific path, the coded location is shifted toward and forward along the path. 3) If many locations and paths contribute to LTP, the shifts reflect the entire history of spatial exploration and provide a map of the environment useful for navigation.

Shifts between the coded and actual positions arise from the following mechanism. During locomotion, cells with place fields overlapping a path being traveled are sequentially activated. A moving rat covers a few centimeters in the 200ms time window for LTP induction. Thus, synapses from presynaptic cells with place fields overlapping a path to postsynaptic cells with fields a few centimeters forward along the path will be potentiated. Subsequently, when the animal is on the learned path, activated cells will excite neurons with place fields ahead of them along the path through the potentiated synapses. This shifts the 'center of gravity' of the firing pattern, and thus the coded location, forward along the path. Similarly synapses from cells with place fields beside the path to those on it will be potentiated, and will then shift the coded position toward the path. For details of the calculation see the section on mathematical results and Abbott and Blum (1995).

These shifts suggest that an animal could navigate by heading from its present location toward the position coded by place cell activity. To illustrate both how a spatial map arises and how it can be used to guide movement, we applied these ideas to navigation in the Morris maze. We should stress that we are not modeling how a rat solves the problem of navigating to a hidden platform in this task. We are studying, instead, how the experiences of the animal while it learns the task are recorded in a cognitive map, and how this recorded information can then be used to help in performance of the task. Similarly, we do not model how place fields arise but instead, study how they are affected by LTP.

In the model, the animal was represented by a point that moved with a velocity composed of two equal parts: a component at a random angle with respect to the current velocity drawn from a uniform distribution of width one radian, and a component parallel to the LTP-induced shift of the coded position. The velocity was normalized to 20cm/s swimming speed. LTP occurred continually. In this form the model had mixed success. If the starting position was gradually moved away from the platform on successive runs an efficient path from any location to the platform was learned robustly. If, however, the starting location was chosen at random around the perimeter, an instability could result. A loop in the trajectory could be reinforced at a location away from the platform. There are several biologically reasonable ways to solve this problem. We chose to introduce a simple reward system. Synapses were only potentiated when the platform was reached and recently activated synapses were potentiated more strongly than synapses activated early in the trial. The amount of LTP was weighted by an exponential factor with a time constant of 4s. This scheme is not supposed to be a realistic description of how the animal solves the navigational task. However, it provides a simple way for the task to be solved in the model so that we can examine how the spatial map forms and what its role might be.

The results of a typical simulation are illustrated in Fig. 1. Early in a trial, paths meander until they get close to the target. Ultimately, the approach to the platform is quite direct and efficient. Fig. 1c shows a simulation of the transfer test in which the platform is removed following training. The path goes directly to the region where the platform had been and stays in that vicinity.

The navigational map that guides the motion in these simulations evolves as a function of the trajectories followed. The map at various stages of the simulation is shown by the arrows in Fig. 1. Initially, all the arrows have zero length since no LTP-induced shift has yet occurred. After a few trials, the arrows point toward the goal, but only over a limited range of the environment near the platform. As subsequent paths lead toward the platform, the range of this directed map extends outward. In Fig. 1c all the arrows point roughly toward the platform indicating that hippocampal place cell activity can provide the information needed to find the platform from any initial location.

Fig. 2 shows how the average time required to find the platform decreases as a function of trial number. There is a decrease in latency over the first 10 trials followed by consistently efficient platform finding behavior. These results resemble the performance of normal animals (Morris *et al.* 1982; Morris *et al.* 1986).

We propose that the difference between ensemble coded position and actual location acts as a cognitive map capable of guiding navigation through the environment. We speculate that this could occur in the following way. Information about the actual location of the animal enters the hippocampus from the entorhinal cortex passing to both CA3 and CA1 regions. The shifted representation that develops due to LTP between place cells is likely to arise in CA3 with its extensive recurrent collaterals. Since CA1 receives input from both CA3 and directly from entorhinal cortex, it could simultaneously represent both the true location and the shifted position. From such a distributed, dual representation it is possible to extract the difference, and this could guide the direction of locomotion (Andersen *et al.* 1985; Zipser and Andersen 1988; Salinas and Abbott 1994b). This idea is completely consistent with the presence of place cells in CA1 but would predict that their activity depends on both the actual location of the animal and on the location coded in CA3. Unfortunately, since these two locations cannot easily be varied independently, it will be difficult to establish whether or not a dual representation exists in CA1. However, if it does, then a downstream network can extract the vector difference needed to guide navigation (Salinas and Abbott, 1995). Of course, the comparison of the coded and actual positions could just as well take place anywhere along the pathway from CA3 output to motor activity (Muller and Kubie; 1989).

Asymmetric synaptic weights develop in our model because of the temporal asymmetry of LTP induction and because place fields are activated sequentially during locomotion. The phase dependence of place cell firing with respect to theta oscillations reported by O'Keefe and Recce (1993) may also play a role. Within each theta cycle, activated place cells tend to fire in the order that the animal encountered their place fields. If we assume that the window for LTP induction does not extend from one theta cycle to the next, this will contribute to the asymmetric potentiation that is central to our model and to the resulting shifts of the coded location in the forward direction along a path. Additional, more complex contributions will arise if the LTP window extends across theta cycles (O'Keefe and Recce, 1993).

Our results show that information about trajectory history can be stored in synaptic strengths and that it can be communicated to subsequent networks through changes in the overall pattern of neuronal activity. Our model is not meant to be a complete description of the mechanisms by which a rat navigates, but one element of a navigational system. This element is particularly interesting because we can relate it to specific neuronal and molecular mechanisms: place cell ensemble coding and properties of the NMDA receptor (Lester *et al.* 1990; Hestrin *et al.* 1990; Jahr and Stevens 1990) that give rise to temporally asymmetric LTP (Levy and Steward 1983; Gustafsson *et al.* 1987). The changes in place cell activity that we have computed arise inevitably if NMDA-mediated LTP occurs between place cells during locomotion. The model predicts that the position of an animal that is moving toward a goal should lag the location decoded from place cell activity. Although we are unable to predict the magnitude of this effect precisely, we expect that it is smaller than the reported 5cm tracking uncertainty (Wilson and McNaughton 1993). Muller and Kubie (1989) report that place fields are best described as centered on future location and speculate about a navigational role for this effect.

Direct experimental tests of our model are feasible. Individual place fields should be altered by learning paths through the environment and the shifts in the location decoded from place cell ensembles may be observable. For example, place fields lying along a path that is frequently traversed in one direction should elongate and move backward along the path. Furthermore, it has recently been shown that place cell pairs correlated during behavior have enhanced correlation during subsequent slow-wave sleep (Wilson and McNaughton 1994). The asymmetric, short-latency, pair-correlation functions measured in the hippocampus during sleep may reflect synaptic weights because the inputs to the hippocampus are less correlated during slow wave sleep than during behavior. The correlation matrix and measured average firing rates can be used to generate navigational maps like those of Fig. 1 directly from experimental data using Eqn. 3 below. The existence of a bias in the shift arrows toward a goal location would provide dramatic evidence of a navigational map in the hippocampus consistent with our model.

Mathematical Results

Let $H(t')$ represent the rate of LTP induction for unit firing rates when presynaptic activity precedes postsynaptic activity by a time t'. If LTP occurs during motion along a path $\vec{X}(t)$ the strength of the synapse from place cell j to place cell i is enhanced by

$$\Delta W_{ij} = \int dt dt'\, H(t') f_i\left(\vec{X}(t+t')\right) f_j\left(\vec{X}(t)\right) \tag{1.1}$$

where f_i is the average firing rate of place cell i as a function of position. This equation allows the synaptic weights to grow without bound but in the simulations we constrained the weights so that the resulting shift is less than unit magnitude. After LTP the firing rate of place cell i is $r_i = f_i(\vec{x}) + \sum_j \Delta W_{ij} f_j(\vec{x})$ when the animal is at location $\vec{x}$. The coded position is given by

$$\vec{p} = \frac{\sum_i \vec{s}_i r_i(\vec{x})}{\sum_i r_i(\vec{x})} = \vec{x} + \Delta\vec{p} \tag{1.2}$$

where $\vec{s}_i$ is the center of the place field for cell i. From this we find

$$\Delta\vec{p} \approx \frac{\sum_{ij}(\vec{s}_i - \vec{x})\Delta W_{ij} f_j}{\sum_i f_i}\ . \tag{1.3}$$

To evaluate this expression we use Gaussians with width 2σ and height R_{max} for the f_i, replace sums over place cells with integrals over their place field centers, and use the approximation $f_i\left(\vec{X}(t+t')\right) \approx f_i\left(\vec{X}(t)\right) + t'\dot{\vec{X}}(t)\cdot\vec{\nabla} f_i\left(\vec{X}(t)\right)$. This gives the result

$$\Delta\vec{p}(\vec{x}) \approx h\pi R_{max}^2 \rho\sigma^2 \int dt\, \left[\vec{X}(t) - \vec{x} + \tau\dot{\vec{X}}(t)\right] \exp\left(-\frac{|\vec{X}(t)-\vec{x}|^2}{4\sigma^2}\right) \tag{1.4}$$

with h the integral of H, ρ the place field density, τ the average LTP window time (the first moment of H) and $\dot{\vec{X}}$ the velocity of the learned path. Other decoding methods give similar results (Salinas and Abbott 1994). For the figures, we integrated Eqn. (1.4) numerically with $h\pi R_{max}^2\rho\sigma^2 = 0.4$, $\sigma = 7$cm, and $\tau = 200$ms.

Acknowledgements

Supported by NSF-DMS9208206, the W.M. Keck Foundation and the McDonnell-Pew Centre for Cognitive Neuroscience at Oxford (LA) and NIH-NS07292 (KB). We thank Marco Idiart and Eve Marder for discussions.

References and Notes

Abbott, L. F. and Blum, K. I. (1994) Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex* **6**, 406-416.

Amaral, D. (1987) Memory: anatomical organization of candidate brain regions, in *Handbook of Physiology Sec. I The Nervous System Vol. V* (ed. F. Plum) pp. 211-294 (Oxford Univ. Press, NY).

Andersen, R. A., Essick, G. K., and Siegel, R. M. (1985) The encoding of spatial location by posterior parietal neurons. *Science* **230**, 456.

Burgess, N., Recce, M., and O'Keefe, J. (1994) A model of hippocampal function. *Neural Networks* **7**, 1065-1081.

Debanne, D., Gahwiler, B. H., and Thompson, S. M. (1994) Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat hippocampus *in vitro Proc. Natl. Acad. Sci. USA* **91**, 1148-1152.

Georgopoulos, A. P., Schwartz, A., and Kettner, R. E. (1986) Neuronal population coding of movement direction. *Science* **233**, 1416-1419.

Gustafsson, B., Wigstrom, H., Abraham, W. C., and Huang, Y. -Y. (1987) Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.* **7**, 774-780.

Hestrin, S., Sah, P., and Nicoll, R. A. (1990) Mechanisms generating the time course of dual component excitatory synaptic currents recorded in hippocampal slices. *Neuron* **5**, 247-253.

Hetherington, P. A. and Shapiro, M. L. (1993) A simple network model simulates hippocampal place fields: II. Computing goal - directed trajectories and memory fields. *Behav. Neurosci.* **107**, 434.

Jahr, C. E. and Stevens, C. F. (1990) A quantitative description of NMDA receptor channel kinetic behavior. *J. Neurosci.* **10**, 1830-1837.

Jarrard, L. E. (1993) On the role of the hippocampus in learning and memory in the rat. *Behav. and Neural Biol.* **60**, 9.

Lester, R. A. J., Clements, J. D., Westbrook, G. L., and Jahr, C. E. (1990) Channel kinetics determine the time course of NMDA receptor-mediated synaptic currents. *Nature* bf 346, 565-567.

Levy, W. B. (1989) A computational approach to hippocampal function, in *Computational Models Of Learning In Simple Neural Systems* (eds. R. D. Hawkins and G. H. Bower) pp. 243-305 (Academic Press, San Diego, CA).

Levy, W. B. and Steward, D. (1983) Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neurosci.* **8**, 791-797.

McNaughton, B. L., Chen, L. L., and Markus, E. J. (1991) 'Dead reckoning', landmark learning, and the sense of direction: a neurophysiological and computational hypothesis. *J. Cogn. Neurosci.* **3**, 190.

Morris, R. G. M., Schenk, F., Garrud, P., Rawlins, J. N. P., and O'Keefe, J. (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681.

Morris, R. G. M., Anderson, E., Lynch, G. S., and Baudry, M. (1986) Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature* **319**, 774-776.

Muller, R. U. and Kubie, J. L. (1989) The firing of hippocampal place cells predicts the future position of freely moving rats. *J. Neurosci.* **9**, 4101-4110.

Muller, R. U., Kubie, J. L., and Saypoff, R. (1991) The hippocampus as a cognitive graph. *Hippocampus* **1**, 243-246.

O'Keefe, J. and Dostrovsky, J. (1971) The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171-175.

O'Keefe, J. and Nadel, L. (1978) *The Hippocampus as a Cognitive Map* (Clarendon, London).

O'Keefe, J. and Recce, M. L. (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317-330.

Salinas, E.and Abbott, L. F. (1994) Vector reconstruction from firing rates. *J. Computational Neurosci.* **1**, 89-107.

Salinas, E. and Abbott, L. F. (1995) Transfer of information between sensory and motor networks. *J. Neurosci.* **15**:6461-6474.

Traub, R. D., Miles, R., Muller, R. U., Gulyas, A. I. (1992) Functional organization of the hippocampal CA3 regions: implications for epilepsy, brain waves and spatial behavior. *Network* **3**, 465.

Wan, H. S., Touretzky, D. S., Redish, A. D. (1994) Towards a computational theory of rat navigation, in *Proceedings of the 1993 Connectionist Models Summer School* (eds. M. C. Mozer, P. Smolensky, D. S. Touretzky, J. L. Elman, A. S. Weigend) pp. 11-19 (Lawrence Erlbaum Associates, Hillsdale, NJ).

Wilson, M. A, and McNaughton, B. L. (1993) Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055-1058.

Wilson, M. A. and McNaughton, B. L. (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676-679.

Worden, R. (1992) Navigation by fragment fitting: a theory of hippocampal function. *Hippocampus* **2**, 165.

Zipser, D. and Andersen, R. A. (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* **331**, 679.

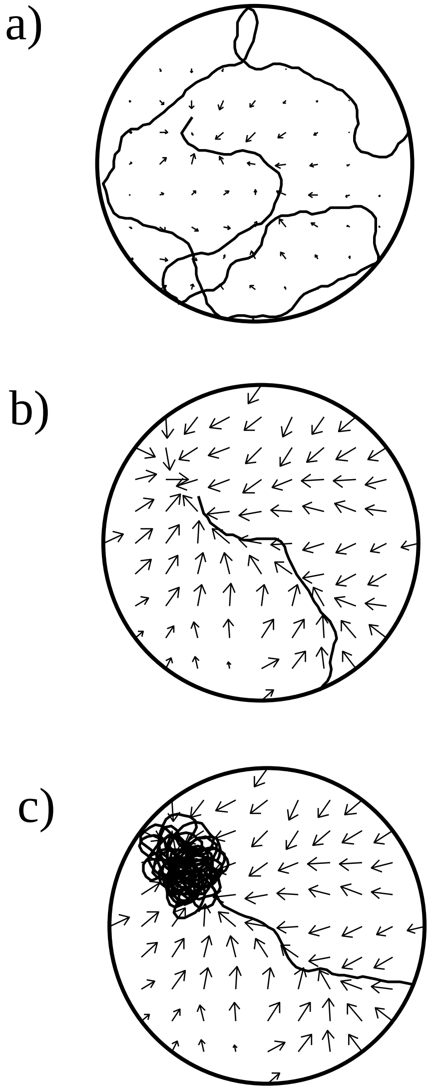

Fig. 1

The path followed and underlying navigational map in a simulated Morris water maze. Starting positions were chosen at random locations along the perimeter of the 1m ‘tank’. Runs proceeded until the computed path intersected the 10cm diameter platform or until 100s had elapsed. At the outer wall an inward radial component was added to the velocity to simulate a rebound. The arrows show the navigational map consisting of the shifts between coded and actual positions plotted on a grid of actual positions (see the section on mathematical results). Longer arrows have been compressed for visual clarity. a) The 2^{nd} path and navigational map of a typical run. b) The 20^{th} path and navigational map of the same run. c) The 21^{st} path and navigational map of the same run, when no platform was present.

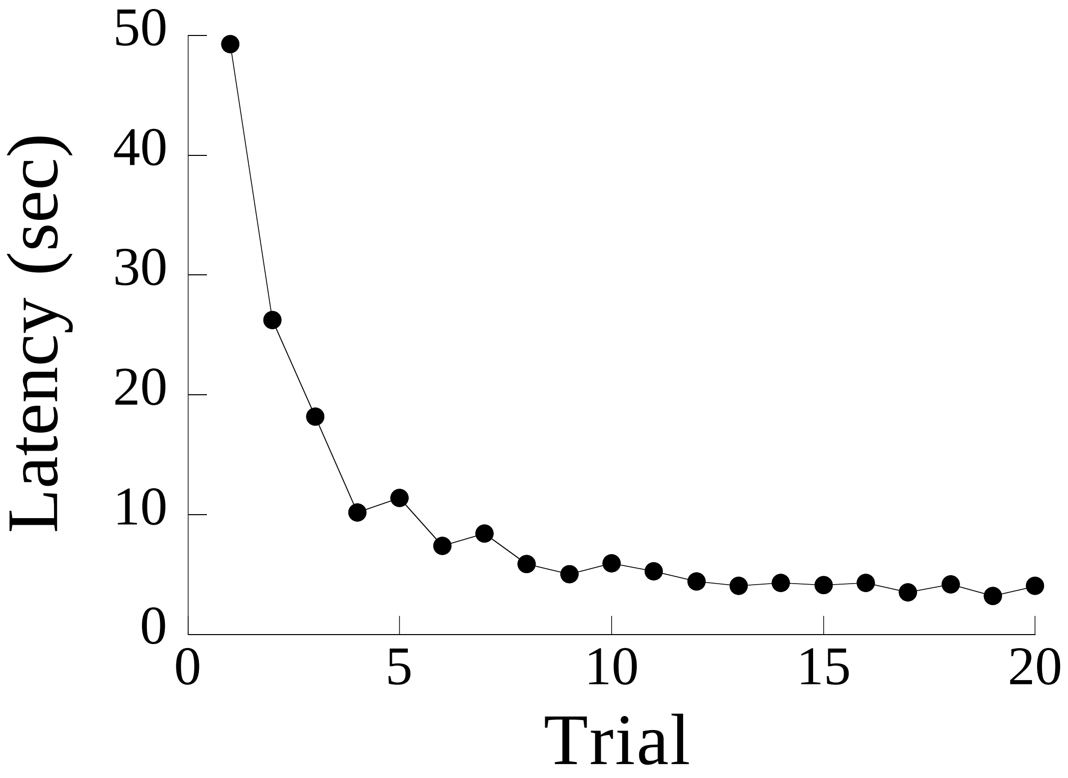

Fig. 2

The escape times averaged over 40 runs as a function of trial number. Each run consisted of 20 trials as in Fig. 1. No decrease in latency was observed when simulated LTP or guidance by the navigational map was removed (results not shown).