LEARNING NAVIGATIONAL MAPS THROUGH POTENTIATION AND MODULATION OF HIPPOCAMPAL PLACE CELLS

Wulfram Gerstner[1] and L.F. Abbott

Volen Center for Complex Systems
Brandeis University
Waltham, MA 02254-9110

Abstract

We analyze a model of navigational map formation based on correlation-based, temporally asymmetric potentiation and depression of synapses between hippocampal place cells. We show that synaptic modification during random exploration of an environment shifts the location encoded by place cell activity in such a way that it indicates the direction from any location to a fixed target avoiding walls and other obstacles. Multiple maps to different targets can be simultaneously stored if we introduce target-dependent modulation of place cell activity. Once maps to a number of target locations in a given environment have been stored, novel maps to previously unknown target locations are automatically constructed by interpolation between existing maps.

---

[1]present address: TU München, Theoretische Physik, 85747 Garching, Germany

1 Introduction

The presence of cells in the hippocampus of the rat that are responsive to spatial location (O'Keefe and Dostrovsky 1971; O'Keefe and Nadel 1978) has led to a number of suggestions about the role of the hippocampus in navigation (O'Keefe and Nadel 1978; McNaughton *et al.* 1991; Muller *et al.* 1991; Traub *et al.* 1992; Worden 1992; Hetherington and Shapiro 1993; Burgess *et al.* 1994; Wan *et al.* 1994; Blum and Abbott, 1996). A recurring idea is that hippocampal place cells provide an environmental map that aids in navigation (O'Keefe and Nadel 1978l; Muller *et al.* 1991; Traub *et al.* 1992; Blum and Abbott, 1995). It has been shown that a navigational map can be formed by the potentiation of synapses between hippocampal place cells arising from the exploration of an environment (Blum and Abbott, 1996). This model of formation of a navigational map relies on two key features of the experimental data; population coding of spatial location by large numbers of place cells (Wilson and McNaughton, 1993) and a temporal asymmetry in the conditions required to produce synaptic long-term potentiation (Levy and Steward 1983; Gustafsson *et al.* 1987; Levy 1989; Debanne *et al.* 1994; Markram and Sakmann, 1995).

A navigation map for a given environment must provide all the information needed to get from any initial point to a specified final target position. If obstacles lie between some of the starting positions and the target, the map must include strategies for getting around them. The basic idea can be visualized in the following way. We start from a spatial map, that is, a detailed representation of the environment including obstacles and walls. The learning rule then adds 'arrows' which give the local direction of motion from any point of the environment towards the target. Such a representation with arrows will be called a navigational map.

In section 3 of this paper, we extend the previous analysis of navigational map formation (Blum and Abbott, 1996) to include an environment with an obstacle. We show how a long-range component of the map, including information about obstacles, develops during random exploration. We also discuss the relationship between the temporal asymmetry of the induction of long-term potentiation and depression (LTP and LTD) and short-range-static and long-range-dynamic components of the map.

To be of practical value, a navigational map should provide directions to a number of learned target locations and, in addition, interpolate between learned locations to guide locomotion to novel targets. The maps discussed in section 3 only provide directions to a single target location. To overcome this limitation, we introduce the idea of target-dependent modulation of place cell activity in section 4. Modulation of receptive fields is an effective way of handling information about multiple spatial locations or directions (Zipser and Andersen, 1988; Pouget and Sejnowski, 1995; Salinas and Abbott, 1995). Experimentally, modulation has been seen, for example, in parietal areas (Andersen and Mountcastle, 1983; Andersen *et al.*, 1985). In section 4 we use the idea of modulation and apply it to hippocampal place cells. We show that modulation of place cell activity by target location allows a network of place cells to simultaneously encode maps to a number of different targets. Modulation by a specific learned target then results in recall of the appropriate map. Modulation corresponding to a target location never previously encountered produces a new map that interpolates between the learned target locations to provide directions to the novel target. In this analysis, we distinguish between knowledge of the location of a target, and knowledge of the path that must be taken from an arbitrary initial position to reach that target. The former is assumed to be known while the latter is provided by the navigational map we study. In the presence of obstacles which prevent a direct approach to the target, getting towards the target can be a nontrivial problem even if the location of the target is known.

Modeling a task like navigation is difficult because many brain regions are likely to be involved. However, it is a worthwhile exercise to explore how much can be achieved by considering a highly reduced model involving known features of hippocampal circuitry and an extremely simple learning procedure. In the model studied, environmental learning is based exclusively on random exploration that results in chance encounters with both targets and obstacles. By using this model, we do not mean to imply that a rat is restricted to such a simplistic scheme when it solves similar problems. Rather we wish to examine how much can be achieved using a limited number of assumptions and guesses about both neuronal circuitry and learning strategies.

2 The Network Model

The model we study is based on the assumption that hippocampal place cell activity provides a population code for spatial location (Wilson and McNaughton, 1993). This means that the firing rates of large numbers of place cells can be combined (by a procedure described in subsection 2.4) to yield a spatial location that we call the encoded location. Initially, we assume that the encoded location is identical to the actual position of the animal in the environment. Thus we start from a purely *spatial* representation. However, long-term potentiation of synapses between place cells (modeling, in particular, those of the CA3 region) that occurs during exploration causes the encoded location to shift slightly. After learning, the location encoded by the collective place cell activity no longer corresponds to the precise location of the animal. Rather, if the learning involved a specific target, place cell activity represents where the animal should go next to get from its present location to the target. The difference between the location encoded by place cell activity and the actual location that evokes that activity provides a *navigational* map directing the animal toward the target from any position in the environment (Blum and Abbott, 1995; Abbott and Blum, 1995; Abbott, 1995).

The basic elements of the model, described below, are: a network of place cells with activities driven by Gaussian functions of the position of the animal, plastic synapses between these place cells, a rule for LTP and LTD induction, and population decoding of the location encoded by place cell activity. A new feature is the target-dependent modulation of place cell activity discussed in section (2.2). The model will be studied in two forms: without modulation in section 3, and with modulation in section 4.

2.1 Place field tuning

We consider a network of N neurons responding to the spatial location of an animal in a two-dimensional environment. The neurons $i = 1, 2, \ldots, N$ of the model network have localized place fields similar to the place cells recorded in the hippocampus. The activity r_i (this may be a firing rate or any other measure of activity relevant for neural coding (Abbott, 1995)) of a given neuron i is high if the animal passes through a specific spatial location. The region of the environment where the activity of a neuron is elevated is its place field. Initially, before learning, we set

$$r_i = A_i f(|\mathbf{x} - \mathbf{x}_i|) \tag{2.1}$$

where $\mathbf{x}$ is a two-dimensional vector representing the location of the animal, $\mathbf{x}_i$ is the center of the place field of neuron i and A_i is an amplitude factor discussed below. We represent the place fields by Gaussian functions

$$f(|\mathbf{x} - \mathbf{x}_i|) = \exp\left(-\frac{|\mathbf{x} - \mathbf{x}_i|^2}{2\sigma_f^2}\right) \tag{2.2}$$

where σ_f is the width of the place field.

We assume that the central locations of the place fields, denoted by $\mathbf{x}_i$ for cell i, are uniformly spaced and that they densely cover the environment. In our analytic calculations, this allows us to replace sums over place cells with integrals over their place field centers. In our simulations, we place these centers on a uniform grid.

2.2 Place field modulation

The amplitude of the response (2.1) is scaled by the factor A_i. In most studies (such as Blum and Abbott, 1995), A_i is a constant representing the maximum response of the cell. In section 3, we will also set A_i equal to a constant value, identical for all cells. However, in section 4 we will explore the idea that the amplitude of the place field response is modulated by the presence of a desired target, e.g., food or the expectation of a food reward. We denote the location of the target by the two-dimensional vector $\mathbf{u}$ and assume that the amplitude of the activity of neuron i is modulated by the factor

$$A_i(\mathbf{u}) = (1 - \alpha) + \alpha\, g(|\mathbf{u} - \mathbf{u}_i|) \tag{2.3}$$

where α is a constant that scales the effect of the modulation. For $\alpha = 0$ the presence or absence of a target does not affect the response of the neuron, whereas for $\alpha > 0$ the target location influences neuron i's activity. Similar to the place field $f(|\mathbf{x} - \mathbf{x}_i|)$ with center $\mathbf{x}_i$, the neuron has a modulation field $g(|\mathbf{u} - \mathbf{u}_i|)$ with center $\mathbf{u}_i$. We will consider two types of modulated cells: cells with aligned place and modulation field centers, $\mathbf{u}_i = \mathbf{x}_i$, and cells with modulation field centers $\mathbf{u}_i$ located at a fixed location, such as near a corner of the environment.

The shape of the modulation field is taken to be either Gaussian

$$g(|\mathbf{u} - \mathbf{u}_i|) = \frac{1}{\sqrt{2\pi}\sigma_g} \exp\left(-\frac{|\mathbf{u} - \mathbf{u}_i|^2}{2\sigma_g^2}\right) \tag{2.4}$$

or triangular

$$g(|\mathbf{u} - \mathbf{u}_i|) = a\,[1 - a|\mathbf{u} - \mathbf{u}_i|]_+ \tag{2.5}$$

where a and σ_g are constants and the notation $[\,.\,]_+$ means that for any quantity s, $[s]_+ = s$ if $s \geq 0$ and $[s]_+ = 0$ if $s < 0$.

Equation (2.3) and the choices (2.4) and (2.5) for the modulation field are theoretical conjectures. At present there are no direct experimental observations which would support the idea of place field modulation by targets.

2.3 Synaptic coupling and learning

All neurons in the network are coupled by synapses that are adjusted during a learning period. Before learning the efficacy J_{ij} of the connection from presynaptic neurons j to postsynaptic neurons i is zero. During learning, the synaptic efficacy increases or decreases according to a 'Hebbian' learning rule. This means that the change in the strength of a given synapse is proportional to the correlation of its pre- and postsynaptic firing rates. In addition, the change in synaptic efficacy depends on the relative timing of presynaptic and postsynaptic activity. Specifically if presynaptic activity occurs at a time t' and postsynaptic activity at time t, the change in synaptic strength is proportional to a factor $H(t - t')$. After a long learning period of duration T, the weight factor for the synapse from neuron j to neuron i is given by (Blum and Abbott, 1996; Abbott and Blum, 1996; Abbott, 1995; Gerstner et al., 1993; Herz et al. 1989)

$$J_{ij} = \int_0^T \int_0^T r_i(t) H(t - t') r_j(t')\, dt\, dt'. \tag{2.6}$$

In accordance with experimental results on long-term potentiation (Levy and Steward 1983; Gustafsson *et al.* 1987; Levy 1989; Debanne *et al.* 1994; Markram and Sakmann, 1995), we assume that synapses are strengthened, if presynaptic activity *precedes* postsynaptic firing. On the other hand, the synaptic efficacy is depressed if the postsynaptic neuron fires *before* the presynaptic neuron. We describe the time-dependence of synaptic plasticity by a function of the form

$$H(t) = \begin{cases} \tau^{-1} \exp(-t/\tau^{-1}) & \text{for} \quad t \geq 0 \\ -\beta\tau^{-1} \exp(t/\tau^{-1}) & \text{for} \quad t < 0 \end{cases} \tag{2.7}$$

that we will refer to as the 'LTP window function'. The parameter β scales the relative importance of LTD. For $\beta = 0$ there is no LTD at all, for $\beta = 1$ the contributions of LTD and LTP are, on average, of equal magnitude. The relative timing of pre- and postsynaptic activity then determines whether a synaptic weight is increased or decreased.

Because of the coupling that arises from learning, the activity r_i of neuron i is influenced by the activity of other neurons. We assume that the coupling remains weak. After learning, the activity of neuron i is no longer given by equation (2.1) but instead by

$$r_i = A_i f(|\mathbf{x} - \mathbf{x}_i|) + \sum_j J_{ij} A_j f(|\mathbf{x} - \mathbf{x}_j|)\,. \tag{2.8}$$

The second term produces the experience-dependent shifts that produce a navigational map (Blum and Abbott, 1996; Abbott and Blum, 1996; Abbott, 1995).

For our simulations, the learning scenario is as follows. The 'animal' moves along straight lines through a two-dimensional environment. To make the model as simple as possible, the environment consists of a large box of extension $L \times L$ where the length L is much larger than the width of the place field ($L \gg \sigma_f$). The animal's movement starts at a random position and proceeds in a random direction. If one of the walls is hit, the direction of movement changes. The new direction is random except that it must lead back into the interior of the box. The movement stops if the target is reached, but learning continues for another 100 time steps while the animal is sitting at the target. Finally the next trial begins at a new random position. A typical set of paths generated in this manner is shown in figure 2.

The idea behind this procedure is that the animal learns only while it is 'hungry' and searching for food. Once it has found the target, it sits at the food source and eats. Afterwards it is satisfied and moves away without further learning. Learning starts again later on when the animal becomes 'hungry' again.

2.4 Population vector coding

We assume that the model network contains a large number of neurons with overlapping place fields. Thus, if the animal is at a location $\mathbf{x}$ many neurons with place field centers close to $\mathbf{x}$ are active at the same time. In order to interpret the resulting pattern of activity, we use a simple vector population decoding scheme (Georgopoulos et al., 1986; Salinas and Abbott, 1994). We take the coded location to be the center of 'mass' of the population activity pattern. In other words, we compute the weighted sum

$$\mathbf{p} = \frac{\sum_i \mathbf{x}_i\, r_i}{\sum_i r_i} \tag{2.9}$$

as the relevant quantity. In the absence of modulation ($\alpha = 0$) and before learning, the population vector represents the *location* of the animal and, if the locations of the spatial fields are dense and uniformly placed,

$$\mathbf{p} = \frac{\sum_i \mathbf{x}_i\, A_i f(|\mathbf{x} - \mathbf{x}_i|)}{\sum_i A_i f(|\mathbf{x} - \mathbf{x}_i|)} \approx \mathbf{x}. \tag{2.10}$$

The approximation $p \approx x$ is good, if there are many cells with overlapping place fields and centers located in such a way that all of the environment (inside *and* outside the box) is uniformly and densely covered. If the centers of all place fields are restricted to lie inside the box, then the 'center of maas' is systematically shifted towards the center of the box. In our simulations, place field centers were assumed to lie on a grid of (11×11) locations including the boundaries (0,0), (0,0.1), ..., (1,1). Place fields had a width of $\sigma = 0.1$. Since the approximation $p \approx x$ is not valid for such a coarse representation of space, we made no use of the approximation in (2.10) in our simulations, but plotted the exact value $\mathbf{p}$ for the inital representation of space (diamonds in Figs. 2-6). For the mathematical analysis presented below, however, we assume densely packed place fields and set $\mathbf{p}$ equal to $\mathbf{x}$.

Under these idealized mathematical conditions, $\mathbf{p} \approx \mathbf{x}$ holds before learning and in the absence of modulation. The population vector changes, however, in the presence of target-induced modulation or through learning-induced changes of synaptic strength. We separate these two effects by calling the shift of the population vector produced directly by modulation $\Delta\mathbf{p}_m$ and that due to learning $\Delta\mathbf{p}_l$. In general, after learning and in the presence of modulation produced by a target at location $\mathbf{u}$, the population vector is

$$\mathbf{p}(\mathbf{x};\mathbf{u},\{X\}) = \mathbf{x} + \Delta\mathbf{p}_m(\mathbf{x};\mathbf{u}) + \Delta\mathbf{p}_l(\mathbf{x};\{X\}) \tag{2.11}$$

where $\{X\}$ represents the family of exploratory paths taken during the training process. The two components of the shift, representing the effects of modulation and the effect of learning, will be discussed in the following sections. Note that, because of these shifts, the population vector no longer encodes the actual position of the animal but a different location nearby. In previous work, (Blum and Abbott, 1996; Abbott and Blum, 1996; Abbott, 1995) and below, we show that this encoded location is where the animal must go if it is to proceed to the target. In other words, $\Delta\mathbf{p}_m + \Delta\mathbf{p}_l$ gives the direction that the animal must follow to reach the target location. When the animal changes its position, the encoded location changes simultaneously always moving slightly ahead of the animal and leading it to the target. The collection of the vectors $\Delta\mathbf{p}_m + \Delta\mathbf{p}_l$ corresponding to all positions $\mathbf{x}$ is thus a navigational map.

3 Navigation To A Single Target

We begin our analysis by studying how a navigationally useful map can arise from the potentiation and depression of synapses between place cells that occur during random exploration of an environment. In this section, we turn off modulation by setting $\alpha = 0$ and study the effect of learning while a single target is present at location $\mathbf{u}$. We concentrate on the population vector $\mathbf{p}(\mathbf{x})$ that arises if the animal is at location $\mathbf{x}$. Before learning, defined as a period of exploration during which synapses are plastic, we have $\mathbf{p}(\mathbf{x}) = \mathbf{x}$ and afterward $\mathbf{p}(\mathbf{x};\mathbf{u}) = \mathbf{x} + \Delta\mathbf{p}_l(\mathbf{x};\mathbf{u})$. We will show that the shifts, $\Delta\mathbf{p}_l$, of the locations encoded by place cell activity that are caused by synaptic changes constitute a useful navigational map. This map includes directions toward a target goal, away from surrounding walls and around obstacles.

The experience-dependent shifts in the encoded population vector can be divided into two terms with different characteristics. We call these static and dynamic components. If the animal sits motionless at some position $\mathbf{x}_1$, neurons with place field centers close to $\mathbf{x}_1$ will be activated and synapses between them will be strengthened due to their correlated firing. If, at a later time, the animal returns to the vicinity of the point $\mathbf{x}_1$, neurons with place fields nearer to $\mathbf{x}_1$ will be more strongly activated than those further from $\mathbf{x}_1$ because they receive additional synaptic input through the previously strengthened synapses. As a result, the population vector for nearby locations is shifted towards $\mathbf{x}_1$. This static learning does not depend on the temporal asymmetry of the LTP window function and it is a local effect that vanishes for distances larger than a few times σ_f.

In addition, there is a dynamic learning term. If the animal moves along a line from $\mathbf{x}_1$ to another location $\mathbf{x}_2$ nearby, synapses from neurons with place fields near $\mathbf{x}_1$ to neurons with fields near $\mathbf{x}_2$ will be strengthened if they are sequentially activated within a time period covered by the LTP window function. However, synapses from neurons with place fields near $\mathbf{x}_2$ to those with fields near $\mathbf{x}_1$ will be weakened because their firing is in the reverse, post-then-pre order. In general, since LTP requires the presynaptic neuron to be active before or during the activity of the postsynaptic neuron, only synapses in the forward direction along the path will increase in strength. As we will see below, the dynamic learning effect gives rise to a long-range component of the navigational map.

3.1 Long-range effect of learning

For the case of a single target in a simple environment, it is possible to perform a fairly general analysis of the effects of LTP and LTD during random exploration on the encoded population vector. As a first step, we study the shift of the population vector evoked when the animal is at a reference location $\mathbf{x}$, $\Delta p_1(\mathbf{x}; \{X\})$, after a learning session which involved movement along a *single* path $X(t)$ with the target fixed at $\mathbf{u}$. For simplicity, we take $X(t)$ to be a straight line along which the animal moves at constant velocity $\mathbf{v}$. We use the parameterization

$$X(t') = \mathbf{x} + D(\mathbf{x})\,\mathbf{v}_\perp + (t' - t_0)\,\mathbf{v} \tag{3.12}$$

where $\mathbf{v}_\perp$ is the unit vector perpendicular to the animals direction of motion (assuming a right-hand rule) and $|D(\mathbf{x})|$ is the perpendicular distance from the reference location $\mathbf{x}$ to the path. Note that $D(\mathbf{x})$ can be positive or negative depending on the direction of motion along the path.

As shown previously (Blum and Abbott, 1996; Abbott and Blum, 1996), learning a path $X(t')$ with the parameters $D(\mathbf{x})$ and $\mathbf{v}$ causes a change of the population vector

$$\Delta \mathbf{p}_1(\mathbf{x}; D, \mathbf{v}) = \gamma\,\sigma_f \left[H_0 \frac{\mathbf{v}_\perp}{|\mathbf{v}|}\, D(\mathbf{x}) + H_1 \frac{\mathbf{v}}{|\mathbf{v}|} \right] \exp\left(-\frac{D^2(\mathbf{x})}{4\sigma_f^2} \right) \tag{3.13}$$

where γ is a constant and $H_0 = \int dt H(t)$ and $H_1 = \int dt H(t) t$ are the mean and first moment of the LTP window function. The first term in the bracket on the right-hand side of (3.13) leads to a change of the population vector that points towards the path, independent of the direction of motion along the path. This is the static learning term discussed above; it depends on $\sigma_f/|\mathbf{v}|$ which is roughly the time the animal spends within any place field that it passes through along the path. The second term in (3.13) is the dynamic learning term; it depends on the direction of the animals motion and generates a component of the population vector pointing along the path. Note that moving the same path with the same speed in the opposite direction changes the sign of the of dynamic learning term, but does not change the contribution of static learning. The magnitude of the dynamic learning term depends on the size of the temporal asymmetry in the LTP window function through the first moment H_1 while the static learning term does not depend on this asymmetry but only on the total magnitude of the LTP window function, H_0.

We now return to the full learning scenario involving multiple paths. For the sake of convenience we restrict our considerations to straight paths with unit velocity $|\mathbf{v}| = 1$. We can then use the direction θ instead of the velocity vector $\mathbf{v}$ as a parameter. θ is the angle between the direction $\mathbf{v}$ and the line connecting the reference position $\mathbf{x}$ and the target $\mathbf{u}$. As before, a second parameter is D, the perpendicular distance of the from the reference point $\mathbf{x}$ introduced above in Eq. (3.12). A given path k with $1 \leq k \leq K$ has therefore parameters D_k and θ_k. After K paths, the effect of learning is

$$\Delta\mathbf{p}_l(\mathbf{x};\mathbf{u}) \propto \sum_{k=1}^{K} \Delta\mathbf{p}_1(\mathbf{x}; D_k, \theta_k) \tag{3.14}$$

The effect of learning increases, of course, with the number of paths and it is therefore sensible to divide by K to get the mean learning effect.

We assume that the learning time T is long enough so that the learning process averages over a large number of paths $K \gg 1$. In this case, the mean shift can be approximated by

$$\begin{aligned} \Delta\mathbf{p}_l(\mathbf{x};\mathbf{u}) &= \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \Delta\mathbf{p}_1(\mathbf{x}; D_k, \theta_k) \\ &= \int dD\ P(D) \int d\theta\ P_D(\theta) \Delta\mathbf{p}_1(\mathbf{x}; D, \theta) \end{aligned} \tag{3.15}$$

where $P(D)$ is the probability of finding a path at distance D from the reference pint $\mathbf{x}$ and $P_D(\theta)$ is the conditional probability that a path at distance D has direction θ. Note that by definition, the direction θ depends on the target location $\mathbf{u}$. More generally, for arbitrary (not necessarily straight) paths we would have

$$\Delta\mathbf{p}_l(\mathbf{x};\mathbf{u}) = \int d\{X\}\ \Delta\mathbf{p}_1(\mathbf{x}; \{X\}) \tag{3.16}$$

where $\mathbf{p}_1$ is the shift due to a single path and the integration covers the ensemble of all possible paths. On the right-hand side of (3.16), the dependence upon the target location only enters due to its effect on the ensemble of allowed paths $\{X\}$ during learning.

We make the following essential observation. Despite the fact that the paths are generated by a random search strategy, some paths are impossible; those that come *from* the target. According to our learning procedure, a path that hits the target is terminated and restarted at a new random location. Thus, there are no paths starting at the target and, as a result, the distribution $P_D(\theta)$ is non-uniform. In the case of straight paths and a circular target of small diameter, the set of allowed paths can be constructed geometrically. It follows from basic symmetry arguments that in an homogeneous and isotropic environment with a sufficiently large number of paths, the shift $\Delta\mathbf{p}_l$ must always point towards the target. Thus we have

$$\Delta\mathbf{p}_l(\mathbf{x},\mathbf{u}) = \frac{\mathbf{u}-\mathbf{x}}{|\mathbf{u}-\mathbf{x}|}\, F(|\mathbf{u}-\mathbf{x}|) \tag{3.17}$$

with some function F that depends on the distance between the location $\mathbf{x}$ and the target $\mathbf{u}$.

If we assume an unbounded environment where all paths originate at infinity and go to infinity unless they hit the target, then the function $F(d)$ falls off as d^{-1} for distances larger than the typical place field size. Thus, learning has a long-range effect that extends local information about the location of the target into a global navigational map. To see how this happens we note that the static learning term does not contribute due to the symmetry of the set of paths (their only asymmetry is with respect to direction, inward or outward from the target, and the static term is not sensitive to direction). Thus, we concentrate on the dynamic learning term. In the sum over all possible paths, all those that miss the target add up to zero. This is because there are two opposite directions of motion for each of these paths and the dynamic effects of learning for the two directions cancel each other. To derive the $1/d$ dependence, it is sufficient to note that, in a two-dimensional environment, the density of paths that hit the target decays with the inverse distance from the target. This is due to the fact that the number of paths entering (minus the number of paths leaving) through a closed loop around the target is the same whatever the size of the loop since all paths start at infinity and can only end at the target. This is a rather general argument that holds for ensembles of arbitrary, not necessarily straight, paths in an unconstrained, homogeneous and isotropic environment. In the case of straight paths, the $1/d$ dependence can also be found from a straightforward, but somewhat lengthy derivation using the explicit formula (3.13) in (3.15) and integrating over all paths that hit the target.

In a finite-sized environment with fixed boundaries and perhaps some obstacles, the vector field of the navigational map generated by learning will no longer be radial, but it will still be long-ranged. Furthermore, all trajectories generated by the vector field will end at the target. A complication in verifying this statement is that, in a finite environment, closed paths (loops) can occur during learning, even if paths are constructed out of straight segments. However, since clockwise and counterclockwise loops are equally likely to occur, learning will average over loops in both directions and, in the limit of large numbers of paths, a rotation free vector field will result. Thus, the sum over learning paths produces a conservative vector field with a sink at the target. It follows, that we can find a potential with a minimum at the target location. If, after learning, the animal moves along the trajectory generated by the vector field, this will correspond to movement downward on the potential surface and it must ultimately result in reaching the target.

Can we say something about the learning time needed to evolve such a vector field? We present a rough estimate. Let us consider a circular environment of radius R and paths which start at the border and are directed towards the target at the center. Each path leads to a minor improvement of the vector field. Relevant changes of the synapses only occur at cells which have place field centers less than σ away from the path (σ is the width of the place fields). If we could chose optimal paths, we would need $(2\pi R)/(2\sigma)$ paths to make sure that all cells with place fields close to the border are hit at least once. What we learn from considerations such as the above is that we have to compare the diameter of a single place field with the length scale of the environment. This is true even in a situation where paths are generated statistically. If we draw paths randomly from the ensemble of all radial paths, we should take 10 times as many paths as above to make sure that the probability that we miss one cell completely is neglilible. Finally, let us consider paths with arbitrary directions, not necessarily pointing towards the target. In this case, the size of the target is an additional parameter that influences the learning time, since a larger target is more likely to be hit than a smaller one. Pathes which miss the target do, after some time, average out. As before, we estimate that we need about $10\,(2\pi R)/(2\sigma)$ succesful trials to build up a reasonable vector field.

3.2 Simulations with a single target

While the analytic treatment of the previous section shows that the learning rules and procedures we are studying will lead to a useful navigational map, it is informative to study more complex environments with computer simulations. For this purpose, we use a square box environment with side length $L = 1$. Place fields have a width of $\sigma_f = 0.1$ and their centers are uniformly spaced on a grid with lattice spacing 0.1. In the simulations with a single target, we use 121 (11×11) neurons.

The 'animal' moves through the box along straight lines at a constant speed of $|\mathbf{v}| = 0.05$ distance units per time step. The movement is randomly scattered back into the box at its borders and stops if a target is reached. The target is square with side length of 0.1. We use an LTP window as in equation (2.7) with time constant $\tau = 10$. The exact form of the LTP window does not matter but a combination of exponential turns out to be computationally convenient. As discussed previously, learning has both static and dynamic components with magnitudes proportional to the mean H_0 and first moment H_1 of the LTP window function respectively. The ratio of H_0 to H_1 and thus the relative contribution of these two components can be adjusted by varying the asymmetry of the LTP window function through the parameter β in (2.7).

The static component of learning causes the population vector to be shifted towards places that are visited more frequently than surrounding areas. We can see this clearly in simulations in which no target is present. In figure 3a we have set $\beta = 0$

to enhance the static component. After random exploration in a box with no target, the dominant features seen in the navigational map are the walls of the box. The arrows in figure 3a, representing the shift in the population encoded location due to exploration and synaptic plasticity, point away from the walls. The encoded location evoked by passing through locations near the borders of the box is shifted towards the interior of the box because the border acts as a repeller during exploration. The shift at the border vanishes, if we remove the static component by setting $\beta = 1$ which makes $H_0 = 0$. In figure 3b, the only remaining shifts seen in the map are those due to random inhomogeneities in the distribution of paths during the exploratory and learning period.

We now consider the map that arises when both a target and an obstacle are added to the environment. The environment, which contains a single target and a barrier separating the left and right halfs of the box, and a number of exploratory paths are shown in figure 2. If a strong static component is present during learning in such an environment, it produces a large border effect as in figure 3a that can interfere with the arrows in the map pointing toward the target, especially for targets near the boundaries. On the other hand, the static component helps to drive the population vectors away from obstacles. Also, the static learning effect is helpful when the animal sits at the target after it has found it. For these reasons, we used in all subsequent simulations a parameter setting of $\beta = 0.8$, thus keeping a small contribution from the static component of learning. The value of 0.8 is somewhat arbitrary, but the precise value of β does not matter. A value of 0.7 and 0.8 is just as good.

As mentioned before, dynamic learning gives rise to a long-range component of the navigational map. The long-range component is seen clearly in figure 4. After a period of random exploration and synaptic modification, the resulting navigational map provides a useful guide to the target from any location in the box. The arrows on the right-hand side of the box do not point directly to the target, but upward and around the barrier directing appropriate motion toward the target. Thus, the representation now contains a record of the experiences during exploration that can easily be recalled through readout of the encoded location to guide future movements. Some of the arrows on the right-hand side of the barrier are rather short. We have checked that the overall structure of the map is better, if we increase the resolution of the underlying network of place cells. In an additional simulations we have used have the lattice spacing (21×21 cells) and place fields with half the width ($\sigma = 0.05$) than before. We find that the map is more regular and that arrows on the right hand side and around the tip of the barrier are longer and always point in the correct direction.

4 Multiple Targets and Modulation

Figure 4 shows that the navigational map we are analyzing can be of use in non-trivial environments with a single target. However, equally important is the ability to navigate to multiple targets or to new locations on the basis of experience with previously learned targets. To accomplish this, we must introduce target-dependent modulation of place cell activity.

Target-dependent modulation allows different pools of neurons to deal with different target locations. When a modulation field is added to the place fields we have been discussing (through the multiplicative amplitude term) neuronal activity depends both on the location of the animal relative to the place field of the cell and on the location of the target relative to the center of its modulation field. Neurons are most active when the target is located inside their modulation field. As a result, targets that are in different regions of the environment are handled by different subsets of neurons. Thus, we might expect a network of modulated neurons to be useful for encoding maps to multiple targets. During learning, the presence of the target in a particular location activates a set of cells which then encode a map to that target within their modified synapses as discussed in the previous section. The presence of a target in a different location then activates a second set of neurons which encode a second map, and so on. Later, the presence of a target at one of the previously learned positions will activate the appropriate group of cells and the map can be recovered.

From the above discussion it might appear that a network capable of encoding multiple targets will require all possible combinations of place and modulation field centers. In fact, we will show in the following that only two types of modulated cells are required. One type is composed of neurons with modulation field centers $\mathbf{u}_i$ that are indepedent of their place field centers $\mathbf{x}_i$. Because we place the modulation field centers for these cells near corners of the environment we call them corner cells. These neurons are capable of developing long-ranged maps to multiple targets. However, when we attempt to interpolate between learned target positions, we find that local errors can arise. This problem is eliminated by adding a set of neurons with modulation field centers equal to their place field centers, $\mathbf{u}_i = \mathbf{x}_i$. We call these target cells. These introduce a local component to the map, $\Delta\mathbf{p}_m$, that is independent of any learning process and arises purely from modulation. In combination, the corner and target cells provide an effective map to multiple targets that can also interpolate between learned locations to find new targets.

4.1 Effect of amplitude modulation during learning

We will now study how target-dependent modulation modifies the learning process. We start by concentrating on the change of synaptic weights that occurs during learning from an ensemble of exploratory paths $\{X\}$ while the target is at a fixed location $\mathbf{u}$. We define J^0_{ij} to be the synaptic efficacies that would arise through learning in the absence of modulation, i.e. with $\alpha = 0$. A modulation term in the amplitude factor of equation (2.1) of the form (2.3) modifies the synaptic weights produced by identical exploration to

$$J_{ij}(\mathbf{u}) = J^0_{ij}\,[(1-\alpha) + \alpha\, g(|\mathbf{u}-\mathbf{u}_i|)]\,[(1-\alpha) + \alpha\, g(|\mathbf{u}-\mathbf{u}_j|)]. \tag{4.18}$$

Note that for $\alpha \to 1$, only those neurons with the center of the modulation field $\mathbf{u}_i$ close to the target $\mathbf{u}$ are affected by learning.

Our purpose in introducing target-dependent modulation is to determine whether navigational maps to multiple targets can be learned and recalled. What happens if several targets are presented sequentially during the training period? We assume that for each target the learning time is long enough to ensure averaging over many paths. If the targets have positions $\mathbf{u}^{(n)}$ labeled by n, the total effect of learning is the superposition of all the individual contributions

$$J_{ij} = \sum_n J_{ij}(\mathbf{u}^{(n)}). \tag{4.19}$$

After learning several target positions $\mathbf{u}^{(n)}$, the network is tested with a new target position $\mathbf{u}$. We want to calculate the shift $\Delta \mathbf{p}_l(\mathbf{x};\mathbf{u})$ of the population vector induced by the weigths (4.19). To linear order in J_{ij} the shift is (Abbott and Blum 1996)

$$\Delta \mathbf{p}_l(\mathbf{x};\mathbf{u}) = \frac{\sum_{i,j}(\mathbf{x}_i - \mathbf{x}) J_{ij} r_j(\mathbf{x})}{\sum_j r_j(\mathbf{x})} \tag{4.20}$$

where J_{ij} is given by (4.18) and (4.19) while r_j is given by (2.1) and (2.3).

To keep the formulas as simple as possible, we consider a densely packed network of place cells and assume that there are no correlations between the centers of the place field and the modulations field. In this case, equation (4.20) can be evaluated and yields

$$\Delta \mathbf{p}_l(\mathbf{x};\mathbf{u}) = \sum_n \Gamma(|\mathbf{u} - \mathbf{u}^{(n)}|)\, \Delta \mathbf{p}_l(\mathbf{x};\mathbf{u}^{(n)}) \tag{4.21}$$

where $\Delta \mathbf{p}_l(\mathbf{x};\mathbf{u}^{(n)})$ is the single target result (3.16) without modulation ($\alpha = 0$) and $\Gamma(.)$ is some factor. Eq. (4.21) shows that the resulting shift is, as expected, a linear superposition of the previously learned navigational maps. The weight of the map pointing towards the target $\mathbf{u}^{(n)}$ is given by the function

$$\Gamma(|\mathbf{u} - \mathbf{u}^{(n)}|) = \int d\mathbf{u}'\ \rho(\mathbf{u}')[(1-\alpha) + \alpha\, g(|\mathbf{u}' - \mathbf{u}^{(n)}|)]\,[(1-\alpha) + \alpha\, g(|\mathbf{u}' - \mathbf{u}|)] \tag{4.22}$$

where $\rho(\mathbf{u}')$ with $\int \rho(\mathbf{u}') d\mathbf{u}' = 1$ is the density of cells with modulation field center $\mathbf{u}'$. Both the utility and limitations of this simple linear summation of individual maps will be seen in the simulation results.

4.2 Local effects of amplitude modulation

In order to estimate the shift $\Delta \mathbf{p}_m$ of the population vector that is introduced by the target directly through modulation, we assume $\alpha \ll 1$ and linearize (2.9) in α to find

$$\Delta \mathbf{p}_m(\mathbf{x},\mathbf{u}) = \frac{\alpha}{\sum_k f(|\mathbf{x} - \mathbf{x}_k|)} \sum_i (\mathbf{x}_i - \mathbf{x}) f(|\mathbf{x} - \mathbf{x}_i|)[g(|\mathbf{u} - \mathbf{u}_i|) - 1]. \tag{4.23}$$

Recall that each neuron is characterized by two parameters, the center of its place field and of its modulation field. For large numbers of uniformly distributed cells, the sum over i can be replaced by a sum over all combinations of $\mathbf{x}_i$ and $\mathbf{u}_i$. Thus, we make the replacement

$$\sum_i \longrightarrow \int d\overline{\mathbf{x}} \int d\overline{\mathbf{u}}\, \rho(\overline{\mathbf{x}}, \overline{\mathbf{u}}) \tag{4.24}$$

where $\rho(\overline{\mathbf{x}}, \overline{\mathbf{u}})$ is the fractional density of cells with place field centers with $\mathbf{x}_i = \overline{\mathbf{x}}$ and $\mathbf{u}_i = \overline{\mathbf{u}}$.

In order to proceed further, we must specify the distribution of place fields. We consider the two cases discussed in the introduction:

(i) Cells with fixed modulation field centers so that the distributions of place fields and modulation fields are uncorrelated, $\rho(\overline{\mathbf{x}}, \overline{\mathbf{u}}) = \rho(\overline{\mathbf{x}})\,\rho(\overline{\mathbf{u}})$. In this case, we can perform the integration and find $\Delta\mathbf{p}_m = 0$.

(ii) Cells for which the center of the modulation field is identical to the center of the place field, $\rho(\overline{\mathbf{x}}, \overline{\mathbf{u}}) = \rho(\mathbf{x}_i)\,\delta(\overline{\mathbf{u}} - \overline{\mathbf{x}})$. If we take a uniform distribution $\rho(\overline{\mathbf{x}}) = \rho_0$ and g to be a Gaussian, Eq. (2.4), then we can do the integrals and find from (4.23)

$$\Delta\mathbf{p}_m \propto (\mathbf{u} - \mathbf{x}) \exp\left[-\frac{(\mathbf{u} - \mathbf{x})^2}{2(\sigma_f^2 + \sigma_g^2)}\right]. \tag{4.25}$$

The shift points from the location $\mathbf{x}$ towards the target $\mathbf{u}$. In other words, the presence of the target induces a radial vector field with direction towards the target. The amplitude of the vectors is significant in the neighborhood of the target only. For locations $\mathbf{x}$ far away from the target, $|\mathbf{x} - \mathbf{u}|^2 \gg \sigma_f^2 + \sigma_g^2$, the effect vanishes. Thus, we have a purely *local* effect that looks much like the local effects of learning, but in this case no exploration or synaptic change is required.

4.3 Simulations with Modulation

The simulations were run as in the case with a single target. We use the same set of place fields, but for each place field there are five neurons with different centers of the modulation field. The total number of neurons is $N = 121 \times 5 = 605$. In each group of five neurons, the first neuron has a modulation field center identical to the place field center ($\mathbf{u}_i = \mathbf{x}_i$). These neurons are the target cells. The other four neurons are the corner cells with modulation field centers at $(.25, .25), (.25, .75), (.75, .25), (.75, .75)$. The shape of the modulation field is triangular, Eq. (2.5) with parameter $a = 2$. To get a significant modulatory effect we set $\alpha = 1$.

Figure 5a shows the local effect of modulation (without learning) due to a single target. The arrows in figure 5a point radially towards the target but their amplitude of the arrows is negligible beyond the width of the modulation field ($1/a = 0.5$). Recall that this effect is due to neurons with aligned place field and modulation field centers. It arises simply because neurons close to the target are modulated to higher activity levels than neurons away from the target and thus the encoded location is shifted toward the target.

In order to study the effect of learning, we temporarily turn off the local effect by disregarding the contribution of the target cells. We are interested in studying both the encoding of multiple target maps and interpolation between learned targets. The network has been trained using four target positions at (.25,.25), (.25,.75), (.75,.25), and (.75,.75). During the learning session, each target has been presented 50 times. The network was then tested with two of the previously known targets at (.25,.25), (.25,.75) and one unknown target at (.25,.50). It is obvious from figures 5b and 5c, that presentation of a previously encountered target generates the desired vector fields pointing towards the target. When the novel target is presented in figure 5d, however, the situation is less clear. As predicted from theoretical considerations, the superposition of the two vector fields maps to known targets yields a combined vector field where all the vectors far from the new target point towards it, but local uncertainties arise. This is seen in the region directly to the right of the target where the arrows are extremely short and a few, in fact, point away from the target.

As mentioned in the introduction, local errors, like those seen in figure 5d, can be eliminated by adding back the effects of the target cells as shown in figure 5a. In figure 6a, we combine the local effect generated by the target cells and the long-range effect generated by learning through the corner cells. We now have a vector field that is correct both in the local neighborhood and far from the target. An accurate map to any one of a set of targets can be recalled by the presence of a given target, and maps to new target locations are generated through interpolation. Additional tests demonstrated accurate interpolation to targets located anywhere in the box.

In a final simulation, we have combined the task of obstacle avoidance with multiple targets (Figure 6b). During learning, two targets (dashed squares in Fig. 6b) at (0.25,0.25) and (0.25,0.75) have been presented and the model rat performed 100 succesful search trial to each of these targets. After learning, we tested with a new target at an intermediate postion (0.25,0.5). The resulting navigational map has indeed a long range component which guides the the rat around the barrier to the target. The long range effect arises from a superposition of the previously learned vector fields. There is also a short range component which guides the animal to the new target once it is close to it. In this simulation, the target cells had a Gaussian modulation field of width $\sigma_g = 0.1$. Thus the effect of the target cell is negligible beyond a distance of 0.2 units from the target.

5 Discussion

Our analysis of map formation by hippocampal place cells has involved only a few basic principles: population encoding of spatial position, correlation-dependent, temporally asymmetric modification of synapses, random exploration of an environment and target-dependent modulation of place cell activity. With these simple components, a network of place cells can develop maps that indicate the direction toward a target around obstacles and away from walls. These maps are extremely easy to follow; the animal merely has to head from its present location toward the location encoded by place cell activity. As it moves, the encoded location will lead the actual location of the animal along the paths stored in the map.

Multiplicative modulation of neurons involved in population coding is a powerful mechanism for increasing computational complexity. We have used this principle and applied it to hippocampal place cells. In our case, target-dependent modulation allows a place cell network to encode multiple maps and recall the appropriate map when a given target appears. Furthermore, knowledge of an environment represented by multiple maps to different targets is sufficient to allow a map to a novel target to be constructed by interpolation. While far from a complete theory of navigation, these results show that significant progress can be made using a small number of established properties of neuronal coding and plasticity and an extremely simple learning strategy. It will be interesting to explore how more sophisticated learning procedures would augment these findings.

In our present model, we have made a clear distinction between an idealized exploration phase and a separate test or navigation phase. During the exploration phase, the synaptic weights were continuously modified according to Eq. (2.6). The change of weights results in a change of neuronal activity and, hence, in a shifted population vector. The shift of the population vector, however, was *not* used to improve target approach during the following exploration trials. Rather, all exploratory paths were based on the same random search strategy outlined in Section 2.2. In contrast, for navigation during the 'test' phase we only considered the shift of the population vector without any random component.

The assumption of a long exploration phase (in simulations we used about 100 independent trials) is, of course, rather artificial, but it has enabled us to address the formation of navigational maps analytically. A further advantage of this approach is that the model rat can not end up trapped in a cyclic path, since contributions from loops will average out.

Real rats use the knowledge from the first few trials immediately to improve preformance in the following search trials. Blum and Abbott (1996) have simulated this by a search strategy where the choice of exploratory paths is influenced by both a random component *and* the previous shift of the population vector. This reduces the number of paths need to form a navigational map to around 10-30 in good agreement with experimental results. Such an approach of strategy improvement during the exploration phase is similar to the well studied models of reinforcement learning (Barto et al. 1983, Dayan 1992) policy iteration in dynamic programming although we have not been able to show the connection analytically.

Our learning scheme relies on a fundamental asymmetry: learning occurs while the animal is 'hungry' and searching for a target, but does not occur while it moves away in a satisfied state. This is the essential effect which makes the population vector shift towards the target. But there is also a symmetric effect: Even if there is no target, the static component of learning drives the population vector away from walls and obstacles. As a result of learning, a smooth navigational map develops. This implies that the animal does not move along the shortest possible connection around a barrier to the target, but on a nicely curved line. The above result holds for a dense and uniform distribution of place fields combined with a random search strategy. Abbott and Blum (1996) have shown that if learning continues during the 'test' phase the animal learns to cut corners and evolves navigational strategies with shorter paths.

Our simple model was not intended to explain all known effects in hippocampus and navigation. There are some obvious points where we can improve the model. We mention three issues: (i) Presently we have started from a spatial representation which was turned into a navigational map by learning. In would be interesting to go one step back and learn the spatial representation itself. In particular, it would be nice to learn where to locate the centers and how to adjust the width of place fields. The picture would be even more complete if we would generate place fields from views of the environment. (ii) Presently place fields have no directional selectivity. Also the model rat was assumed to turn instantaneoulsy when it hits a wall and starts in a new random direction. What is the relevance of direction? This is another topic which should be addressed. (iii) Presently, the learning mode is triggered by a 'hungry' animal in search for a target. Therefore unrewarded exploration cannot generate a navigation map, except for the stationary learning effect which drives the animal away from walls. Can we generalize the argument so a to include 'latent learning'? One possibility is to introduce an attention variable whose value triggers the learning mode.

In our present model, we have studied two mechanisms, viz. learning and modulation, which can shift the population vector and thus guide navigation. If the same principles are used in real systems, it should be possible to observe experimentally the following effects.

(i) Modification of lateral connectivity during learning leads in our model to a deformation of the place field and a change of place cell activity. As a result, the population vector (the weighted sum over the activities of many place cells) is shifted away from the walls and towards the target.

(ii) We have proposed modulation of place cell activity in the presence of a target. This can be tested on the single neuron level. Also, if this is true, the population vector should be shifted towards the target even without exploration merely by showing the target or by providing information about the target location.

It would be interesting to find out how this compares with experiments. It seems that experiments of Speakman and O'Keefe are in contradiction with item (ii).

6 References and Notes

Abbott, L.F. (1995) Decoding Neuronal Firing and Modeling Neural Networks. Quart. Rev. Biophys. **27**, 291-331.

Abbott, L. F. and Blum, K. I. (1996) Functional significance of long-term potentiation for sequence learning and prediction. Cerebral Cortex (in press).

Andersen, R.A., Essick, G.K., Siegel, R.M. (1985) Encoding of spatial location by posterior parietal neurons. Science **230**, 450-458.

Andersen, R.A., Mountcastle, V.B. (1983) The influence of the angle of gaze upon the excitability of light-sensitive neurons of the posterior parietal cortex. J. Neurosci. **3**, 532-548.

Barto, A. G., Sutton, R. S., and Anderson C. W. (1983) Neuronlike adaptive elements that can solve difficult learning problems. IEEE Trans. SMC **13**:835-846

Blum, K. I. and Abbott, L. F. (1996) A model of spatial map formation in the hippocampus of the rat. Neural Comp. **8**, xxx-xxx.

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

Dayan, P. (1992) The convegence of TD(λ) for general λ. Machine learning **8**:341-362

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.

Gerstner, W., Ritz, R., and van Hemmen, J. L. (1993) Why spikes? Hebbian learning and retrieval of time–resolved excitation patterns. *Biol. Cybern.* **69**, 503–515.

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.

Herz, A. V. M., Sulzer, B., Kühn, R., and van Hemmen, J. L. (1989) Hebbian learning reconsidered: Representation of static and dynamic objects in associative neural nets.", *Biol. Cybern.* **60**, 457–467.

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.

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.

Markram, H. and Sakmann, B. Action potentials propagating back into dendrites trigger changes in efficacy of single-axon synapses between layer V pyramidal neurons. *Soc. Neurosci. Abstr.* **21**, 2007 (1995).
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.

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).

Pouget, A., Sejnowski, T.J. (1994) A neural model of the cortical representation of egocentric distance. Cerebral Cortex **4**, 314-329.

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 Coded Information from Sensory to 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.

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.

Acknowledgments: Research supported by National Science Foundation grant NSF-DMS9503261 and the W.M. Keck Foundation (LA). W.G. has been supported by the Human Frontier Science Foundation (Short-term fellowship program, SF380/95).

7 Figure Captions

Figure 1: Place field tuning. The activity r_i of a cell i with the center of the place field at $\mathbf{x_i}$ is shown as a function of the location $\mathbf{x}$ of the animal. The activity is high if the animal is close to $\mathbf{x_i}$. The amplitude of the place cell activity is modulated in the presence of a target. If the target is at at the optimal location $\mathbf{u_i}$, the activity is given by the solid line. If the target is at different location which is not optimal for cell i, then the activity is always lower (dashed line).
Figure 2:
Four samples of exploratory paths. The environment consists of a square box with a barrier (dotted) in the center. Each exploration starts at a random location, crosses the interior along straight lines, is scattered at the borders, and ends at the target indicated by a small square at (.25,.25). Note that in one trial the movement is trapped for a long time on the right side of the barrier, before the model rat moves to the left and finds the target.
Figure 3:
Static and dynamic components of learning. The shift $\Delta\mathbf{p}_l$ of the population vector after a long exploration period whithout targets is indicated by pointers on a grid of 81 (9×9) positions. Each pointer consists of a line connecting the position of the population vector before learning (diamonds) with the position after learning.
a) The static component of learning shifts the population vector always away from the walls. The asymmetry parameter β in Eq. (2.7) has been set to zero.
b) If the static component of learning vanishes ($\beta = 1$), then a random exploration leads to small statistical shifts of the population vector, but there is no systematic effect.
Figure 4:
Learning a single target position. The pointers show the shift of the population vector after a learning session of 100 exploration trials. The field of pointers forms a navigational map which guides the animal from arbitrary initial positions to the target.
Figure 5:
Multiple targets and amplitude modulation.
a) Effect of target cells. The modulation field of a target cell shifts the population vector of surrounding locations towards the target at (.25,.50). Note that the effect is short-range and occurs without learning.
b) - d) Effect of corner cells, no target cells present. Four different target locations, have been learned during a learning session consisting of 50 exploration trials for each target. If one of the known targets is presented after learning, the population vector is shifted towards the correct target position. In b) the target, indicated by a small square, is at (.25,.25), in c) it is at (.25,.75). d) A target at a new position between the previously known locations evokes a shift field which is the superposition of the shifts in b) and c). Note that the effect is long-range, but close to the target information is missing.

Figure 6:
Unknown target position. A network of 5×121 neurons (4 corner cells and 1 target cell at each of the 11×11 grid locations) has been trained with four different target positions close to the corners of the box. A target at a previously unknown position evokes a shift field pointing towards the target. a) The shifts are a superposition of the local effects of the target cells (Fig. 5a) and the long-range effects of the corner cells (Fig. 5d). b) The same principles also work in the presence of an obstacle.