Parallel Hopfield Networks

Robert C. Wilson
rcwilson@seas.upenn.edu
Department of Bioengineering, University of Pennsylvania, Philadelphia,
PA 19103, USA

We introduce a novel type of neural network, termed the parallel Hopfield network, that can simultaneously effect the dynamics of many different, independent Hopfield networks in parallel in the same piece of neural hardware. Numerically we find that under certain conditions, each Hopfield subnetwork has a finite memory capacity approaching that of the equivalent isolated attractor network, while a simple signal-to-noise analysis sheds qualitative, and some quantitative, insight into the workings (and failures) of the system.

1 Introduction

The Hopfield network (Hopfield, 1982) is a milestone in computational neuroscience due to its conceptual simplicity, effectiveness as an associative memory, and analytic tractability. Despite some biologically implausible assumptions (such as a symmetric weight matrix, full connectivity, uniform axon delays, and neurons with both excitatory and inhibitory character), the concept of the attractor is a powerful idea, and there is evidence that attractors exist in the brain (Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005). In the light of this, many authors have worked to remove the nonbiological constraints on Hopfield networks (Derrida, Gardner, & Zippelius, 1987; Sompolinsky, 1986).

Of particular relevance to this letter is the work of Herz and coworkers (Herz, Li, & van Hemmen, 1991), who extended the Hopfield formalism to include nonuniform transmission delays between neurons. In particular, they showed that for transmission delays with certain properties, there exists a Lyapunov function for the dynamics of the delayed Hopfield network. In contrast to the regular Hopfield network, however, memories, rather than being just static patterns of activity, are now represented by spatiotemporal patterns of spiking.

Precise temporal relationships between spikes on different neurons are also important in the synfire chain model of Abeles (1991), which postulates that synchrony across groups of neurons is the computational currency of the brain. This hypothesis has yet to be proved experimentally, but a growing body of evidence suggests that precisely timed spikes do exist in

the brain (Foster & Wilson, 2006; Ikegaya et al., 2004; Shmiel et al., 2005; Meister, 1996).

Extending the concept of synchrony to include networks with nonuniform axon delays leads to the polychronization network of Izhikevich (2005). In this model, the presence of nonuniform delays means that synchronous spiking no longer propagates through the system. Instead the delays allow specific asynchronous spike patterns to travel through the network if the axon delays are such that for each neuron in the pattern, the asynchronicity in the input spikes is cancelled by the axon delays and the spikes arrive at the next cell in the sequence in synchrony.

A related model is the concurrent recall network (CRN) of Wills (2004). This model is similar to the polychronization network in that it uses nonuniform axon delays to store asynchronous memories, but it uses very different neurons. In this model, each neuron has a set of conjunction detectors, each of which receives a set of delayed inputs from presynaptic neurons. The activation rule then has two stages. First, a conjunction detector will activate if it receives enough synchronous input to pass a threshold. Then the neuron will fire if a certain number (usually set to 1) of the conjunction detectors are active within a given time window (e.g., 1 ms).

Interestingly, recent experiments (Gasparani & Magee, 2006) have found that dendritic sodium spikes can carry out a very similar function. In particular, a dendritic spike will occur only when the neuron receives synchronous input at a localized spot on the dendrite and, furthermore, the dendritic spike will cause a somatic neural spike with about 80% probability.

Wills (2004) used these elements to create a network capable of storing and recalling many spatiotemporal spike patterns. Furthermore, the conjunction detector activation rule allowed enough separation between different memories such that many different patterns could be recalled concurrently—hence, the name *concurrent recall networks*.

A diagram demonstrating the output of a CRN is shown in Figures 1A and 1B. In Figure 1A, we show the spike pattern produced by a CRN recalling just one periodic memory. The gray boxes indicate what we term the CRN mask, which are the time points at which we expect the neurons to fire if the CRN network is performing properly. In Figure 1B, we show the same network simultaneously recalling two CRN memories with different spatiotemporal characteristics. For clarity we denote the CRN mask and spikes associated with the second memory by gray circles and black "plus" marks, respectively.

One disadvantage of these precise timing networks is that, in relation to the Hopfield model, they are rather ferromagnetic in nature, in that each memory (spatiotemporal pattern) is either activated or not activated. A more general approach would be to have a network where each spatiotemporal pattern can be activated in different ways, each time activating a subset of the neurons at the mask points. In this letter, we create such a network. In particular, by perturbing the weights of the
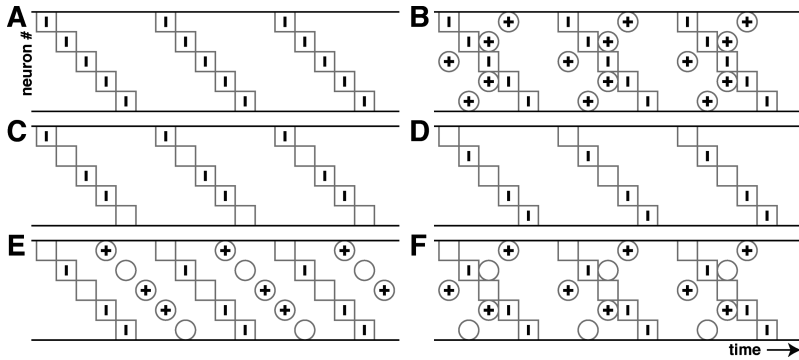
Figure 1: Diagram illustrating the behavior of concurrent recall networks and parallel Hopfield networks. (A) Activity from five neurons in a CRN recalling one periodic memory. Time is on the x-axis and neuron number on the y-axis. The small black bars correspond to spikes, and the gray boxes correspond to the CRN mask. For the CRN network, all of the mask points are occupied with a spike. (B) The same CRN recalling two memories simultaneously. For clarity, spikes associated with the first memory are denoted by solid lines, and spikes from the second memory are associated with solid plus symbols. The CRN mask for the first memory consists of gray boxes, and the gray circles denote the CRN mask for the second memory. (C) Five neurons from a PHN with one subnetwork activated. Notice that this time, the CRN mask boxes are not all filled with spikes and that the same pattern of spiking propagates through time. (D) The same PHN with the same subnetwork activated (hence the same CRN mask) but with a different Hopfield pattern. (E) Double activation of the same subnetwork (hence the same CRN mask for each) with different Hopfield patterns. (F) Concurrent activation of two different subnetworks with two different CRN masks.

connections between neurons, we extend the CRN model in order to store multiple, random, Hopfield-like patterns on each of the spatiotemporal CRN memories. In this case, each CRN memory effectively acts as an independent subnetwork approximating the dynamics of an isolated Hopfield network.

The resulting model has the ability of the CRN to store and simultaneously recall multiple spatiotemporal memories while retaining the ability of the Hopfield network to act as an error-correcting associative memory. The combination of these two functions allows the network to effect the behavior of many independent Hopfield networks in parallel; hence, we term these networks *Parallel Hopfield Networks* (PHNs).

This idea is illustrated in Figures 1C through 1F. As before, the boxes and circles indicate the CRN mask points, where a neuron would be active if the network was a pure CRN network. Notice that unlike Figures 1A and

1B, the networks in these panels do not have all neurons active at the CRN mask points; instead only a subset of them produces spikes.

In Figure 1C, we show a network with a single subnetwork activated with a particular pattern. In Figure 1D, we show the same network with the same subnetwork activated, but in a different way.

Figures 1E and 1F illustrate the ability of the network to access several different memories simultaneously. In Figure 1E, the same subnetwork is activated twice at two different onset times and with two different patterns. The two subnetworks are independent, and so the network is effectively running two different Hopfield network simulations (of the same Hopfield network, but with different initial conditions) in parallel.

Clearly, this is a rather trivial example of parallelism, and for the rest of this letter, we focus on the more interesting case shown in Figure 1F. Here we activate two different subnetworks (with, in general, different CRN masks and different connection weights) that correspond to quite different Hopfield networks. In this case, we demonstrate that the network can be thought of as approximating the dynamics of many different Hopfield networks in parallel.

The remainder of this letter is laid out as follows. In section 2, we outline the technical details of the PHN. In section 3, we present a simple signal-to-noise analysis that sheds some light on where we can expect to see the parallel Hopfield behavior and when we can expect the network to fail. Results of numerical simulations are presented in section 4. We discuss the biological plausibility of these networks in section 5 and conclude in section 6.

2 The Parallel Hopfield Network

In this section, we describe a nuts-and-bolts approach to constructing a PHN. Our aim is not to justify how such a network could occur naturally through learning, but how, as an engineer, one might go about building it.

2.1 Connectivity. We begin by specifying a set of M periodic CRN memories. Note that although a PHN could be implemented in a feedforward manner (the major limitation being the maximum number of iterations available to each Hopfield subnetwork), for simplicity we require that the CRN memories be periodic.

Since the memories are periodic, each subnetwork, labeled μ, is fully specified by a period T_μ; a set of spike times t_i^μ ($0 \le t_i^\mu < T_\mu$), one for each neuron i; and a set of P randomly generated Hopfield patterns, $\eta_i^{\mu p}$, where

$$\eta_i^{\mu p} = \begin{cases} 1 & \text{with probability } b \\ 0 & \text{otherwise} \end{cases} \tag{2.1}$$

denotes the activity of neuron i in Hopfield pattern p on subnetwork μ. Note that we are using [0, 1] neurons rather than the more usual [−1, 1] neurons to emphasize the asymmetry between active and inactive cells.

We then connect the network in the following way. For each subnetwork μ, we create one conjunction detector per neuron. For a given neuron i, this conjunction detector (labeled μ) is wired up to receive inputs from all the other neurons in the network. The axon delay from neuron j to conjunction detector μ on neuron i is given by

$$
\tau_{ij}^{\mu} = \begin{cases} \left(t_i^{\mu} - t_j^{\mu}\right) \bmod T_{\mu} & \text{if } t_i^{\mu} \neq t_j^{\mu} \\ T_{\mu} & \text{if } t_i^{\mu} = t_j^{\mu} \end{cases}, \tag{2.2}
$$

and the weight is

$$
W_{ij}^{\mu} = \begin{cases} \dfrac{1}{b(1-b)N} \displaystyle\sum_{p=1}^{P} \left(\eta_i^{\mu p} - b\right)\left(\eta_j^{\mu p} - b\right) & \text{if } i \neq j \\ 0 & \text{if } i = j \end{cases}, \tag{2.3}
$$

which is the same weight matrix as used in Tsodyks and Feigelman (1988) and Buhmann, Divko, and Schulten (1989) and is balanced on average. This form of the weight matrix ensures that the all-zero state is a stable state of the subnetwork, which is important, as we do not want subnetworks to become spontaneously active.

Note also that this construction requires that for $M > 1$, there are multiple connections from neuron j to neuron i, all going through different conjunction detectors and all with potentially different axon delays and weights. This overfull connectivity is clearly unrealistic, and in section 5, we discuss ways in which it might be removed. For now, however, we press on with it, as it simplifies the analysis.

2.2 Dynamics. First, we compute the activation function for each conjunction detector μ on neuron i at time t,

$$
A_i^{\mu}(t) = \sum_{j=1}^{N} W_{ij}^{\mu} x_j\left(t - \tau_{ij}^{\mu}\right), \tag{2.4}
$$

where W_{ij}^{μ} is given in equation 2.3 and $x_j(t - \tau_{ij}^{\mu})$ is the activity of neuron j at time $t - \tau_{ij}^{\mu}$. Note that unlike the original CRN case (Wills, 2004), we work in discrete time.

The conjunction detector activation rule then takes the following form. A conjunction detector fires if its activation $A_i^{\mu}(t)$ is above some threshold θ_i^{μ}, and the neuron fires if any of its conjunction detectors fire. We denote

the activation of the conjunction detector by y_i^μ, where

$$y_i^\mu(t) = \begin{cases} 1 & \text{when } A_i^\mu(t) \geq \theta_i^\mu \\ 0 & \text{otherwise} \end{cases}, \tag{2.5}$$

and then the neuron is active if any one of the conjunction detectors is active,

$$x_i(t) = y_i^1 \text{ OR } y_i^2 \text{ OR, } \ldots, \text{ OR } y_i^M$$

$$= 1 - \prod_{\mu=1}^{M} \left(1 - y_i^\mu\right). \tag{2.6}$$

As noted in section 1, this form of activation rule is a highly simplified caricature of a dendritic spike (Gasparani & Magee, 2006).

3 Analysis

Given its complexity, it seems unlikely that a complete theoretical analysis of the PHN is feasible. However, by making a series of approximations, we can shed some light on the workings (and failures) of the system.

   We begin with a simple signal-to-noise approach that gives some qualitative insight into the network, before we construct a more detailed, though approximate, self-consistent mean field equation for erroneous spiking in the network.

   3.1 Signal-to-Noise Analysis. Consider a PHN running in the parallel Hopfield state, that is, with one or more subnetworks activated recalling a pattern. If such a state exists, then we can divide the output of the network into two types based on whether it occurred at a CRN mask point. In the ideal case, we would like the output at the CRN mask points to be exactly the same as one of the stored Hopfield patterns, and the activity away from the mask points should be zero.

   There are therefore three ways in which the network could fail: two possible bit flips away from the activated Hopfield pattern at the mask points, $1 \rightarrow 0$ or $0 \rightarrow 1$, and spurious spiking away from the mask points. Our goal is to derive expressions for the initial probabilities of these events; i.e., the probability of a failure at the next time step, given that the network is set up in the parallel Hopfield state. We consider the case of spurious spiking first.

3.1.1 Off Mask Point Failure. We begin by writing the expression for the activation of conjunction detector μ on neuron i at time t:

$$A_i^\mu(t) = \frac{1}{b(1-b)N} \sum_{p=1}^{P} \sum_{j=1}^{N} (\eta_i^{\mu p} - b)(\eta_j^{\mu p} - b) x_j(t - \tau_{ij}^\mu). \tag{3.1}$$

Since we are considering activity away from the mask points, if we assume that the different CRN masks and Hopfield patterns are all uncorrelated with one another, then it follows that all of the terms in the double sum are independent random variables. If we write $\Pr[x_j(t - \tau_{ij}^\mu) = 1] = f$, then, as $N \to \infty$, we have that $A_i^\mu(t)$ is a gaussian random variable with mean 0 and standard deviation αf, where $\alpha = P/N$ is the loading of the subnetwork.

This allows us to compute the probability that a conjunction detector becomes active. Since the conjunction detector will fire only if $A_i^\mu(t) \geq \theta_i^\mu$, we have

$$p_{i\mu}^1 = 1 - p_{i\mu}^0 = 1 - \Phi\left(\frac{\theta_i^\mu}{\sqrt{\alpha f}}\right), \tag{3.2}$$

where $\Phi(x)$ is given by

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-z^2/2} dz. \tag{3.3}$$

To get the probability that a neuron will fire, we note that it will fail to fire only if all of its conjunction detectors are silent. Using this fact and assuming that all conjunction detectors have the same threshold, θ, we can write the probability of generating spurious spikes as

$$f^{sp} = 1 - \left[\Phi\left(\frac{\theta}{\sqrt{\alpha f}}\right)\right]^M. \tag{3.4}$$

3.1.2 Mask Point Failures. We can take a similar approach for the activity at the mask points. However, it is now not always the case that $x_j(t - \tau_{ij}^\mu)$ is uncorrelated with the Hopfield pattern, $\eta_j^{\mu p}$. Without loss of generality, we consider the activity at a mask point belonging to subnetwork 1, and we begin by computing the activation of the first conjunction detector, $\mu = 1$.

If the subnetwork is perfectly recalling the first Hopfield pattern, then $x_j(t - \tau_{ij}^1) = \eta_j^{1,1}$, and we can write $A_i^1(t)$ as the sum of a signal, s, and noise, r, terms:

$$A_i^1(t) = s + r, \tag{3.5}$$

where

$$s = \frac{1}{b(1-b)N} \sum_{j=1}^{N} \left(\eta_i^{1,1} - b \right)\left(\eta_j^{1,1} - b \right)\eta_j^{1,1} = \eta_i^{1,1} - b \tag{3.6}$$

and

$$r = \frac{1}{b(1-b)N} \sum_{j=1}^{N} \sum_{p=2}^{P} \left(\eta_i^{1,p} - b \right)\left(\eta_j^{1,p} - b \right)\eta_j^{1,1}. \tag{3.7}$$

In analogy to the case of spurious spiking, as $N \to \infty$, r becomes a gaussian random variable with mean 0 and standard deviation αb.

From this, we can write the probability that the conjunction detector will fire given $\eta_i^{1,1}$, what the neuron should do if the subnetwork is correctly recalling the pattern. We write $p_{i\mu=1}^{I \to J}$ as the probability that conjunction detector 1 on neuron i will be in state J ($\in [0,1]$) given that it should be in state I:

$$
\begin{aligned}
p_{i\mu=1}^{0 \to 0} &= \Phi\left(\frac{\theta_i^1 + b}{\sqrt{\alpha b}} \right); \quad p_{i\mu=1}^{0 \to 1} = 1 - \Phi\left(\frac{\theta_i^1 + b}{\sqrt{\alpha b}} \right) \\
p_{i\mu=1}^{1 \to 0} &= \Phi\left(\frac{\theta_i^1 + b - 1}{\sqrt{\alpha b}} \right); \quad p_{i\mu=1}^{1 \to 1} = 1 - \Phi\left(\frac{\theta_i^1 + b - 1}{\sqrt{\alpha b}} \right)
\end{aligned}
\tag{3.8}
$$

Next, we note that the other conjunction detectors will behave in the same way as in the spurious spiking case (except for the rare events when two or more CRN masks overlap). Again, noting that a neuron fails to fire only when all of its component conjunction detectors are silent, we can write the probabilities for the neural output given the expected state of the neuron. If we assume homogeneous thresholds such that $\theta_i = \theta$ for all i, then we can write, $p^{I \to J}$ as the probability that a neuron is in state J given that it should be in state I. Therefore,

$$
\begin{aligned}
p^{0 \to 0} &= p_{i\mu=1}^{0 \to 0} \prod_{\mu=2}^{M} p_{i\mu}^0 = \Phi\left(\frac{\theta + b}{\sqrt{\alpha b}} \right)\left[\Phi\left(\frac{\theta}{\sqrt{\alpha f}} \right) \right]^{M-1} \\
p^{0 \to 1} &= 1 - p_i^{0 \to 0} = 1 - \Phi\left(\frac{\theta + b}{\sqrt{\alpha b}} \right)\left[\Phi\left(\frac{\theta}{\sqrt{\alpha f}} \right) \right]^{M-1} \\
p^{1 \to 0} &= p_{i\mu=1}^{1 \to 0} \prod_{\mu=2}^{M} p_{i\mu}^0 = \Phi\left(\frac{\theta + b - 1}{\sqrt{\alpha b}} \right)\left[\Phi\left(\frac{\theta}{\sqrt{\alpha f}} \right) \right]^{M-1} \\
p^{1 \to 1} &= 1 - p_i^{1 \to 0} = 1 - \Phi\left(\frac{\theta + b - 1}{\sqrt{\alpha b}} \right)\left[\Phi\left(\frac{\theta}{\sqrt{\alpha f}} \right) \right]^{M-1}.
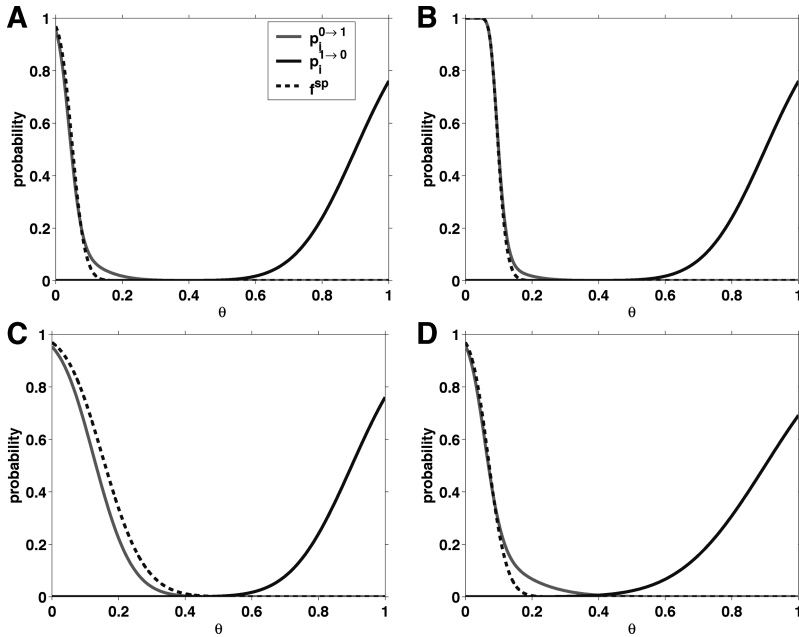\end{aligned}
\tag{3.9}
$$

Figure 2: Signal-to-noise analysis of the parallel Hopfield networks. In each panel, we plot the three different failure probabilities as a function of θ for different parameter settings. Unless otherwise stated, $b = 0.1$, $f = 0.01$, $M = 5$, and $\alpha = 0.2$. In all of the panels, $p^{0\to1}$ is denoted by the solid gray line, $p^{1\to0}$ by the solid black line, and f^{sp} by the dashed black line. (A) The default settings. (B) $M = 50$. (C) $f = 0.1$. (D) We increase α to 0.4, which is close to the maximum memory capacity of the Hopfield subnetworks.

3.1.3 Evaluation for Specific Parameter Settings. To get a handle on the relative sizes of these quantities, we plot the three failure probabilities (f^{sp}, $p^{0\to1}$ and $p^{1\to0}$) as a function of θ for four different parameter settings in Figure 2. In all of the panels, $p^{0\to1}$ is denoted by the solid gray line, $p^{1\to0}$ by the solid black line, and f^{sp} by the dashed black line. Unless otherwise stated, the default parameter settings are $b = 0.1$, $f = 0.01$, $M = 5$, and $\alpha = 0.2$.

In Figure 2A, we use the default settings. The first thing to notice is that there is a range of θ values over which all of the failure probabilities are approximately zero and where we might hope to find the parallel Hopfield behavior. Note also that $f^{sp} \approx p^{0\to1}$, which suggests that failure of the network due to $0 \to 1$ bit flipping is likely to occur at a similar point in parameter space as failure due to excess spurious spiking.

In Figure 2B, we increase M to 50 while keeping all other parameters the same. This leaves $p^{1\to0}$ approximately unchanged, while f^{sp} and $p^{0\to1}$ are

both increased. In Figure C, we set $f = 0.1$, which significantly reduces the range where the failure probability is close to zero. Finally in Figure 2D, we increase α to 0.4, which is close to the maximum memory capacity of the equivalent, isolated Hopfield network.

3.2 Self-Consistent Mean Field Equation for Spurious Spiking. We can gain further insight into the effects of spurious spiking by turning equation 3.4 into a self-consistent equation for f^{sp}. To do this we make three assumptions. First, we assume that the number of memories we are trying to recall is small, so that the mask points represent a fairly small proportion of the total output of the network. Next, we assume that f^{sp} changes slowly relative to the longest axon delays in the system, such that all inputs to each conjunction detector have the same probability of firing spuriously. Finally, we assume that all of the spurious spiking activity is asynchronous and apolychronous; none of the subnetworks is activated by chance. This assumption is justified for low levels of spurious spiking as the choice of weight matrix in equation 2.3 ensures that the all-zero pattern is a stable state of the subnetworks (Buhmann et al., 1989).

Given these assumptions, we can write an expression for f as

$$f = f^{sp} + f^{gen} - f^{sp} f^{gen}, \tag{3.10}$$

which gives the self-consistent equation for f^{sp} as

$$f^{sp} = 1 - \left[\Phi \left(\frac{\theta}{\sqrt{\alpha(f^{sp} + f^{gen} - f^{sp} f^{gen})}} \right) \right]^M . \tag{3.11}$$

The stable fixed points of equation 3.11 are found easily by iteration. In Figure 3, we show cobweb diagrams illustrating the iterative solution to the equation. We can identify three different regimes depending on the parameter settings. In Figure 3A, $f^{sp} \approx 0$ is the only stable fixed point of the system. In Figure 3B, there are two stable fixed points at $f^{sp} \approx 0$ and $f^{sp} \approx 0.8$; which one the network falls into is determined by the initial conditions. For the parameter settings in Figure 3C, we expect the network to always fail by proliferation, where the probability of spurious spiking is high. Finally in Figure 3D, we show the equilibrium values of f^{sp} for different values of f^{gen}. Stable fixed points lie on the solid lines, while the dotted line represents the locus of unstable equilibria. It is clear from the diagram that the system undergoes two saddle-node bifurcations at $f^{gen} \approx 0.1$ and 0.15.

Below $f^{gen} \approx 0.1$, only the silent state, with low levels of spurious spiking, is stable. At $f^{gen} \approx 0.1$, the proliferation state also becomes stable, and the fixed point we converge to will depend on the initial conditions in the network. At $f^{gen} \approx 0.15$, the silent state destabilizes, and only the
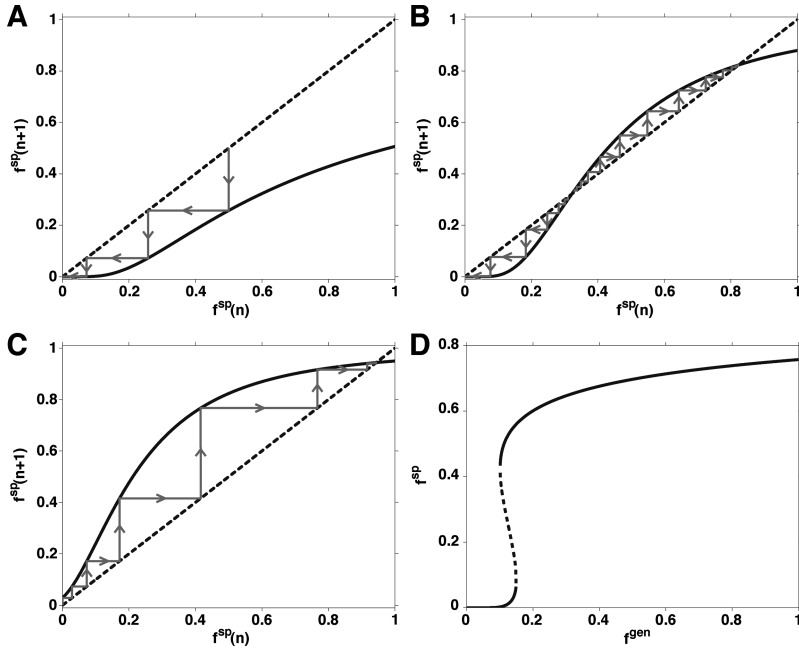
Figure 3: Self-consistent solutions for spurious spiking activity. In all panels, $\theta = 0.5$ and $b = 0.1$. (A) $\alpha = 0.2$, $M = 5$, $f^{gen} = 0.01$. For these parameter settings, the self-consistent equation has only one fixed point at $f^{sp} \approx 0$. (B) $\alpha = 0.2$, $M = 15$, $f^{gen} = 0.01$. Here there are two stable fixed points at high and low values of f^{sp} separated by an unstable fixed point that acts as a threshold to proliferation. (C) $\alpha = 0.3$, $M = 15$, $f^{gen} = 0.1$. In this case, the only stable state is proliferation. (D) We show the fixed points of equation 3.11 as a function of f^{gen} for $\alpha = 0.2$, $M = 10$. The solid lines represent stable fixed points, while the unstable fixed points lie on the dotted line. In this case, the network is stable with respect to proliferation up to about $f^{gen} = 0.15$.

proliferation state is stable. If the initial rate of spurious spiking is always zero, then this is the maximum level of genuine spiking that the network can tolerate before it will fail by proliferation, f^{gen}_{thresh}. Note that in general, f^{gen}_{thresh} is a function of θ, α, and M.

All that remains is to relate f^{gen} to the other parameters in the system. To do this, we note that the rate of genuine spiking associated with the activation of a single subnetwork μ is given by

$$f^{gen}_\mu = \frac{a_\mu}{T_\mu}, \tag{3.12}$$

where we have introduced the variable a_μ to denote the fraction of neurons in subnetwork μ that are firing.

If we assume that the spike times associated with different subnetworks are uncorrelated and that all subnetworks are in the same state such that $a_\mu = a$ for all μ, then when M_{recall} subnetworks are activated, we arrive at the following expression for f^{gen},

$$f^{gen} = 1 - \prod_{\mu=1}^{M_{recall}} \left(1 - f_\mu^{gen}\right) \tag{3.13}$$

$$= 1 - \left(1 - \frac{a}{T_{eff}}\right)^{M_{recall}},$$

where we have introduced T_{eff} as some kind of average period given by

$$\frac{1}{T_{eff}} \approx \frac{1}{M_{recall}} \sum_\mu \frac{1}{T_\mu}. \tag{3.14}$$

The only unknown quantity, then, is a, the activity level in the subnetworks. All else being equal, this determines whether a network will proliferate; subnetwork states with higher values of a, such as the spin-glass states described in the next section, are more likely to lead to proliferation. Thus, given T_{eff}, the threshold for proliferation, f_{thresh}^{gen}, translates directly into a minimal condition for the activity levels in the subnetworks: $a < a_{thresh}$.

4 Simulations

In this section, we present the results of some numerical experiments that demonstrate the existence of parallel Hopfield behavior in these networks. All simulations were run on a G5 Mac and a Dell Precision 530 desktop running Linux. Typical run times varied between 10 minutes and more than 12 hours depending on the complexity of the network under consideration. Networks were set up according to the prescription of section 2. Each memory had its own distinct period to keep intersubnetwork correlations to a minimum. Initial conditions were set by choosing the set of memories and patterns to recall and then presenting one full period of each memory to the network. Where two CRN memories overlapped and were contradictory (e.g., memory 1 requiring a spike and memory 2 not), the spike was always assumed to "win."

In the experiments presented here, no initial noise was added to the system, although interference between different memories is effectively a source of noise for the subnetwork, and informal investigations with small amounts of noise lead to little noticeable difference in performance. In all of the experiments, we set the number of neurons, N, equal to 500 and the bias, b, is fixed at 0.1. (Demo code for implementing parallel

Hopfield networks in Matlab can be found on the PHN Web site at http://www.seas.upenn.edu/~rcwilson/parallel_hopfield/.)

4.1 Order Parameters. Throughout this letter, we find it useful to characterize the behavior of the subnetworks in terms of various order parameters. In particular, we find the overlap, m, and the activity, a, to be most useful.

We define the overlap of pattern p on subnetwork μ with the current activity in subnetwork μ, x_i^μ as

$$m_{\mu p} = \frac{1}{N(b(1-b))} \sum_i \left(\eta_i^{\mu p} - b \right) x_i^\mu. \tag{4.1}$$

Note that $m_{\mu p}$ takes the value 1 when recall is perfect and zero when the recalled pattern is uncorrelated with the input pattern. For simplicity, we report only the overlap with the input pattern for each subnetwork and hence, drop the superscript p.

The activity, a_μ, of each subnetwork was encountered in the previous section and is defined formally as

$$a_\mu = \frac{1}{N} \sum_i x_i^\mu \tag{4.2}$$

and will take the value b for perfect recall.

4.2 Examples of Behavior Types. If the subnetworks are truly acting as independent Hopfield networks, then we can expect the network to exhibit at least four types of behavior. Three of these are characteristic of the isolated Hopfield networks: *extinction*, where activity in the network goes to zero; an *associative memory* regime, where each subnetwork almost perfectly recalls one of the stored patterns; and a *spin-glass state*, where each subnetwork falls into a stable state that is uncorrelated with any of the stored patterns. The fourth behavior type, *proliferation*, is not seen in the isolated networks.

In Figure 4 we present raster plots of the network output in each of the four regimes. In Figure 4A, we set $M = 5$, $T_{eff} = 50$, $\theta = 0.7$, and $\alpha = 0.4$. Running the network with these parameters produces a good example of extinction behavior, where the network is unable to sustain prolonged activity. In Figure 4B, we have $M = 5$, $T_{eff} = 50$, $\theta = 0.5$, and $\alpha = 0.05$. Such settings are well within the associative memory regime, and the network produces sustained activation without proliferating. In Figure 4C, we move into the spin-glass regime by increasing α to 0.5. Although we have not yet extracted the subnetwork activity, it is clear that in this case, the level of sustained activity is higher than in the associative memory case. Finally,
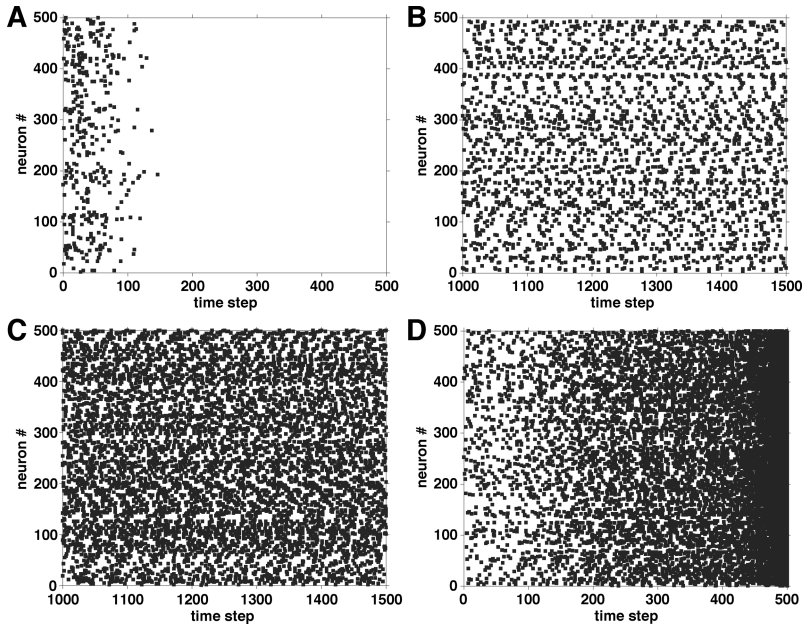
Figure 4: Experimental results demonstrating the four behavior types. In all cases, we have $b = 0.1$, $T_{eff} = 50$, and $M = 5$. (A) $\theta = 0.7$ and $\alpha = 0.4$ produce extinction behavior with the network unable to sustain activity for very long. (B) $\theta = 0.5$, $\alpha = 0.05$ produce an associative memory network, and the network is able to maintain a fixed level of activity for a long period of time. (C) $\theta = 0.5$, $\alpha = 0.5$ give rise to a spin-glass network. The level of activity is constant but higher than in the associative case. (D) $\theta = 0.5$, $\alpha = 0.65$ lead to proliferation where the activity of the network blows up.

in Figure 4D, we increase α to 0.65 to demonstrate proliferation behavior, where the activity of the network "explodes."

In Figure 5 we focus on the subnetwork activity in the associative memory and spin-glass examples. In Figure 5A, we show the activity extracted at one set of mask points in the associative memory case. Clearly, this subnetwork is able to maintain the same pattern of activity over a large number of time steps. In Figure 5B, we show the values of the overlap, m, and mean activity level, a, order parameters computed from the subnetwork activity for all five subnetworks. In this associative memory network, the subnetworks maintain high levels (≈ 1) of overlap with the pattern we are trying to recall and $a \approx b$, as we expect for perfect recall.

In Figure 5C, we show the activity extracted at one set of mask points in the spin-glass regime. This clearly has higher levels of activity, and when we compute the order parameters (shown in Figure 5D), a is indeed higher. As
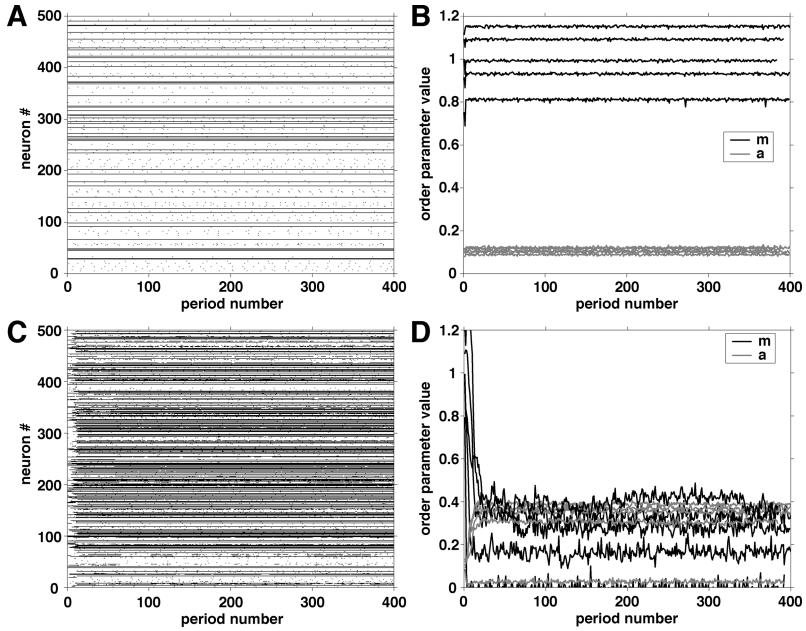
Figure 5: A closer look at the associative memory and spin-glass examples. (A) The activity at one set of CRN mask points for the network shown in Figure 4B. This subnetwork clearly sustains the same pattern over a prolonged period of time, and the order parameters in this case (B) clearly show that all of the subnetworks maintain macroscopic overlap, m, with the input pattern. (C) Activity at one set of CRN mask points for the network shown in Figure 4C. This subnetwork exhibits stable spin-glass ordering characterized by higher levels of activity and low levels of overlap (D) with any of the stored patterns.

expected for this regime, none of the stored patterns has significant overlap with the induced network state.

There is one more point to note from Figure 5D: one of the subnetworks in this simulation falls into the all-zero state—hence, the lines at m and $a \approx 0$. This spontaneous extinction of subnetwork activity is not at all described by our simple theory and is one of the behaviors that makes a full analysis so difficult.

4.3 Comparison with Isolated Networks. To gain insight into how well subnetworks are approximating the activity of the isolated Hopfield networks, we performed a series of simulations and compared the extracted order parameter values with those computed for the equivalent, isolated Hopfield network using the mean field equations developed in Buhmann et al. (1989). The results of some of these are shown in Figure 6.
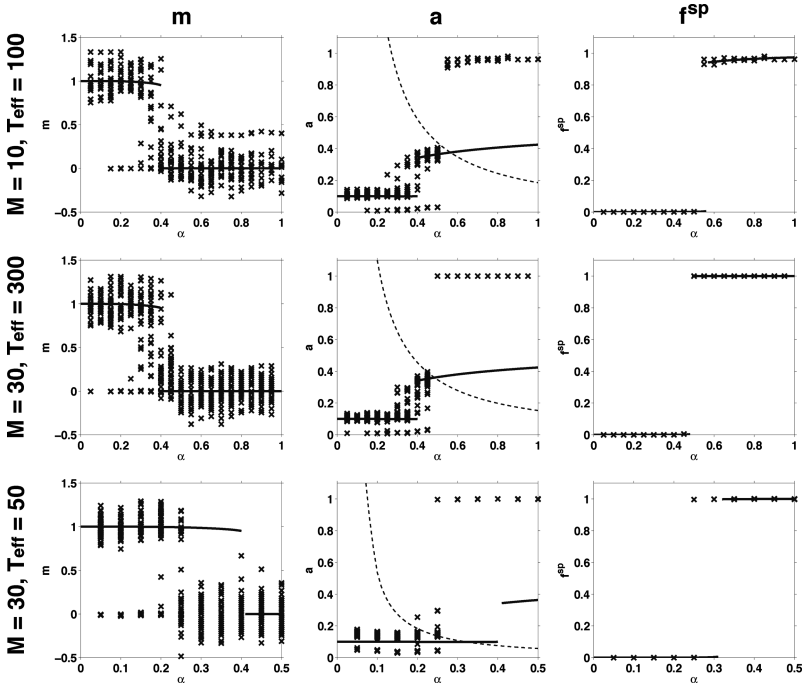
Figure 6: Quantitative comparison with the theory. Values of different order parameters (m, a, and f^{sp}) as a function of α. The results of the theory are represented by the black line. For m and a, these are the mean field values of the equivalent isolated Hopfield networks, while for f^{sp}, they are the fixed-point solutions to equation 3.11. The black dashed lines in the plots of a are the computed values of a_{thresh} as a function of α. The black crosses are the experimental results. Each column corresponds to a different order parameter—m on the left, a in the middle, and f^{sp} on the right—and different rows correspond to different parameter settings. In all cases, we have $\theta = 0.5$, $b = 0.1$, and the number of neurons is 500. Values of M and T_{eff} are given on the left of each row.

The results are arranged on a grid, with different parameter settings for each row and a different order parameter in each column. In the top row, we have, $M = 10$, $T_{eff} = 100$, and $\theta = 0.5$; in the middle, we increase M to 30 and T_{eff} to 300; and in the bottom row, we keep $M = 30$, but reduce T_{eff} to 50.

Each plot shows the value of an order parameter (m, a, or f^{sp}) for different values of α. In the plots of m and a, the black lines denote the expected order parameter values of the isolated Hopfield subnetworks; for f^{sp}, the black line denotes the steady-state solution to equation 3.11, assuming that the initial value of f^{sp} is zero. The black dashed line in the plots of a is the computed value of a_{thresh} as a function of α. In all cases,

the black crosses represent the experimental results. To compute the order parameters a and m from the experimental data, we take their mean values over the previous 10 periods for each subnetwork. For f^{sp}, the mean is taken over $10T_{eff}$ time steps.

To present the results, rather than compute the mean value of each order parameter across all of the subnetworks and repeat experiments, we have chosen to plot the raw data, as this gives better insight into the failures of the subnetworks when one might settle into a state that is not expected. In the first two rows, for values of $\alpha < 0.5$, there is good quantitative agreement between theory and experiment, not only in the values of the order parameters but also in the positions of the boundaries between the different types of behavior. Above $\alpha = 0.5$, the subnetworks no longer behave like their isolated counterparts, as the PHN fails by proliferation. Thus a, in particular, is in strong disagreement with the theory for isolated subnetworks. However, in both cases, the simple mean field theory for f^{sp} and a_{thresh} correctly predicts the point at which proliferation occurs. Thus, for these cases, the very simple theory has genuine predictive power.

In the third row, the theory does not correctly predict the onset of proliferation. This is most likely due to fluctuations in a that arise due to finite size effects in the simulation and that could momentarily raise a above the threshold required for proliferation. Of course, finite size effects should be present in all of the simulations, but they are more apparent here for two reasons. First, we expect the fluctuations to be greater in magnitude since f^{gen} is six times higher in this case. Second, the difference in slopes at the intersection between the lines for a and a_{thresh} is much smaller in the third row than in the other two cases. Therefore, small changes in a lead to much larger changes in the point of intersection, magnifying the effects of the fluctuations.

5 Biological Plausibility

The main problem with PHNs, as presented here, is that they do not seem particularly biologically plausible. In this section, we identify some of the more glaring problems and suggest possible solutions. We emphasize, however, that the main contribution of this letter is to introduce an analytically tractable model that illustrates a possible function for precisely timed spikes in the brain: as a tool that allows computations to be multiplexed.

First, the network exhibits many of the same problems as the traditional Hopfield model: neurons with both excitatory and inhibitory character; ultrahigh connectivity, implied by the full connectivity for each memory; infinite precision in the weights; and connectivity that satisfies extended symmetry (Herz et al., 1991) for each memory. These are clearly problematic, but we conjecture that as with the Hopfield network, many of these constraints will turn out to be unnecessary for the subnetworks to act as

attractors. For example, Hopfield networks can be made to obey Dale's law (Eccles, 1964) by pruning away incompatible connections. Taken further, more extreme pruning (Derrida et al., 1987) could reduce the connectivity to more realistic levels and can also remove the constraint of symmetric connections. Other work (Sompolinsky, 1986) has shown that requiring finite precision in the weights also does not remove the attractor properties of the Hopfield network.

Problems more specific to the PHN include the necessity for periodic memories, the assumption of clocked dynamics, and most important, the question of how such behavior could be learned. In fact, it is easy to see that the necessity for periodic memories is not a necessity at all and that the parallel Hopfield dynamics could easily be implemented in a feedforward network. The only limitation in this case would be on the maximum number of iterations available to each Hopfield subnetwork.

The assumption of clocked dynamics is more problematic, as there is no guarantee that in continuous time, the spikes will remain precisely timed over long periods. Nevertheless, two examples suggest that this might be possible: first, for the case of synfire chains, it has been shown that the synchronous state is an attractor (Diesmann, Gewaltig, & Aertsen, 1999), although these networks do not have conjunction detector–like elements. Second, simulations of the original CRN network (Wills, 2004) have shown that stable patterns of precisely timed spikes can propagate through these networks in continuous time.

Finally, we consider the need for a local and biologically plausible learning rule. We suppose that the connectivity of the network is sufficiently dense that there is a good chance of finding a neuron with a set of inputs close together on its dendrite firing in synchrony for an arbitrary pattern. This clearly will not be the case if we require full connectivity for each subnetwork, but for diluted subnetworks, this might be possible.

When the pattern is presented in the learning phase, synchronous input to this neuron would cause a dendritic spike. If the activity of the neuron is clamped so that it is independent of the dendritic spike during this learning phase, then the dendritic spike can act as a labeling event—tagging all of the synapses close to the site of the dendritic spike initiation point as being ready for learning. Note that this will also include synapses that did not fire to cause the dendritic spike.

The presence or absence of a neuronal spike can then be signaled by a backpropagating action potential (Stuart, Spruston, Sakmann, & Hausser, 1997), which means that all of the information required to learn the PHN (tagging of the conjunction detector and the correlation of synaptic activity with output of neuron) is present locally at the synapse. All that is then required is for a tagged synapse to alter its strength in the usual Hebbian fashion.

Clearly, this learning rule pushes the bounds of known biology, but it is certainly possible given the current knowledge.

6 Conclusion

We have introduced a new type of neural network, termed the Parallel Hopfield Network (PHN), that is capable of simultaneously effecting the dynamics of multiple Hopfield networks in the same piece of neural hardware. The key to the parallel behavior is the presence of conjunction detectors on the neurons. These elements (which are simplified models of spiking dendrites) are effective in two ways: first, they reduce interference between the activities in different subnetworks, and second, they reduce the incidence of spurious spiking, and hence failure, by the proliferation failure mode.

Outside the proliferation regime, simulations on networks with 500 neurons show that the activity of the subnetworks closely approximates that of the equivalent isolated Hopfield networks, having the same values for the order parameters m and a, while the transition to the proliferation regime is fairly well predicted by a simple mean field theory.

Acknowledgments

References

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.

Buhmann, J., Divko, R., & Schulten, K. (1989). Associative memory with high information content. *Physical Review A*, *39*, 2689–2692.

Derrida, B., Gardner, E., & Zippelius, A. (1987). An exactly soluble asymmetric neural network model. *EuroPhys. Lett.*, *4*, 167–173.

Diesmann, M., Gewaltig, M.-O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, *402*, 529–533.

Eccles, J. C. (1964). *The physiology of synapses*. Berlin: Springer-Verlag.

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*, 680–683.

Gasparani, S., & Magee, J. C. (2006). State-dependent dendritic computation in hippocampal CA1 pyramidal neurons. *Journal of Neuroscience*, *26*(7), 2088–2100.

Herz, A. V. M., Li, Z., & van Hemmen, J. L. (1991). Statistical mechanics of temporal association in neural networks with transmission delays. *Physical Review Letters*, *66*(10), 1370–1373.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. In *PNAS*, *79*, 2554–2558.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Fester, D., et al. (2004). Synfire chains and cortical songs: Temporal modules of cortical activity. *Science*, *304*, 559–564.

Izhikevich, E. M. (2005). Polychronization: Computation with spikes. *Neural Computation*, *18*, 245–282.

Meister, M. (1996). Multineuron codes in retinal signaling. *PNAS*, *93*, 609–614.

Shmiel, T., Drori, R., Shmiel, O., Ben-Shaul, Y., Nadasdy, Z., Shemesh, M., et al. (2005). Neurons of the cerebral cortex exhibit precise interspike timing in correspondence to behavior. *PNAS*, *102*, 18655–18657.

Sompolinsky, H. (1986). Neural networks with non-linear synapses and static noise. *Physical Review A*, *34*, 2571–2574.

Stuart, G., Spruston, N., Sakmann, B., & Hausser, M. (1997). Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends in Neuroscience*, *20*, 125–131.

Tsodyks, M., & Feigelman, M. (1988). The enhanced storage capacity in neural networks with low activity level. *EuroPhys. Lett.*, *6*, 101–105.

Wills, S. (2004). *Computation with spiking neurons*. Unpublished doctoral dissertation, Cambridge University.

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Attractor dynamics in hippocampal representation of the local environment. *Science*, *308*, 873–876.