LETTER Communicated by Matthijs van der Meer

Neural Mechanism to Simulate a Scale-Invariant Future

Karthik H. Shankar
shankark@bu.edu
Inder Singh
inders@bu.edu
Marc W. Howard
marc777@bu.edu
Center for Memory and Brain, Initiative for the Physics and Mathematics of Neural Systems, Boston University, Boston, MA 02215, U.S.A.

Predicting the timing and order of future events is an essential feature of cognition in higher life forms. We propose a neural mechanism to nondestructively translate the current state of spatiotemporal memory into the future, so as to construct an ordered set of future predictions almost instantaneously. We hypothesize that within each cycle of hippocampal theta oscillations, the memory state is swept through a range of translations to yield an ordered set of future predictions through modulations in synaptic connections. Theoretically, we operationalize critical neurobiological findings from hippocampal physiology in terms of neural network equations representing spatiotemporal memory. Combined with constraints based on physical principles requiring scale invariance and coherence in translation across memory nodes, the proposition results in Weber-Fechner spacing for the representation of both past (memory) and future (prediction) timelines. We show that the phenomenon of phase precession of neurons in the hippocampus and ventral striatum correspond to the cognitive act of future prediction.

1 Introduction

The brain encodes external stimuli in real time and represents information about the current spatial location and temporal history of recent events as activity distributed over neural networks. It is often useful for us to make decisions based on nonlocal events by predicting events to occur at distant future and remote locations. Optimal prediction could indeed be the guiding principle underlying the functional organization of the brain (Friston, 2010; Bialek, 2012; Palmer, Marre, Berry & Bialek, 2015). In order to anticipate events that might occur in the future after a given time or at a given distance from the current location, the brain needs to simulate how the current state of memory will have changed after waiting for a given amount of time or after moving through a given amount of distance. In this

Neural Computation **28, 2594–2627** (2016)
doi:10.1162/NECO_a_00891

letter, we propose that the brain can swiftly and nondestructively perform space/time translation operations on the memory state so as to anticipate events to occur at various future moments or remote locations.

The rodent brain contains a rich and detailed representation of current spatial location and temporal history. Some neurons—place cells—in the hippocampus fire in circumscribed locations within an environment, referred to as their place fields. Early work excluded confounds based on visual (Save, Cressant, Thinus-Blanc, & Poucet, 1998) or olfactory cues (Muller & Kubie, 1987), suggesting that the activity of place cells is a consequence of some form of path integration mechanism guided by the animal's velocity. Similarly, some neurons in the hippocampus—time cells—fire during a circumscribed period of time within a delay interval (Pastalkova, Itskov, Amarasingham, & Buzsaki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011; Gill, Mizumori, & Smith, 2011; Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013; Eichenbaum, 2014). By analogy to place cells, a set of time cells represents the animal's current temporal position relative to past events. Some researchers have long hypothesized a deep connection between the hippocampal representations of place and time (Eichenbaum & Cohen, 2014; Eichenbaum, 2000).

Motivated by the spatial and temporal memory represented in the hippocampus, we hypothesize that the translation operation required to anticipate the events at a distant future engages this part of the brain (Hasselmo, Giocomo, & Zilli, 2007; Schacter, Addis, & Buckner, 2007). We hypothesize that theta oscillations, a well-characterized rhythm of 4 Hz to 8 Hz in the local field potential observed in the hippocampus, may be associated with a translation operation. In particular, we hypothesize that sequential translations of different magnitudes take place at different phases within a cycle of theta oscillation, such that a timeline of anticipated future events is swept out in a single cycle (see Figure 1a).

Theta oscillations are prominently observed during periods of navigation (Vanderwolf, 1969). Critically, there is a systematic relationship between the animal's position within a neuron's place field and the phase of the theta oscillation at which that neuron fires (O'Keefe & Recce, 1993), known as phase precession. This suggests that the phase of firing of the place cells conveys information about the anticipated future location of the animal. This provides a strong motivation for our hypothesis that the phase of theta oscillation would be linked to a translation operation.

This letter develops a computational mechanism for a translation operation of a spatial/temporal memory representation constructed from a two-layer neural network model (Shankar & Howard, 2012) and links it to theta oscillations by imposing certain constraints based on some neurophysiological observations and some physical principles we expect the brain to satisfy. Since the focus here is to understand the computational mechanism of a higher-level cognitive phenomenon, the imposed constraints and the

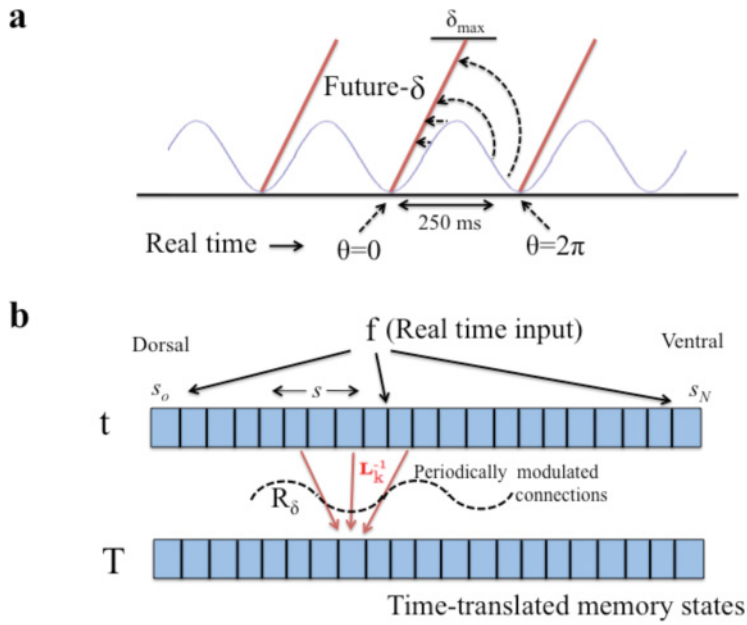

Figure 1: (a) Theta oscillations of 4–8 Hz are observed in the voltage recorded from the hippocampus. We hypothesize that within a theta cycle, a time line of future translations of magnitude δ is constructed, with δ sweeping from zero to some large value within each cycle. (b). Two-layer network with theta-modulated connections. The **t** layer receives external input **f** in real time and encodes its Laplace transform. The Laplace transform is inverted via a synaptic operator $\mathbf{L}_k^{-1}$ to yield an estimate of the function **f** on the **T** layer nodes. By periodically modulating the weights in $\mathbf{L}_k^{-1}$, denoted by the operator $\mathbf{R}_\delta$, the memory state represented in **T** layer is translated to represent its future states.

resulting derivation should be viewed at a phenomenological level, and not as emerging from biophysically detailed neural interactions.

We start with a memory representation constructed by a two-layer neural network (see Figure 1b) where the first layer (**t**) is a set of leaky integrators that encode the external stimuli in real time and the second layer (**T**) reconstructs the history of events through feedforward connections from the first layer. Time cells and place cells observed in the hippocampus can be directly derived from this simple network (Shankar & Howard, 2012; Howard et al., 2014). Further, this network mechanism intrinsically provides the feature of scale invariance that is widely observed in behavioral memory experiments. Hence, this network model is very adept at representing spatial and temporal memory.

At a mathematical level, the first layer encodes the Laplace transform of externally observed stimuli in real time, and the second layer approximately inverts the Laplace transform to represent a fuzzy estimate of the actual stimulus history (Shankar & Howard, 2012). This mathematical observation plays a crucial role in the ability of this memory network to simulate the future. The translation operation needed to simulate the future can be performed in the Laplace domain as an instantaneous point-wise product of the activity of nodes in the first layer with a scalar. We propose that the translation operation is achieved by modulating the connection weights between the two layers within each theta cycle (see Figure 1b). The translated representations can then be used to predict events at distant future and remote locations. In constructing the translation operation, we impose two physical principles we expect the brain to satisfy. The first principle is scale invariance, the requirement that all scales (temporal or spatial) represented in the memory are treated equally in implementing the translation. The second principle is coherence, the requirement that at any moment, all nodes forming the memory representation are in sync, translated by the same amount.

Further, to implement the computational mechanism of translation as a neural mechanism, we impose certain phenomenological constraints based on neurophysiological observations. First, along the dorsoventral axis of the hippocampus, the size of place fields increases systematically from the dorsal to the ventral end (Jung, Wiener, & McNaughton, 1994; Kjelstrup et al., 2008). In light of this observation we hypothesize that the nodes representing different temporal and spatial scales of memory are ordered along the dorsoventral axis. Second, the phase of theta oscillation is not uniform along the dorsoventral axis; the phase advances from the dorsal to the ventral end like a traveling wave (Lubenov & Siapas, 2009; Patel, Fujisawa, Berényi, Royer, & Buzsáki, 2012) with a phase difference of about π from one end to the other. Third, the synaptic weights change as a function of phase of the theta oscillation throughout the hippocampus (Wyble, Linster, & Hasselmo, 2000; Schall, Kerber, & Dickson, 2008). In light of this observation we hypothesize that the change in the connection strengths between the two layers required to implement the translation operation depends only on the local phase of the theta oscillation at any node (neuron).

In section 2 we impose these physical principles and phenomenological constraints to derive quantitative relationships for the distribution of scales of the nodes representing the memory and the theta-phase dependence of the translation operation. This yields specific forms of phase precession in the nodes representing the memory, as well as the nodes representing future prediction. Section 3 compares these forms to neurophysiological phase precession observed in the hippocampus and ventral striatum. Section 3 also makes explicit neurophysiological predictions that could verify our hypothesis that theta oscillations implement the translation operation to construct a time line of future predictions.

2 Mathematical Model

We start with a basic overview of the two-layer memory model and summarize the relevant details from previous work (Shankar & Howard, 2012, 2013; Howard et al., 2014) to serve as a background. Following that, we derive the equations that allow the memory nodes to be coherently time-translated to various future moments in synchrony with the theta oscillations. Finally we derive the predictions generated for various future moments from the time-translated memory states.

2.1 Theoretical Background. The memory model can be implemented as a two-layer feedforward network of nodes whose activities represent continuous firing rate (see Figure 1b). In this artificial neural network, the **t** layer holds the Laplace transform of the recent past, and the **T** layer reconstructs a temporally fuzzy estimate of past events (Shankar & Howard, 2012, 2013). Let the stimulus at any time τ be denoted as $\mathbf{f}(\tau)$. The nodes in the **t** layer are leaky integrators parameterized by their decay rate s and are all independently activated by the stimulus. The nodes are assumed to be arranged with regard to their s values. The nodes in the **T** layer are in one-to-one correspondence with the nodes in the **t** layer and hence can also be parameterized by the same s. The feedforward connections from the **t** layer into the **T** layer are prescribed to satisfy certain mathematical properties:

$$\frac{d}{d\tau}\mathbf{t}(\tau, s) = -s\mathbf{t}(\tau, s) + \mathbf{f}(\tau), \tag{2.1}$$

$$\mathbf{T}(\tau, s) = \mathbf{L}_{\mathrm{k}}^{-1}\,\mathbf{t}(\tau, s). \tag{2.2}$$

The activities $\mathbf{t}(\tau, s)$ and $\mathbf{T}(\tau, s)$, respectively, denote the continuous firing rate of the nodes in the **t** and **T** layers. For mathematical convenience, we have taken s to continuously parameterize the nodes of the two layers, while in reality the nodes will be parameterized by discrete values of s. However, it is clear from equations 2.1 and 2.2 that the two layers represent a continuous firing rate linear neural network, with the first layer having finite time constants of integration and the second layer having an infinite time constant—meaning the inputs from the first layer are instantaneously integrated to yield the activity of the second layer.

By integrating equation 2.1, note that the **t** layer encodes the Laplace transform of the entire past of the stimulus function leading up to the present. The s values distributed over the **t** layer represent the (real) Laplace domain variable. The fixed connection between the **t** layer and **T** layer denoted by the operator $\mathbf{L}_{\mathrm{k}}^{-1}$ (in equation 2.2), is constructed to approximate an inverse Laplace transformation. In effect, the Laplace-transformed stimulus history, which is distributed over the **t** layer nodes,

is inverted by $\mathbf{L}_k^{-1}$ such that a fuzzy (or coarsely grained) estimate of the actual stimulus value from various past moments is represented along the different **T** layer nodes. More precisely, treating the s values of the nodes as continuous, the $\mathbf{L}_k^{-1}$ operator can be succinctly expressed as

$$\mathbf{T}(\tau, s) = \frac{(-1)^k}{k!} s^{k+1} \mathbf{t}^{(k)}(\tau, s) \equiv \mathbf{L}_k^{-1}\, \mathbf{t}(\tau, s). \tag{2.3}$$

Here $\mathbf{t}^{(k)}(\tau, s)$ corresponds to the kth derivative of $\mathbf{t}(\tau, s)$ with regard to s. It can be proven that the $\mathbf{L}_k^{-1}$ operator executes an approximation to the inverse Laplace transformation, and the approximation becomes increasingly accurate for larger values of k (Post, 1930). Equation 2.3 prescribes the synaptic connections between the two layers. Each node in the **T** layer gets a feedforward input from the corresponding node in the **t** layer (with the same s) and k neighbors from either sides of that **t** layer node. The exact strength of each synapse between the two layers can be derived from a discrete approximation to the kth derivative (Shankar & Howard, 2013). For instance, if the s values of the nodes are uniformly spaced, with $k = 2$, the connections can be described as on-center, off-surround. However, we shall later see that it is optimal to choose the s values from a geometric progression that would require a different set of synaptic connections between the two layers. The appendix illustrates the synaptic connection for $k = 2$. However, for the purposes of this letter, we assume that the spacing between values of s is close enough that the analytic expression for $\mathbf{L}_k^{-1}$ given by equation 2.3 is accurate. Hence all derivations here are mathematically tractable and do not require us to numerically study the artificial neural network via simulations.

To emphasize the properties of this memory representation, consider the stimulus $\mathbf{f}(\tau)$ to be a Dirac delta function at $\tau = 0$. This input could reflect the presentation of an external stimulus at $\tau = 0$. From equations 2.1 and 2.3, the **T** layer activity following the stimulus presentation ($\tau > 0$) turns out to be

$$\mathbf{T}(\tau, s) = \frac{s}{k!} (s\tau)^k e^{-[s\tau]}. \tag{2.4}$$

Note that nodes with different s values in the **T** layer peak in activity after different delays following the stimulus; hence, the **T** layer nodes behave like time cells. In particular, a node with a given s peaks in activity at a time $\tau = k/s$ following the stimulus. Moreover, viewing the activity of any node as a distribution around its appropriate peak time (k/s), we see that the shape of this distribution is exactly the same for all nodes to the extent τ is rescaled to align the peaks of all the nodes. In other words, the activity of different nodes of the **T** layer represents a fuzzy estimate of the past information from different timescales, and the fuzziness associated

with them is directly proportional to the timescale they represent, while maintaining the exact same shape of fuzziness. For this reason, the **T** layer represents the past information in a scale-invariant fashion.

This two-layer memory architecture is also amenable to represent one-dimensional spatial memory analogous to the representation of temporal memory in the **T** layer (Howard et al., 2014). Rather than an external stimulus presented by the experimenter, if the stimulus **f** corresponds to contact with a fixed landmark encountered at a particular location in a one-dimensional spatial arena, then the **t** layer nodes can be made to represent the Laplace transform of the landmark treated as a function of position. By modifying equation 2.1 to

$$\frac{d}{d\tau}\mathbf{t}(\tau, s) = v\left[-s\mathbf{t}(\tau, s) + \mathbf{f}(\tau)\right], \tag{2.5}$$

where v is the velocity of motion, the temporal dependence of the **t** layer activity can be converted to spatial dependence. By employing the $\mathbf{L}_{\mathrm{k}}^{-1}$ operator on this modified **t** layer activity, it is straightforward to construct a layer of nodes (equation 2.2) at different distances from the landmark. Thus, the two-layer memory architecture can be trivially extended to yield place cells in one dimension.

In what follows, rather than referring to translation operations separately on spatial and temporal memory, we simply consider time translations with an implicit understanding that all the results derived can be trivially extended to 1D spatial memory representations.

2.2 Time-Translating the Memory State. The two-layer architecture naturally lends itself to time translations of the memory state in the **T** layer, which we later exploit to construct a time line of future predictions. The basic idea is that if the current state of memory represented in the **T** layer is used to anticipate the present (via some prediction mechanism), then a time-translated state of the **T** layer can be used to predict events that will occur at a distant future using the same prediction mechanism. Time translation essentially signifies the means to mimic the **T**-layer activity at a distant future based on its current state. This mechanism must be nondestructive in the sense that it should not overwrite the current activity in the **t** layer because the real-time external inputs are encoded in memory through the **t** layer.

Let δ be the amount by which we intend to time-translate the state of the **T** layer. So at any time τ, the aim is to access $\mathbf{T}(\tau + \delta, s)$ while still preserving the current **t**-layer activity, $\mathbf{t}(\tau, s)$. This can be easily achieved because the **t** layer represents the stimulus history in the Laplace domain. Noting that the Laplace transform of a δ-translated function is simply the product of

$e^{-s\delta}$ and the Laplace transform of the untranslated function, we see that

$$\mathbf{t}(\tau+\delta, s) = e^{-s\delta}\mathbf{t}(\tau, s). \tag{2.6}$$

Now noting that $\mathbf{T}(\tau+\delta, s)$ can be obtained by employing the $\mathbf{L}_k^{-1}$ operator on $\mathbf{t}(\tau+\delta, s)$ analogous to equation 2.3, we obtain the δ-translated $\mathbf{T}$ activity as

$$\begin{aligned}\mathbf{T}_\delta(\tau, s) \equiv \mathbf{T}(\tau+\delta, s) &= \mathbf{L}_k^{-1}\,\mathbf{t}(\tau+\delta, s)\\ &= \left[\mathbf{L}_k^{-1}\mathbf{R}_\delta\right]\,\mathbf{t}(\tau, s),\end{aligned} \tag{2.7}$$

where $\mathbf{R}_\delta$ is just a diagonal operator whose rows and columns are indexed by s and the diagonal entries are $e^{-s\delta}$. The δ-translated activity of the $\mathbf{T}$ layer is now subscripted by δ as $\mathbf{T}_\delta$ so as to distinguish it from the untranslated $\mathbf{T}$-layer activity given by equation 2.3. In this notation, the untranslated state $\mathbf{T}(\tau, s)$ from equation 2.3 can be expressed as $\mathbf{T}_0(\tau, s)$. In a continuous firing rate neural network implementation, the time-translated $\mathbf{T}$ activity can be obtained from the current $\mathbf{t}$-layer activity if the connection weights between the two layers given by $\mathbf{L}_k^{-1}$ are modulated by $\mathbf{R}_\delta$. Here we emphasize the analytic results that can be derived from this continuous firing rate network. This computational mechanism of time translation can be implemented as a neural mechanism in the brain by imposing certain phenomenological constraints and physical principles.

Observation 1. Anatomically, along the dorsoventral axis of the hippocampus, the width of place fields systematically increases from the dorsal end to the ventral end (Jung et al., 1994; Kjelstrup et al., 2008). Figure 2 schematically illustrates this observation by identifying the s-axis of the two-layer memory architecture with the dorsoventral axis of the hippocampus, such that the scales represented by the nodes are monotonically arranged. Let there be $N+1$ nodes with monotonically decreasing s values given by $s_o, s_1, ..., s_N$.

Observation 2. The phase of the theta oscillations along the axis is nonuniform, representing a traveling wave from the dorsal to the ventral part of the hippocampus with a net phase shift of π (Lubenov & Siapas, 2009; Patel et al., 2012). The oscillations in Figure 2 symbolize the local field potentials at different locations of the s-axis. The local phase of the oscillation at any position on the s-axis is denoted by θ_s, which ranges from $-\pi$ to $+\pi$ by convention. However, as a reference, we denote the phase at the dorsal end as θ_o, ranging from 0 to 2π, with the understanding that the range $(\pi, 2\pi)$ is mapped on to $(-\pi, 0)$. The x-axis in Figure 2 is time within a theta oscillation labeled by the phase θ_o.

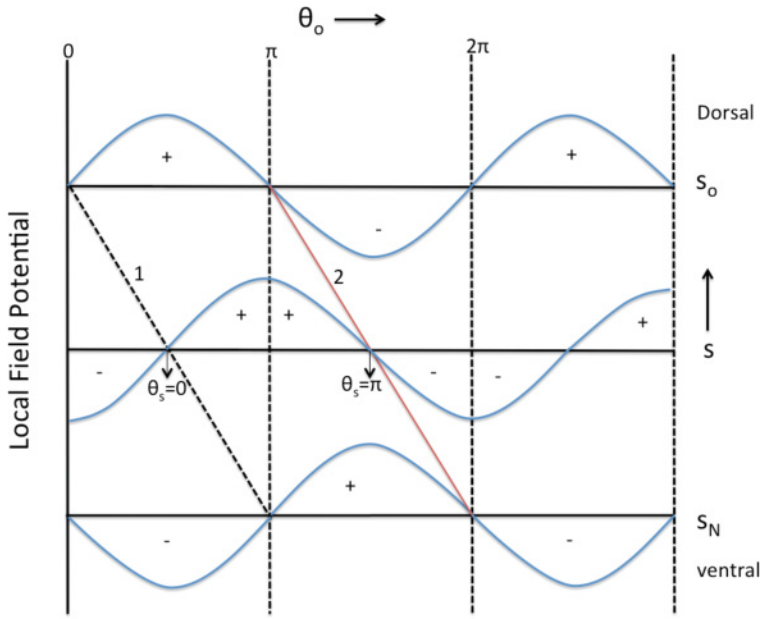

Figure 2: Traveling theta wave along the s-axis. The x-axis is real time. Each point along the dorsoventral axis corresponds to a different value of s_n. The blue curves show the theta oscillation for several different values of s. Lines 1 and 2 connect the positions where the local phases θ_s are 0 and π, respectively.

In this convention the value of θ_s discontinuously jumps from $+\pi$ to $-\pi$ as we move from one cycle of oscillation to the next. In Figure 2, the solid diagonal (red) line labeled 2 denotes all the points where this discontinuous jump happens. The diagonal dashed line labeled 1 denotes all the points where $\theta_s = 0$. It is straightforward to infer the relationship between the phase at any two values of s. Taking the nodes to be uniformly spaced anatomically, the local phase θ_s of the nth node is related to θ_o (for $0 < \theta_s < \pi$) by[1]

$$\theta_s/\pi = \theta_o/\pi - n/N. \tag{2.8}$$

Observation 3. Synaptic weights in the hippocampus are modulated periodically in synchrony with the phase of theta oscillation (Wyble et al., 2000; Schall et al., 2008). Based on this observation, we impose the constraint

[1]Since the s values of the nodes are monotonically arranged, we can interchangeably use s or n as subscripts to θ.

that the connection strengths between the **t** and **T** layers at a particular value of s depend only on the local phase of the theta oscillations. Thus the diagonal entries in the $\mathbf{R}_\delta$ operator should depend only on θ_s. We take these entries to be of the form $\exp[-\Phi_s(\theta_s)]$, where Φ_s is any continuous function of $\theta_s \in (-\pi, +\pi)$. Heuristically, at any moment within a theta cycle, a **T** node with a given s value will be roughly translated by an amount $\delta = \Phi_s(\theta_s)/s$.

Principle 1: *Preserve scale invariance*. Scale invariance is an extremely adaptive property for a memory to have; in many cases biological memories seem to exhibit scale invariance (Balsam & Gallistel, 2009). As the untranslated **T**-layer activity already exhibits scale invariance, we impose the constraint that the time-translated states of **T** should also exhibit scale invariance. This consideration requires the behavior of every node to follow the same pattern with respect to their local theta phase. This amounts to choosing the functions Φ_s to be the same for all s, which we refer to as Φ.

Principle 2: *Coherence in translation*. Since the time-translated memory state is going to be used to make predictions for various moments in the distant future, it would be preferable if all the nodes are time-translated by the same amount at any moment within a theta cycle. Otherwise, different nodes would contribute to predictions for different future moments, leading to noise in the prediction. However, such a requirement of global coherence cannot be imposed consistently along with principle 1 of preserving scale invariance.[2] But in the light of prior work (Hasselmo, Bodelón, & Wyble, 2002; Hasselmo, 2012), which suggests that retrieval of memory or prediction happens only in one-half of the theta cycle,[3] we impose the requirement of coherence only to those nodes that are all in the positive half of the cycle at any moment. That is $\delta = \Phi(\theta_s)/s$ is a constant along any vertical line in the region bounded between the diagonal lines 1 and 2 shown in Figure 2. Hence, for all nodes with $0 < \theta_s < \pi$, we require

$$\frac{\Phi\left(\theta_o - \pi n/N\right)}{s_n} = \text{constant}. \tag{2.9}$$

For coherence as expressed in equation 2.9 to hold at all values of θ_o between 0 and 2π, $\Phi(\theta_s)$ must be an exponential function so that θ_o can be functionally decoupled from n; consequently s_n should also have an

[2]This is easily seen by noting that each node will have a maximum translation inversely proportional to its s-value to satisfy principle 1.

[3]This hypothesis follows from the observation that while both synaptic transmission and synaptic plasticity are modulated by theta phase, they are out of phase with one another. That is, while certain synapses are learning, they are not communicating information, and vice versa. This led to the hypothesis that the phases where plasticity is optimal are specialized for encoding, whereas the phases where transmission is optimal are specialized for retrieval.

exponential dependence on n. So the general solution to equation 2.9 when $0 < \theta_s < \pi$ can be written as

$$\Phi(\theta_s) = \Phi_o e^{b\theta_s}, \tag{2.10}$$

$$s_n = s_o(1+c)^{-n}, \tag{2.11}$$

where c is a positive number. In this letter, we take $c \ll 1$, so that the analytic approximation for the $\mathbf{L}_k^{-1}$ operator given in terms of the kth derivative along the s axis in equation 2.3 is valid.

Thus, the requirement of coherence in time translation implies that the s values of the nodes—the timescales represented by the nodes—are spaced out exponentially, which we refer to as a Weber-Fechner scale because of its similarity to the behavioral Weber-Fechner law (Fechner, 1860/1912). Remarkably, this result strongly resonates with a requirement of the exact same scaling when the predictive information contained in the memory system is maximized in response to long-range correlated signals (Shankar & Howard, 2013). This feature allows this memory system to represent scale-invariantly coarse grained past information from timescales exponentially related to the number of nodes.

The maximum value attained by the function $\Phi(\theta_s)$ is at $\theta_s = \pi$, and the maximum value is $\Phi_{\max} = \Phi_o \exp(b\pi)$, such that $\Phi_{\max}/\Phi_o = s_o/s_N$ and $b = (1/\pi)\log(\Phi_{\max}/\Phi_o)$. To ensure continuity around $\theta_s = 0$, we take equation 2.10 to hold true even for $\theta_s \in (-\pi, 0)$. However, since notationally θ_s makes a jump from $+\pi$ to $-\pi$, $\Phi(\theta_s)$ would exhibit a discontinuity at the diagonal line 2 in Figure 2 from $\Phi_{\max}$ (corresponding to $\theta_s = \pi$) to $\Phi_{\min} = \Phi_o^2/\Phi_{\max}$ (corresponding to $\theta_s = -\pi$).

Given these considerations, at any instant within a theta cycle, referenced by the phase θ_o, the amount δ by which the memory state is time-translated can be derived from equations 2.8 and 2.10 as

$$\delta(\theta_o) = (\Phi_o/s_o)e^{b\theta_o}. \tag{2.12}$$

Analogous to having the past represented on a Weber-Fechner scale, the translation distance δ into the future also falls on a Weber-Fechner scale as the theta phase is swept from 0 to 2π. In other words, the amount of time spent within a theta cycle for larger translations is exponentially smaller.

To emphasize the properties of the time-translated $\mathbf{T}$ state, consider the stimulus to be a delta function at $\tau = 0$. From equation 2.7, we can express the $\mathbf{T}$-layer activity analogous to equation 2.4:

$$\mathbf{T}_\delta(\tau, s) \simeq \frac{s}{k!}\left[s\tau + \Phi\left(\theta_s\right)\right]^k e^{-\left[s\tau + \Phi(\theta_s)\right]}. \tag{2.13}$$

Notice that equations 2.8 and 2.12 specify a unique relationship between δ and θ_s for any given s. The right-hand side above is expressed in terms of θ_s rather than δ so as to shed light on the phenomenon of phase precession.

Since $\mathbf{T}_\delta(\tau, s)$ depends on both τ and θ_s only via the sum $\left[s\tau + \Phi\left(\theta_s\right)\right]$, a given node will show identical activity for various combinations of τ and θ_s.[4] For instance, a node would achieve its peak activity when τ is significantly smaller than its timescale (k/s) only when $\Phi(\theta_s)$ is large—meaning $\theta_s \simeq +\pi$. And as τ increases toward the timescale of the node, the peak activity gradually shifts to earlier phases all the way to $\theta_s \simeq -\pi$. An important consequence of imposing principle 1 is that the relationship between θ_s and τ on any iso-activity contour is scale invariant. That is, every node behaves similarly when τ is rescaled by the timescale of the node. We pursue the analogy of this phenomenon of phase precession further with neurophysiological findings in the next section (see Figure 4).

2.3 Timeline of Future Prediction. At any moment, $\mathbf{T}_\delta$ (see equation 2.13) can be used to predict the stimuli expected at a future moment. Consequently, as δ is swept through within a theta cycle, a timeline of future predictions can be simulated in an orderly fashion, such that predictions for closer events occur at earlier phases (smaller θ_o) and predictions of distant events occur at later phases. In order to predict from a time-translated state $\mathbf{T}_\delta$, we need a prediction mechanism. Here we consider the simple Hebbian mechanism for learning and prediction, wherein an event is learned (or an association formed in long-term memory) by increasing the connection strengths between the neurons representing the currently experienced stimulus and the neurons representing the recent past events ($\mathbf{T}_0$). Neurally, this simply amounts to a long-term potentiation between the **T**-layer neurons and those neurons instantly activated by the external stimulus. Because the **T**-layer activity contains temporal information about the preceding stimuli, simple associations between $\mathbf{T}_0$ and the current stimulus are sufficient to encode and express well-timed predictions (Shankar & Howard, 2012).

In general, the input **f** would be vector-valued in order to represent the different stimuli that could be presented, and each component would feed into its corresponding **t** and **T** layers. A layer of nodes coding for the future prediction **p** of various stimuli would be postsynaptically connected to the **T** layer to facilitate the Hebbian associations. During learning, the prediction nodes will inherit the input signal so that presynaptic activity in the **T** layer multiplied by the activity in the prediction neurons will contribute

[4] While representing timescales much larger than the period of a theta cycle, τ can essentially be treated as a constant within a single cycle. In other words, θ_s and τ in equation 2.7 can be treated as independent, although in reality, the phases evolve in real time.

to the increment in connection strength. Thus, the prediction nodes learn associations between a representation of the history of prior stimulus presentations in **T** and the current input. Under these assumptions, the output of a continuous firing rate model for the prediction of any stimulus is the scalar product of the learned states of the **T** layer with the current state of the **T** layer. And the scalar product of learned connection strengths and the time-translated memory state $\mathbf{T}_\delta$ can be understood as the future prediction of a stimulus at a time δ in the future.

To make this more concrete, consider a thought experiment where a conditioned stimulus CS is consistently followed by another stimulus, A or B, after a time τ_o. Later when CS is repeated (at a time $\tau = 0$), the subsequent activity in the **T** nodes can be used to generate predictions for the future occurrence of A or B. During learning, the connections to the node corresponding to A in **p** will be incremented by the state of $\mathbf{T}_0$ when A is presented; the connections to the node corresponding to B will be incremented by the state of $\mathbf{T}_0$ when B is presented.[5] After training, the prediction for the stimulus at a future time δ will depend on the time τ following the presentation of CS and the training delay τ_o. With Hebbian learning the prediction will be proportional to the sum of $\mathbf{T}_\delta$ activity of each node multiplied by the learned state $\mathbf{T}_0$:

$$\mathbf{p}_\delta(\tau, \tau_o) = \sum_{n=\ell}^{N} \mathbf{T}_\delta\left(\tau, s_n\right)\ \mathbf{T}_0\left(\tau_o, s_n\right)/s_n^w. \tag{2.14}$$

In order to preserve scale invariance while allowing different association strengths for different s nodes, we let the learning rate vary across nodes by a factor s_n^w (for any w). Since δ and θ_o are monotonically related (see equation 2.12), the prediction $\mathbf{p}_\delta$ for various future moments happens at various phases of a theta cycle.

Recall that all the nodes in the **T** layer are coherently time-translated only in the positive half of the theta cycle. Hence, for computing future predictions based on a time-translated state $\mathbf{T}_\delta$ only coherent nodes should contribute. In Figure 2 the region to the right of diagonal line 2 does not contribute to the prediction. The lower limit ℓ in the summation over the nodes given in equation 2.14 is the position of the diagonal line 2 in Figure 2 marking the position of discontinuity where θ_s jumps from $+\pi$ to $-\pi$. In the neural network implementation, this amounts to temporarily blocking the synaptic inputs into the prediction nodes from those **T**-layer nodes for which the local theta phase transitioned from $+\pi$ to $-\pi$; the synaptic inputs are unblocked at the beginning of the next cycle ($\theta_o = 0$). One could consider

[5] For simplicity, we assume that A and B are each presented the same number of times and set the learning rate to the inverse of that number.

a modification where only the set of **T** nodes in the positive phase of the theta cycle contribute to the prediction, in which case the upper limit in the summation given in equation 2.14 will be different from N; however the timeline of predictions generated will have the same qualitative properties (Shankar & Howard, 2015).

In the limit when $c \to 0$, the s values of neighboring nodes are very close, and the summation can be approximated by an integral. Defining $x = s\tau_o$ and $y = \tau/\tau_o$ and $v = \delta/\tau_o$, the above summation can be rewritten as

$$\mathbf{p}_\delta(\tau, \tau_o) \simeq \frac{\tau_o^{w-2}}{k!^2} \int_{x_{\min}}^{x_u} x^{2k+1-w}(y+v)^k e^{-x(1+y+v)}\, dx. \tag{2.15}$$

Here $x_{\min} = s_N\tau_o$, and $x_u = s_o\tau_o$ for $0 < \theta_o < \pi$ and $x_u = \Phi_{\max}\tau_o/\delta$ for $\pi < \theta_o < 2\pi$. The integral can be evaluated in terms of lower incomplete gamma functions to be

$$\mathbf{p}_\delta(\tau, \tau_o) \simeq \frac{\tau_o^{w-2}}{k!^2} \frac{\left[(\tau+\delta)/\tau_o\right]^k}{\left[1+(\tau+\delta)/\tau_o\right]^C} \times \left(\Gamma\left[C, (\tau_o+\tau+\delta)U\right] - \Gamma\left[C, (\tau_o+\tau+\delta)s_N\right]\right), \tag{2.16}$$

where $C = 2k + 2 - w$ and $\Gamma[.,.]$ is the lower incomplete gamma function. For $\theta_o < \pi$ (i.e., when $\delta < \Phi_{\max}/s_o$), $U = s_o$ and for $\theta_o > \pi$ (i.e., when $\delta > \Phi_{\max}/s_o$), $U = \Phi_{\max}/\delta$.

Figure 3 provides a graphical representation of some key properties of equation 2.16. The figure assumes that the CS is followed by A after $\tau_o = 3$ and followed by B after $\tau_o = 7$. The left panel shows the predictions for both A and B as a function of δ immediately after presentation of CS. The prediction for A appears at smaller δ and with a higher peak than the prediction for B. The value of w affects the relative sizes of the peaks. The right panel shows how the prediction for B changes with the passage of time after presentation of the CS. As τ increases from zero and the CS recedes into the past, the prediction of B peaks at smaller values of δ, corresponding to more imminent future times. In particular, when τ_o is much smaller than the largest (and larger than the smallest) timescale represented by the nodes the shape of $\mathbf{p}_\delta$ remains the same when δ and τ are rescaled by τ_o. Under these conditions, the time line of future predictions generated by $\mathbf{p}_\delta$ is scale invariant.

Because δ is in one-to-one relationship with θ_o, as a predicted stimulus becomes more imminent the activity corresponding to that predicted stimulus should peak at earlier and earlier phases. Hence a timeline of future predictions can be constructed from $\mathbf{p}_\delta$ as the phase θ_o is swept from 0 to 2π. Moreover, the cells representing $\mathbf{p}_\delta$ should show phase precession with respect to θ_o. Unlike cells representing $\mathbf{T}_\delta$, which depend directly on their local

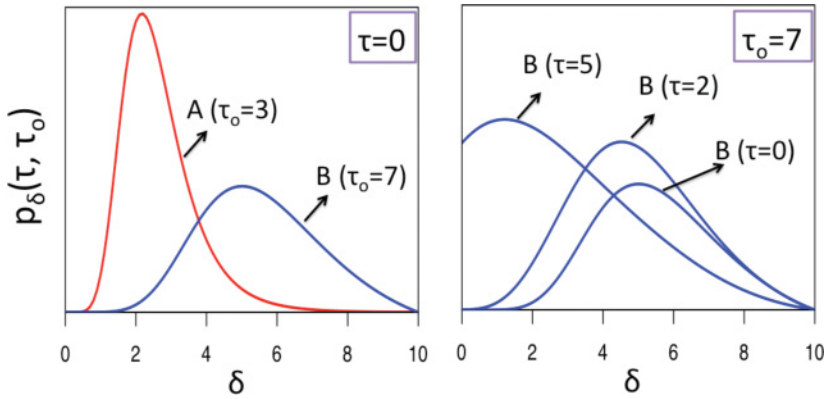

Figure 3: Future timeline. Equation 2.16 is plotted as a function of δ. During training, the CS was presented at $\tau_o = 3$ before A and $\tau_o = 7$ before B. (Left) Immediately after presentation of the CS, the predictions for A and B are ordered on the δ-axis. Note that the prediction for B approximates a rescaled version of that for A. (Right) The prediction for B is shown for varying times after the presentation of CS. With the passage of time, the prediction of B becomes stronger and more imminent. In this figure, $\Phi_{\max} = 10$, $\Phi_o = 1$, $k = 10$, $s_o = 10$, $s_N = 1$, and $w = 1$.

theta phase, θ_s, the phase precession of cells representing $\mathbf{p}_\delta$ should depend on the reference phase θ_o at the dorsal end of the s-axis. The next section elaborates on connections between this formalism and neurophysiological findings.

3 Comparisons with Neurophysiology

The mathematical development focused on two entities $\mathbf{T}_\delta$ and $\mathbf{p}_\delta$ that change their value based on the theta phase (see equations. 2.13 and 2.16). In order to compare these to neurophysiology, we need to have some hypothesis linking them to the activity of neurons from specific brain regions. We emphasize that although the development in the preceding section was done with respect to time, all of the results generalize to one-dimensional position as well (see equation 2.5). The overwhelming majority of evidence for phase precession comes from studies of place cells (but see Pastalkova et al., 2008). Here we compare the properties of $\mathbf{T}_\delta$ to phase precession in hippocampal neurons and the properties of $\mathbf{p}_\delta$ to a study showing phase precession in ventral striatum (van der Meer & Redish, 2011).[6]

[6]This is not meant to preclude the possibility that $\mathbf{p}_\delta$ could be computed at other parts of the brain as well.

Due to various analytic approximations, the activity of nodes in the **T** layer as well as the activity of the nodes representing future prediction (see equations 2.13 and 2.16) is expressed as smooth functions of time and theta phase. However, what we experimentally observe are discrete spikes. In order to facilitate a comparison of the model to neurophysiology, we adopt a simple stochastic spike-generating method. In this simplistic approach, the activity of the nodes given by equations 2.13 and 2.16 is taken to be proportional to the instantaneous probability for generating a spike. The probability of generating a spike at any instant is taken to be the instantaneous activity divided by the maximum activity achieved by the node if the activity is greater than 60% of the maximum activity. In addition, we add spontaneous stochastic spikes at any moment with a probability of 0.05. For all of the figures in this section, the parameters of the model are set as $k = 10$, $\Phi_{\max} = 10$, $w = 2$, $\Phi_o = 1$, $s_N = 1$, $s_o = 10$.

This relatively coarse level of realism in spike generation from the analytic expressions from a continuous firing rate neural network is probably appropriate to the resolution of the experimental data. There are some experimental challenges associated with exactly evaluating the model. First, the theta phase has to be estimated from a noisy signal. Second, phase precession results are typically shown as averaged across many trials. It is not necessarily the case that the average is representative of an individual trial (although this is the case at least for phase-precessing cells in medial entorhinal cortex; Reifenstein, Kempter, Schreiber, Stemmler, & Herz, 2012). Finally, the overwhelming majority of phase precession experiments use extracellular methods, which cannot perfectly identify spikes from individual neurons.

3.1 Hippocampal Phase Precession. It is clear from equation 2.13 that the activity of nodes in the **T** layer depends on both θ_s and τ. Figure 4 shows phase precession data from a representative cell (see Figure 4a; see also Mehta, Lee, & Wilson, 2002) and spikes generated from equation 2.13 (see Figure 4b). The model generates a characteristic curvature for phase precession, a consequence of the exponential form of the function Φ (see equation 2.10). The example cell chosen in Figure 4 shows roughly the same form of curvature as that generated by the model. While it should be noted that there is some variability across cells, careful analyses have led researchers to conclude that the canonical form of phase precession resembles this representative cell. For instance, a detailed study of hundreds of phase-precessing neurons (Yamaguchi, Aota, McNaughton, & Lipa, 2002) constructed averaged phase-precession plots using a variety of methods and found a distinct curvature that qualitatively resembles this neuron. Because of the analogy between time and one-dimensional position (see equation 2.5), the model predicts the same pattern of phase precession for time cells and place cells.

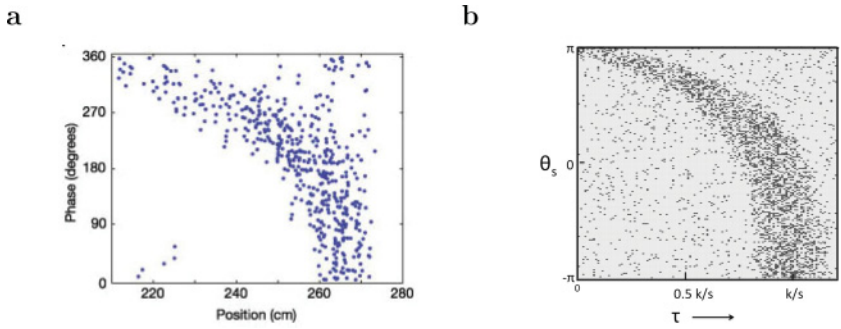

Figure 4: (a) Neurophysiological data showing phase precession. Each spike fired by a place cell is shown as a function of its position along a linear track (x-axis) and the phase of local theta (y-axis). After Mehta et al. (2002), (b) Simulated spikes from a node in the **T** layer described by equation 2.13 as a function of τ and local phase θ_s. The curvature is a consequence of equation 2.10. See text for details.

The **T**-layer activity represented in Figure 4a is scale invariant; note that the x-axis is expressed in units of the scale of the node (k/s). It is known that the spatial scale of place fields changes systematically along the dorsoventral axis of the hippocampus. Place cells in the dorsal hippocampus have place fields on the order of a few centimeters, whereas place cells at the ventral end have place fields as large as a few meters (see Figure 5a; Jung et al., 1994; Kjelstrup et al., 2008). However, all of them show the same pattern of precession with respect to their local theta phase: the phase measured at the same electrode that records a given place cell (see Figure 5). Recall that at any given moment, the local phase of theta oscillation depends on the position along the dorsoventral axis (Lubenov & Siapas, 2009; Patel et al., 2012), denoted as the s-axis in the model.

Figure 5a shows the activity of three different place cells in an experiment where rats ran down a long track that extended through open doors connecting three testing rooms (Kjelstrup et al., 2008). The landmarks controlling a particular place cell's firing may have been at a variety of locations along the track. Accordingly, Figure 5b shows the activity of cells generated from the model with different values of s and with landmarks at various locations along the track (described in the caption). From Figure 5, it can be qualitatively noted that phase precession of different cells depends on only the local theta phase and is unaffected by the spatial scale of firing. Further, studies (Royer, Sirota, Patel & Buzsáki, 2010) show that the hippocampal neurons show a ramping activity toward a landmark (associated with phase precession) irrespective of whether the journey was reward bound.

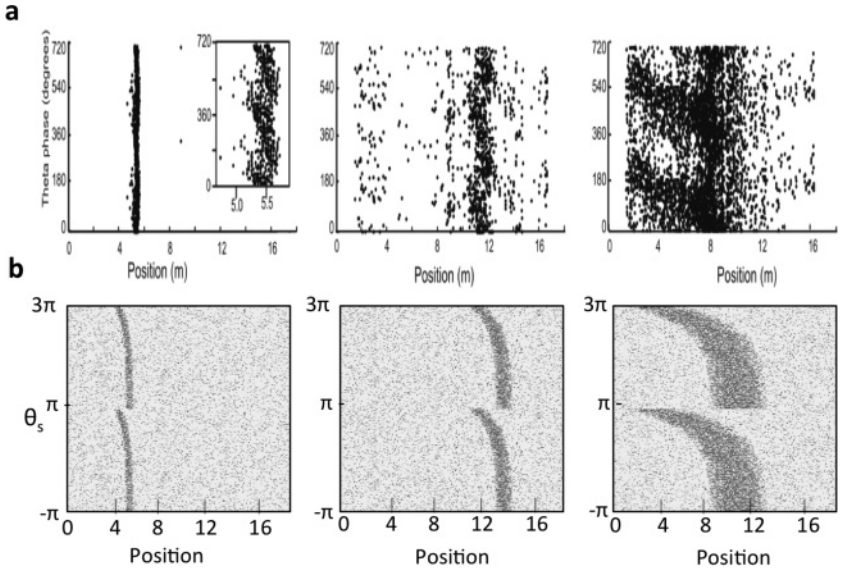

Figure 5: Place cells along the dorsoventral axis of the hippocampus have place fields that increase in size. (a) The three panels show the activity of place cells recorded at the dorsal, intermediate and ventral segments of the hippocampus, when a rat runs along an 18 m track. After (Kjelstrup, et al., 2008). Each spike the cell fired is shown as a function of position and the local theta phase at the cell's location when it fires (recall that theta phase is not constant across the dorsoventral axis). Regardless of the width of the place field, neurons at all locations along the dorsoventral axis phase precess through the same range of local theta phases. (b) According to the model, phase precession extends over the same range of values of local theta θ_s regardless of the value of s, which sets the scale for a particular node. As a consequence, cells with different values of s show time and place fields of different size but phase precess over the same range of local theta. The s values of the nodes are set to .1, .22, and .7, for the left, middle, and right panels, respectively, and they are assumed to respond to landmarks at location 4, 11, and 3 meters, respectively, from one end of the track.

3.2 Prediction of distant rewards via phase precession in the ventral striatum. We compare the future predictions generated by the model (see equation 2.16) to an experiment that recorded simultaneously from the hippocampus and nucleus accumbens, a reward-related structure within the ventral striatum (van der Meer & Redish, 2011). Here the rat's task was to learn to make several turns in sequence on a maze to reach two locations where a reward was available. Striatal neurons fired over long stretches of the maze, gradually ramping up their firing as a function of distance along the path and terminating at the reward locations (see Figure 6a, bottom).

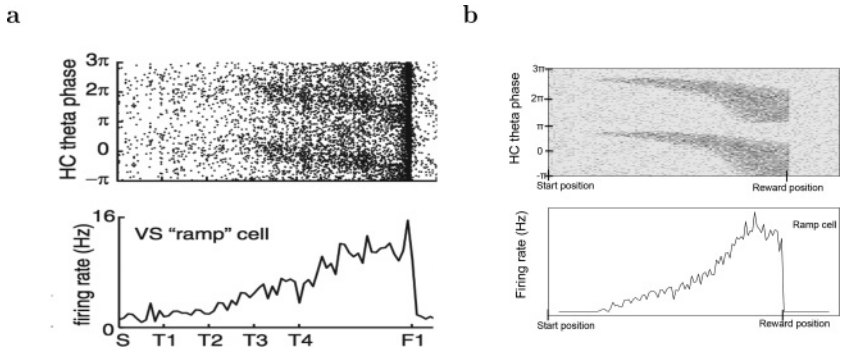

Figure 6: (a) A representative ramping cell in the ventral striatum. On each trial, the animal started the maze at S, made a series of turns (e.g., T1, T2), and received reward at F1 on 75% of trials. The total distance between S and F1 is on the order of a few meters. Position along the track is represented linearly on the x-axis for convenience. In the top panel, the spikes are shown as a function of theta phase at the dorsal hippocampus and position. The bottom panel shows the firing rate as a function of position, which is seen to gradually ramp up toward the reward location. (b) The activity of prediction node generated by the model is plotted with regard to the reference phase θ_o and position in the top panel, and the average activity within a theta cycle is plotted against position in the bottom panel.

Many striatal neurons showed robust phase precession relative to the theta phase at the dorsal hippocampus (see Figure 6a, top). Remarkably, the phase of oscillation in the hippocampus controlled firing in the ventral striatum to a greater extent than the phase recorded from within the ventral striatum. After the reward at the expected location (F1), there was another ramp up to a secondary reward location (F2), accompanied again by phase precession (not shown in Figure 6a).

This experiment corresponds reasonably well to the conditions assumed in the derivation of equation 2.16. In this analogy, the start of the trial (start location S) plays the role of the CS and the reward plays the role of the predicted stimulus. However, there is a discrepancy between the methods and the assumptions of the derivation. The ramping cell (see Figure 6a) abruptly terminates after the reward is consumed, whereas equation 2.16 would gradually decay back toward zero. This is because of the way the experiment was set up—there were never two rewards presented consecutively at the same location. As a consequence, having just received a reward at F1 strongly predicts that there will not be another reward at F1 in the next few moments.[7] In light of this consideration, we force the prediction

[7]We did not explicitly model the reward at F2 (not shown in Figure 6), but because the reward at F2 is always preceded by a reward at F1, one would expect that to be

generated in equation 2.16 to be zero beyond the reward location and let the firing be purely stochastic. The top panel of Figure 6b shows the spikes generated by model prediction cells with respect to the reference theta phase θ_o, and the bottom panel shows the ramping activity computed as the average firing activity within a complete theta cycle around any moment.

The model correctly captures the qualitative pattern observed in the data. According to the model, the reward starts being predicted at the beginning of the track. Initially, the reward is far in the future, corresponding to a large value of δ. As the animal approaches the location of the reward, the reward moves closer to the present along the δ-axis, reaching zero near the reward location. The ramping activity is a consequence of the exponential mapping between δ and θ_o in equation 2.10. Because the proportion of the theta cycle devoted to large values of δ is small, the firing rate averaged across all phases will be small, leading to an increase in activity closer to the reward.

3.3 Testable Properties of the Mathematical Model. Although the model aligns reasonably well with known properties of theta phase precession, a number of features of the model have, to our knowledge, not yet been evaluated. At a coarse level, the correspondence between time and one-dimensional space implies that time cells should exhibit phase precession with the same properties as place cells. While phase precession has been extensively observed and characterized in hippocampal place cells, there is much less evidence for phase precession in hippocampal time cells (but see Pastalkova et al., 2008). Recent evidence suggests that nonspatial stimuli are subject to hippocampal phase precession (Robinson, 2016), a strong requirement of this position (see also Lenck-Santini, Fenton, & Muller, 2008).

The mathematical assumption of scale invariance leads to several strong quantitative predictions at the cellular and ensemble levels. For instance, at the cellular level, plots of phase as a function of position for different values of s are rescaled versions of one another (see Figure 4b). In order to evaluate this prediction clearly, it is necessary to establish the origin of the coordinate system. For instance, in the Kjelstrup experiment (see Figure 5), there were many landmarks along the track. Unless we know which landmark (or landmarks) is controlling the firing of a particular cell, we cannot expect rescaling to hold. The experimental setup of Gothard and colleagues (Gothard, Skaggs, & McNaughton, 1996; Gothard, Hoffman, Battaglia, & McNaughton, 2001) provides a simple framework to study this question. By moving the start box on a linear track, they found that some place cells followed the location of the start box. This enables one to directly estimate the spatial scale of a particular cell and thus estimate s.

---

the dominant source of prediction and result in another rapid ramp between F1 and F2 accompanied by phase precession. This is consistent with the observations from this experiment.

At the ensemble level, scale invariance implies Weber law scaling of reconstructed position relative to a landmark. That is, if one attempts to reconstruct a position from firing at different values of theta phase from a population of neurons, one would expect that the error associated in the position goes up proportional to the distance from the landmark. Again, because this prediction requires a zero to anchor the coordinate system, it requires control of the population activity by a single landmark.

According to the model, the pattern of phase precession is related to the distribution of s values represented along the dorsoventral axis. While it is known that a range of spatial scales is observed along the dorsoventral axis, their actual distribution is not known. The Weber-Fechner scale of equation 2.10 is a strong prediction of the framework developed here. Moreover, since $\Phi_{\max}/\Phi_o = s_o/s_N$, the ratio of the largest to smallest scales represented in the hippocampus places constraints on the form of phase precession. The larger this ratio is, the larger will be the value of b in equation 2.10, and the curvature in the phase precession plots will emerge only at larger values of the local phase θ_s. Neurophysiological observation of this ratio could help evaluate the model.

The form of $\mathbf{p}_\delta$ (see equation 2.16) leads to several features in the pattern of phase precession of the nodes representing future prediction. It should be possible to observe phase precession for cells that are predicting any stimulus, not just a reward. In addition, the model's assumption that a time line of future predictions is aligned with global theta phase has measurable consequences. We reconsider the thought experiment from the previous section (see Figure 3), where a CS predicts an outcome after a delay τ_o. Immediately after the CS is presented, the value of δ at which the prediction peaks is monotonically related to τ_o. Because δ is monotonically related to the reference phase θ_o, the prediction cells will begin to fire at later phases when τ_o is large, and as time passes, they will fire at earlier and earlier phases all the way to $\theta_o = 0$. In other other words, the entry phase (at which the firing activity begins) should depend on τ_o, the prediction timescale. This is illustrated in Figure 7 with $\tau_o = 3$ and $\tau_o = 7$, superimposed on the same graph to make visual comparison easy. The magnitude of the peak activity would in general depend on the value of τ_o except when $w = 2$ (as assumed here for visual clarity). Experimentally manipulating the reward times and studying the phase precession of prediction cells could help test this feature. This qualitative prediction is shared with any other model that assigns future trajectories to specific theta phases; the quantitative predictions of this approach are shared with any model that assigns points along the future trajectory to theta phase with a Weber law compression.

4 Discussion

This letter has presented a neural hypothesis for implementing translations of temporal and 1D spatial memory states so that future events can be

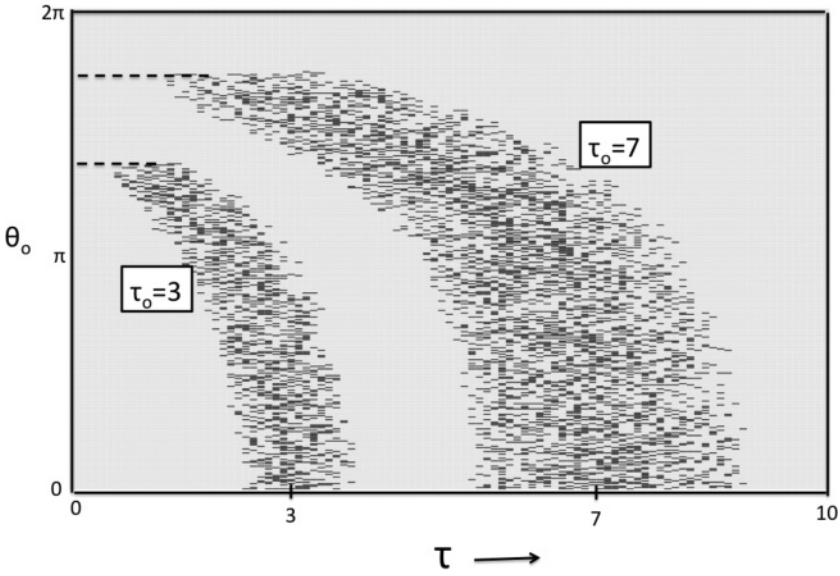

Figure 7: Changing τ_o affects the phase at which prediction cells start firing. At early times, the magnitude of translation required to predict the $\tau_o = 3$ outcome is smaller than that required to predict the $\tau_o = 7$ outcome. Consequently, the cell begins to fire at a larger θ_o for $\tau_o = 7$. Parameter values are the same as the other figures as given in the beginning of this section, except for clarity, the background probability of spiking has been set to zero.

quickly anticipated without destroying the current state of memory. The hypothesis assumes that time cells and place cells observed in the hippocampus represent time or position as a result of a two-layer architecture that encodes and inverts the Laplace transform of external input. It also assumes that sequential translations to progressively more distant points in the future occur within each cycle of theta oscillations. Neurophysiological constraints were imposed as phenomenological rules rather than as emerging from a detailed circuit model. Further, imposing scale invariance and coherence in translation across memory nodes resulted in Weber-Fechner spacing for the representation of both the past (spacing of s_n in the memory nodes) and the future (the relationship between δ and θ_o). Apart from providing cognitive flexibility in accessing a time line of future predictions at any moment, the computational mechanism described qualitative features of phase precession in the hippocampus and in the ventral striatum. Additionally, we have pointed out certain distinctive features of the model that can be tested with existing technology.

4.1 Mechanistic Implementation of Neural Network. Implementing the network in a biologically detailed way requires solving several problems. We briefly discuss the specifics involved in the implementation of the Laplace transform, equation 2.5, its inversion, equation 2.3, and the translation operation, equation 2.7.

4.1.1 Neural Implementation of the Laplace Transform. Equation 2.5 requires a set of integrators with rate constants s aligned on an anatomical axis, presumably the dorsoventral axis of the hippocampus. In order to implement Weber-Fechner scaling, each cell along the axis must set its rate constant to a constant ratio of its near neighbor (see equation 2.11). This suggests some sort of anatomical constraint in the early stages of brain development.

Further, there is also a question of how neurons could have long functional time constants that can be modulated by external control. Considerations from behavioral and neurophysiological experiments suggest time constants ($1/s$) should go up to the order of minutes (Howard, Shankar, Aue, & Criss, 2015; Mello, Soares, & Paton, 2015), much longer than the membrane time constants of actual neurons. Moreover, the modulation by velocity v requires that cells should be able to modulate their time constants rapidly in response to an external signal. It appears that these issues can be addressed in a biophysically feasible manner (Tiganj, Hasselmo, & Howard, 2015). It is known that neurons in regions around the medial temporal lobe, most notably entorhinal cortex and perirhinal cortex, can show stable persistent firing (Egorov, Hamam, Fransén, Hasselmo, & Alonso, 2002; Navaroli, Zhao, Boguszewski, & Brown, 2011), corresponding to an essentially infinite time constant. It is known that this property is a consequence of intracellular calcium ions, and its stability is due to buffering of calcium (Fransén, Tahvildari, Egorov, Hasselmo, & Alonso, 2006; Loewenstein & Sompolinsky, 2003). It can be computationally shown that a persistent firing model neuron will exhibit an exponentially decaying firing rate if the calcium buffering is eliminated (Tiganj et al., 2015). Also, rapid modulation of the rate constant can be implemented by any mechanism that rapidly changes the neuronal gain factor (Chance, Abbott, & Reyes, 2002; Poirazi, Brannon, & Mel, 2003; Silver, 2010).

4.1.2 Inverting the Laplace Transform with Neurons. To implement the inversion of Laplace transform as the connectivity $\mathbf{L}_k^{-1}$ between the $\mathbf{t}$ and $\mathbf{T}$ layers (see equation 2.3), it is essential to compute the weights corresponding to the discretized kth derivative of the $\mathbf{t}$ layer activity. For a given set of s values, the connection strengths can be explicitly evaluated as shown in Shankar and Howard (2013). The appendix elaborates on the connection strengths for $k = 2$. In a straightforward artificial neural network model, each node in the $\mathbf{T}$ layer will be connected to its corresponding node (with same s) in the $\mathbf{t}$ layer and to k nodes on either sides in the $\mathbf{t}$ layer. For larger values of k, perturbations in these connection strengths could significantly

change the computed value of the *k*th derivative. However, if the perturbations in the local connectivity from the **t**-layer nodes into each **T**-layer node along the dorsoventral axis are the same, then the the memory representation on the **T** layer will remain scale-invariant and the individual nodes will exhibit phase precession. However, the shape of the phase precession plot would change depending on the pattern of perturbations in the connectivity. This is because scale-invariant perturbations in the $\mathbf{L}_k^{-1}$ operator can be viewed as linear combinations of multiple $\mathbf{L}_k^{-1}$ operators (Shankar, 2015). Given that we cannot expect the perturbations in the connectivity to be exactly identical along the dorsoventral axis, it is reasonable to expect different cells to exhibit marginally different shapes of phase precession plots.

Although a straightforward artificial neural network implementation of the mathematical model specifies one set of feedforward connections, the same computation could be implemented by multiple different neural circuits. The most obvious neurophysiological solution would be to provide a set of Mexican hat receptive fields (which approximate a second derivative; Marr & Hildreth, 1980) in series. But other solutions are possible. For instance, inhibitory recurrent connections within a region (Ferster & Miller, 2000) could contribute to computation of the derivative. Given the complexity of dendritic computation (London & Häusser, 2005), it seems quite possible that higher-order derivatives could be computed by a dendritic arbor with the correct properties.

4.1.3 Neural Models of Phase Precession and Translation. Taken literally, the translation operator, equation 2.7, specifies that the synaptic weights in $\mathbf{L}_k^{-1}$ change as a function of theta phase. At a coarse level, this is not an unreasonable requirement: it is known that the size of observed postsynaptic EPSPs changes as a function of theta phase at many synapses in the hippocampal formation (Schall et al., 2008). On the other hand, the pattern of these changes is relatively complex, and we have not made a specific neurophysiological linking hypothesis to the $\mathbf{L}_k^{-1}$ connections.

Thus far, phase precession of place cell activity has been modeled by single-cell mechanisms as well as network mechanisms. The oscillator model originally proposed by O'Keefe and Recce (1993) posits that phase precession is a consequence of interaction of two types of oscillations in a place cell: a somatic oscillation synchronized with the local theta frequency and a dendritic oscillation whose frequency and amplitude increase (proportional to the velocity) when the animal is in the place field of that cell. The sum of the two oscillations yields a net membrane potential wavelet with a frequency higher than the theta frequency. The cell fires at the peaks of the membrane potential which synchronize with progressively earlier phases of the LFP as the animal moves from the beginning of the place field to the end. There are variants that model the interaction between the somatic

and dendritic oscillations differently (Kamondi, Acsady, Wang, & Buzsáki, 1998; Bose, Booth, & Recce, 2000; Magee, 2001; Lengyel, Szatmary, & Érdi, 2003) but still produce the phenomenon of phase precession. Intracellular recordings support the hypothesis that phase precession results from a spatially dependent dendritic input and a theta-frequency modulatory input on the soma (Kamondi et al., 1998; Harvey, Collman, Dombeck, & Tank, 2009). At a coarse level, it seems that it might be possible to identify the spatially dependent dendritic input with $\mathbf{L}_k^{-1}$ and the global modulatory signal at theta frequency with the somatic input. However, it is an open question whether this simple linking hypothesis is sufficient to lead to a detailed biophysical circuit model to implement equation 2.7.

In addition to single-cell models of phase precession, other authors have proposed network-level models. Since the observation that place cells with overlapping place fields fire at a constant phase difference within each individual theta cycle, Skaggs, McNaughton, Wilson, and Barnes (1996) proposed that place cells are connected in a network such that the activity spreads from place cells with place fields earlier along the path to place cells with place fields farther along the path, and these connections could be formed by the animals repeatedly traversing the path (Jensen & Lisman, 1996; Blum & Abbott, 1996; Mehta et al., 2002). Because these models depend on having the experience of a trajectory, they cannot account for translations along novel trajectories. Translation in the current framework does not depend on previous experience with that part of space, so this is an important point of distinction between network models and this model.

4.2 Translations without Theta Oscillations. Although this letter has focused on sequential translation within a theta cycle, translation may also be accomplished by other neurophysiological mechanisms. Sharp wave ripple (SWR) events last for about 100 ms and are often accompanied by replay events—sequential firing of place cells corresponding to locations different from the animal's current location (Davidson, Kloosterman, & Wilson, 2009; Dragoi & Tonegawa, 2011; Foster & Wilson, 2006; Pfeiffer & Foster, 2013; Jadhav, Kemere, German, & Frank, 2012). Notably, some experimentalists have also observed preplay events during SWRs, sequential activation of place cells that correspond to trajectories that have never been previously traversed, as though the animal is planning a future path (Dragoi & Tonegawa, 2011; Ólafsdóttir, Barry, Saleem, Hassabis & Spiers, 2015). Because untraversed trajectories could not have been used to learn and build sequential associations between the place cells along the trajectory, the preplay activity could potentially be a result of a translation operation on the overall spatial memory representation.

Sometimes during navigation, a place cell corresponding to a distant goal location gets activated (Pfeiffer & Foster, 2013), as though a finite distance translation of the memory state had occurred. More interesting, sometimes a

reverse replay is observed in which place cells are activated in reverse order, spreading back from the present location (Foster & Wilson, 2006). This is suggestive of translation into the past (as if δ were negative), to implement a memory search. In parallel there is behavioral evidence from humans that under some circumstances, memory retrieval consists of a backward scan through a temporal memory representation (Hacker, 1980; Hockley, 1984; Singh, Oliva, & Howard, 2016). Mathematically, as long as the appropriate connection strength changes prescribed by the $\mathbf{R}_\delta$ operator can be specified, there is no reason translations with negative δ or discontinuous shift in δ could not be accomplished in this framework. Whether these computational mechanisms are reasonable in light of the neurophysiology of sharp wave ripples is an open question.

4.3 Simulation of Unexperienced Trajectories. In contrast to associative models that require previous learning of trajectories, this approach can simulate paths that have not yet been traversed. This is possible because s forms something like a coordinate system. A particular cell's value of s carries with it meaningful metric relationships with the s values of other cells. This property, rather than previous experience with a particular stimulus, is what allows translation. There is, however, an important distinction to be made between translation and prediction. A state of $\mathbf{T}$ can be translated by δ to a state that has not been previously experienced. For example, immediately after experiencing a CS for the first time, $\mathbf{T}$ can be translated to a state with the CS 100 s in the past—a state that has never been experienced before. However, if the US has never been presented, this translated state of $\mathbf{T}$ does not enable prediction of the US in $\mathbf{p}_\delta$. The prediction of the US in $\mathbf{p}$ depends on learning an association between a state of $\mathbf{T}$ and the outcome US. Consider the limitations of this approach. If one has learned a sequence A B C, and then A is repeated, one can construct the translated states with A sequentially further in the past. This will cause B and then C to be predicted in $\mathbf{p}_\delta$. However, translation is not sufficient to cause information about B or C to appear in $\mathbf{T}$.

The hypothesis of a coordinate system enables, in principle at least, extension from the simple one-dimensional case to two-dimensional spatial environments, and perhaps even higher-dimensional semantic spaces. The coordinate system for, say, two dimensions, requires a set of s values arranged in a sheet (Howard et al., 2014). Specifying a particular movement in two dimensions corresponds to setting the v from equation 2.5 appropriately for each unit. Simulating a trajectory now corresponds to applying an appropriate δ for each unit. To take a simple concrete example, imagine a rat on a T-maze and four sets of $\mathbf{T}$ cells corresponding to the four cardinal directions. Suppose that a stimulus A is experienced at the stem of the T. As the animal moves toward the choice point, the set of cells in $\mathbf{T}$ corresponding to the North direction is sequentially activated, representing A sequentially farther from the present, while the state of the cells representing

the orthogonal directions is unaffected by that movement. After the choice point, the activated set of cells depends on the direction of movement. If the animal turns left, the west cells become sequentially activated; if instead the animal turns right, the east cells become sequentially activated. Simulating movements in 2D amount to setting the values of δ appropriately for the appropriate population of cells. For instance, starting from the choice point, setting δ to be nonzero for the west cells but zero for the east cells would correspond to simulation of a leftward trajectory. Indeed, hippocampal neurons have been observed to sweep out trajectories along one of two possible paths when an animal approaches a choice point on a T-maze (Johnson & Redish, 2007). While there are undoubtedly serious problems to be solved in extending the 1D translation framework described here to higher dimensions, simulation is possible because of the coordinate system established by the ordered relationships among the values of s.

4.4 Computational Advantages of the Laplace Domain. Other proposals for coding temporal memory have received attention (Shouval, Agarwal, & Gavornik, 2013; White, Lee, & Sompolinsky, 2004). However, these approaches do not yield sequentially activated time cells or place coding. The property of the **T** layer that different nodes represent the stimulus value from various delays (past moments) is reminiscent of a shift register (or delay line or synfire chain). However, the two-layer network encoding and inverting the Laplace transform of stimulus has several significant computational advantages over a shift register representation.

4.4.1 Independence of Timescales and Resource Conservation. In the two-layer network architecture, the spacing of s values of the nodes can be chosen freely. This property means that disruptions at one timescale need not affect all longer timescales. By choosing exponentially spaced s-values (Weber-Fechner scaling) as in equation 2.11, the **T** layer can represent memory from exponentially long timescales compared to a shift register with an equal number of nodes, thus making it extremely resource conserving. Although information from longer timescales is more coarsely grained, it turns out that this coarse graining is optimal to represent and predict long-range correlated signals (Shankar & Howard, 2013).

4.4.2 Scale Invariance. The memory representation of this two-layer network is naturally scale invariant (see equation 2.4). To construct a scale-invariant representation from a shift register, the shift register would have to be convolved with a scale-invariant coarse-graining function at each moment, which would be computationally very expensive. Moreover, it turns out that any network that can represent such scale-invariant memory can be identified with linear combinations of multiple such two-layer networks (Shankar, 2015).

4.4.3 Flexible Computation in the Laplace Domain. Because translation can be trivially performed when we have access to the Laplace domain, the two-layer network enables translations by an amount δ without sequentially visiting the intermediate states (with translation magnitudes less than δ). This can be done by directly changing the connection strengths locally between the two layers as prescribed by a diagonal $\mathbf{R}_\delta$ operator for any chosen δ.[8] Consequently the physical time taken for the translation can be decoupled from the magnitude of translation. One could imagine a shift register performing a translation operation by an amount δ either by shifting the values sequentially from one node to the next for δ time steps or by establishing nonlocal connections between far-away nodes. The latter would make the computation very cumbersome because it would require every node in the register to be connected to every other node (since this should work for any δ), which is in stark contrast with the local connectivity required by our two-layer network to perform any translation.

4.4.4 Neural Representation of Predictions. The computational function of $\mathbf{p}_\delta$ (see equation 2.16) is to represent an ordered set of events predicted to occur in the future. Although we focused on ventral striatum here because of the availability of phase precession data from that structure, it is probable that many brain regions represent future events as part of a circuit involving frontal cortex and basal ganglia, as well as the hippocampus and striatum (Schultz, Dayan, & Montague, 1997; Ferbinteanu & Shapiro, 2003; Tanaka et al., 2004; Feierstein, Quirk, Uchida, Sosulski, & Mainen, 2006; Mainen & Kepecs, 2009; Takahashi et al., 2011; Young & Shapiro, 2011). There is evidence that theta-like oscillations coordinate the activity in many of these brain regions (Jones & Wilson, 2005; Lansink, Goltstein, Lankelma, McNaughton, & Pennartz, 2009; van Wingerden, Vinck, Lankelma, & Pennartz, 2010; Fujisawa & Buzsáki, 2011). For instance, 4 Hz oscillations show phase coherence between the hippocampus, prefrontal cortex, and ventral tegmental area, a region that signals the presence of unexpected rewards (Fujisawa & Buzsáki, 2011). A great deal of experimental work has focused on the brain's response to future rewards, and indeed the phase-precessing cells in Figure 6 appear to be predicting the location of the future reward. The model suggests that $\mathbf{p}_\delta$ should predict any future event, not just a reward. Indeed, neurons that appear to code for predicted stimuli have been observed in the primate inferotemporal cortex (Sakai & Miyashita, 1991) and prefrontal cortex (Rainer, Rao, & Miller, 1999). Moreover, theta phase coherence between prefrontal cortex and hippocampus is essential for learning the temporal relationships between stimuli (Brincat & Miller,

[8] In this letter, we have considered sequential translations of various values of δ, since the aim was to construct an entire future time line rather than to discontinuously jump to a distant future state.

2015). Predictions of future events may be represented in a variety of brain regions.

Appendix: An Illustration of the Connection Strengths between the Two Layers

This appendix provides an explicit recipe for computing $\mathbf{L}_{\mathrm{k}}^{-1}$ for the distribution of s_n derived in the text with $k = 2$. A more general derivation can be found in Shankar and Howard (2013). When $s_n = s_o(1+c)^{-n}$, the connection strengths in the $\mathbf{L}_{\mathrm{k}}^{-1}$ operator take a special form: for every n, the local connectivity from the $(2k+1)$**t**-nodes to the nth **T**-node has an identical form multiplied by s_n. For example, with $k = 2$, the connection strengths to the nth **T** node from the **t** nodes in the local neighborhood are given by

$$\mathbf{t}_{n+2} \to s_n \frac{(c+1)^5}{c^2(c+2)^2},$$

$$\mathbf{t}_{n+1} \to s_n \frac{-(c+1)^2}{c},$$

$$\mathbf{t}_{n} \to s_n \frac{c^4+3c^3+c^2-4c-2}{c^2(c+2)^2},$$

$$\mathbf{t}_{n-1} \to s_n \frac{1}{c^2+c},$$

$$\mathbf{t}_{n-2} \to s_n \frac{1}{c^2(c+1)(c+2)^2}.$$

The factor s_n can be treated as a postsynaptic weight, and the rest (which only depend on c) can be treated as a presynaptic weight that is constant along the dorsoventral axis.

Acknowledgments

We gratefully acknowledge helpful discussions with Michael Hasselmo, John Lisman, Sam McKenzie, Ehren Newman, Jon Rueckemann, Shantanu Jadhav, Dan Bullock, Dave Redish, Matt van der Meer, and Adam Johnson. This work was supported by AFOSR award FA9550-12-1-0369, NSF PHY 1444389 and Boston University's Initiative for the Physics and Mathematics of Neural Systems.

References

Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neuroscience*, *32*(2), 73–78.

Bialek, W. (2012). *Biophysics: Searching for principles*. Princeton, NJ: Princeton University Press.

Blum, K. I., & Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, *8*(1), 85–93.

Bose, A., Booth, V., & Recce, M. (2000). A temporal mechanism for generating the phase precession of hippocampal place cells. *Journal of Computational Neuroscience*, *9*(1), 5–30.

Brincat, S. L., & Miller, E. K. (2015). Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nature Neuroscience*, *18*(4), 576–581. doi:10.1038/nn.3954

Chance, F. S., Abbott, L. F., & Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron*, *35*(4), 773–782.

Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, *63*(4), 497–507. doi:10.1016/j.neuron.2009.07.027

Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, *469*(7330), 397–401.

Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002). Graded persistent activity in entorhinal cortex neurons. *Nature*, *420*(6912), 173–178.

Eichenbaum, H. (2000). A cortical-hippocampal system for declarative memory. *Nature Reviews Neuroscience*, *1*(1), 41–50.

Eichenbaum, H. (2014). Time cells in the hippocampus: A new dimension for mapping memories *Nature Reviews Neuroscience*, *15*(11), 732–744. doi:10.1038/nrn3827

Eichenbaum, H., & Cohen, N. J. (2014). Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron*, *83*(4), 764–770.

Fechner, G. (1860/1912). *Elements of psychophysics*, vol. 1. Boston: Houghton Mifflin.

Feierstein, C. E., Quirk, M. C., Uchida, N., Sosulski, D. L., & Mainen, Z. F. (2006). Representation of spatial goals in rat orbitofrontal cortex. *Neuron*, *51*(4), 495–507.

Ferbinteanu, J., & Shapiro, M. L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron*, *40*(6), 1227–1239.

Ferster, D., & Miller, K. D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, *23*(1), 441–471.

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*(7084), 680–683.

Fransén, E., Tahvildari, B., Egorov, A. V., Hasselmo, M. E., & Alonso, A. A. (2006). Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. *Neuron*, *49*(5), 735–746.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*, 127–138.

Fujisawa, S., & Buzsáki, G. (2011). A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. *Neuron*, *72*(1), 153–165. doi:10.1016/j.neuron.2011.08.018

Gill, P. R., Mizumori, S. J. Y., & Smith, D. M. (2011). Hippocampal episode fields develop with learning. *Hippocampus*, *21*(11), 1240–1249. doi:10.1002/hipo.20832

Gothard, K. M., Hoffman, K. L., Battaglia, F. P., & McNaughton, B. L. (2001). Dentate gyrus and CA1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. *Journal of Neuroscience*, *21*(18), 7284–7292.

Gothard, K. M., Skaggs, W. E., & McNaughton, B. L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience*, *16*(24), 8027–8040.

Hacker, M. J. (1980). Speed and accuracy of recency judgments for events in short-term memory. *Journal of Experimental Psychology: Human Learning and Memory*, *15*, 846–858.

Harvey, C. D., Collman, F., Dombeck, D. A., & Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, *461*(7266), 941–946.

Hasselmo, M. E. (2012). *How we remember: Brain mechanisms of episodic memory*. Cambridge, MA: MIT Press.

Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, *14*, 793–817.

Hasselmo, M. E., Giocomo, L. M., & Zilli, E. A. (2007). Grid cell firing may arise from interference of theta frequency membrane potential oscillations in single neurons. *Hippocampus*, *17*(12), 1252–1271.

Hockley, W. E. (1984). Analysis of response time distributions in the study of cognitive processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(4), 598–615.

Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, *34*(13), 4692–4707. doi:10.1523/JNEUROSCI.5808-12.2014

Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, *122*(1), 24–53.

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, *336*(6087), 1454–1458. doi:10.1126/science.1217230

Jensen, O., & Lisman, J. E. (1996). Hippocampal CA3 region predicts memory sequences: Accounting for the phase precession of place cells. *Learning and Memory*, *3*, 279–287.

Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, *27*(45), 12176–12189.

Jones, M. W., & Wilson, M. A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol*, *3*(12), e402. doi:10.1371/journal.pbio.0030402

Jung, M. W., Wiener, S. I., & McNaughton, B. L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, *14*(12), 7347–7356.

Kamondi, A., Acsady, L., Wang, X. J., & Buzsáki, G. (1998). Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: Activity-dependent phase-precession of action potentials. *Hippocampus*, *8*, 244–261.

Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., ... Moser, M. B. (2008). Finite scale of spatial representation in the hippocampus. *Science*, *321*(5885), 140–143.

Kraus, B. J., Robinson, R. J., II, White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal "time cells": Time versus path integration. *Neuron*, *78*(6), 1090–1101. doi:10.1016/j.neuron.2013.04.015

Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L., & Pennartz, C. M. A. (2009). Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biology*, *7*(8), e1000173. doi:10.1371/journal.pbio.1000173

Lenck-Santini, P. P., Fenton, A. A., & Muller, R. U. (2008). Discharge properties of hippocampal neurons during performance of a jump avoidance task. *Journal of Neuroscience*, *28*(27), 6773–6786.

Lengyel, M., Szatmary, Z., & Érdi, P. (2003). Dynamically detuned oscillations account for the coupled rate and temporal code of place cell firing. *Hippocampus*, *13*(6), 700–714.

Loewenstein, Y., & Sompolinsky, H. (2003). Temporal integration by calcium dynamics in a model neuron. *Nature Neuroscience*, *6*(9), 961–967.

London, M., & Häusser, M. (2005). Dendritic computation. *Annual Review of Neuroscience*, *28*, 503–532.

Lubenov, E. V., & Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, *459*(7246), 534–539.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, *71*(4), 737–749.

Magee, J. C. (2001). Dendritic mechanisms of phase precession in hippocampal CA1 pyramidal neurons. *Journal of Neurophysiology*, *86*(1), 528–532.

Mainen, Z. F., & Kepecs, A. (2009). Neural representation of behavioral outcomes in the orbitofrontal cortex. *Curr. Opin. Neurobiol*, *19*(1), 84–91. doi:10.1016/j.conb.2009.03.010

Marr, D., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London B*, *207*(1167), 187–217.

Mehta, M. R., Lee, A. K., & Wilson, M. A. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature*, *417*(6890), 741–746.

Mello, G. B., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, *25*(9), 1113–1122.

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, *7*(7), 1951–1968.

Navaroli, V. L., Zhao, Y., Boguszewski, P., & Brown, T. H. (2011). Muscarinic receptor activation enables persistent firing in pyramidal neurons from superficial layers of dorsal perirhinal cortex. *Hippocampus*, *22*, 1391–1404. doi:10.1002/hipo.20975

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

Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D., & Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *eLife*, *4*, e06063.

Palmer, S. E., Marre, O., Berry, M. J., II, & Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences USA*, *112*(22), 6908–6913. doi:10.1073/pnas.1506855112

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*(5894), 1322–1327.

Patel, J., Fujisawa, S., Berényi, A., Royer, S., & Buzsáki, G. (2012). Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron*, *75*(3), 410–417. doi:10.1016/j.neuron.2012.07.015

Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, *497*(7447), 74–79. doi:10.1038/nature12112

Poirazi, P., Brannon, T., & Mel, B. W. (2003). Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron*, *37*(6), 977–987.

Post, E. (1930). Generalized differentiation. *Transactions of the American Mathematical Society*, *32*, 723–781.

Rainer, G., Rao, S. C., & Miller, E. K. (1999). Prospective coding for objects in primate prefrontal cortex. *Journal of Neuroscience*, *19*(13), 5493–5505.

Reifenstein, E. T., Kempter, R., Schreiber, S., Stemmler, M. B., & Herz, A. V. M. (2012). Grid cells in rat entorhinal cortex encode physical space with independent firing fields and phase precession at the single-trial level. *Proceedings of the National Academy of Sciences*, *109*(16), 6301–6306. doi:10.1073/pnas.1109599109

Robinson, N. T. M. (2016). *The role of medial entorhinal cortex activity in hippocampal CA1 spatiotemporally correlated sequence generation and object selectivity for memory function*. Unpublished doctoral dissertation, Boston University.

Royer, S., Sirota, A., Patel, J., & Buzsáki, G. (2010). Distinct representations and theta dynamics in dorsal and ventral hippocampus. *Journal of Neuroscience*, *30*(5), 1777–1787.

Sakai, K., & Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. *Nature*, *354*(6349), 152–155.

Save, E., Cressant, A., Thinus-Blanc, C., & Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. *Journal of Neuroscience*, *18*(5), 1818–1826.

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews, Neuroscience*, *8*(9), 657–661.

Schall, K. P., Kerber, J., & Dickson, C. T. (2008). Rhythmic constraints on hippocampal processing: State and phase-related fluctuations of synaptic excitability during theta and the slow oscillation. *Journal of Neurophysiology*, *99*(2), 888–899.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.

Shankar, K. H. (2015). Generic construction of scale-invariantly coarse grained memory. *Lecture Notes in Artificial intelligence*, *8955*, 175–184.

Shankar, K. H., & Howard, M. W. (2012). A scale-invariant internal representation of time. *Neural Computation*, *24*(1), 134–193.

Shankar, K. H., & Howard, M. W. (2013). Optimally fuzzy temporal memory. *Journal of Machine Learning Research*, *14*, 3753–3780.

Shankar, K. H., & Howard, M. W. (2015). Scale-free memory to swiftly generate fuzzy future predictions. In *Proceedings of the Fifth International Conference on Fuzzy and Neuro Computing* (pp. 185–194). New York: Springer.

Shouval, H. Z., Agarwal, A., & Gavornik, J. P. (2013). Scaling of perceptual errors can predict the shape of neural tuning curves. *Physical Review Letters*, *110*, 168102. http://link.aps.org/doi/10.1103/PhysRevLett.110.168102

Silver, R. A. (2010). Neuronal arithmetic. *Nature Reviews Neuroscience*, *11*(7), 474–489.

Singh, I., Oliva, A., & Howard, M. W. (2016). *Visual memories are stored along a logarithmically-compressed representation of the past*. Manuscript submitted for publication.

Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, *6*, 149–172.

Takahashi, Y. K., Roesch, M. R., Wilson, R. C., Toreson, K., O'Donnell, P., Niv, Y., & Schoenbaum, G. (2011). Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nature Neuroscience*, *14*(12), 1590–1597. doi:10.1038/nn.2957

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., & Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience*, *7*(8), 887–893. doi:10.1038/nn1279

Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2015). A simple biophysically plausible model for long time constants in single neurons. *Hippocampus*, *25*(1), 27–37.

van der Meer, M. A. A., & Redish, A. D. (2011). Theta phase precession in rat ventral striatum links place and reward information. *Journal of Neuroscience*, *31*(8), 2843–2854. doi:10.1523/JNEUROSCI.4869-10.2011

Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement of the rat. *Electroencephalography and Clinical Neurophysiology*, *26*, 407–418.

van Wingerden, M., Vinck, M., Lankelma, J., & Pennartz, C. M. (2010). Theta-band phase locking of orbitofrontal neurons during reward expectancy. *Journal of Neuroscience*, *30*(20), 7078–7087.

White, O. L., Lee, D. D., & Sompolinsky, H. (2004). Short-term memory in orthogonal neural networks. *Physical Review Letters*, *92*(14), 148102.

Wyble, B. P., Linster, C., & Hasselmo, M. E. (2000). Size of CA1-evoked synaptic potentials is related to theta rhythm phase in rat hippocampus. *Journal of Neurophysiology*, *83*(4), 2138–2144.

Yamaguchi, Y., Aota, Y., McNaughton, B. L., & Lipa, P. (2002). Bimodality of theta phase precession in hippocampal place cells in freely running rats. *Journal of Neurophysiology*, *87*(6), 2629–2642.

Young, J. J., & Shapiro, M. L. (2011). Dynamic coding of goal-directed paths by orbital prefrontal cortex. *Journal of Neuroscience*, *31*(16), 5989–6000. doi:10.1523/JNEUROSCI.5436-10.2011

Received May 19, 2016; accepted July 12, 2016.