Systems/Circuits

A Unified Mathematical Framework for Coding Time, Space, and Sequences in the Hippocampal Region

Marc W. Howard,[1] Christopher J. MacDonald,[1] Zoran Tiganj,[1] Karthik H. Shankar,[1] Qian Du,[2] Michael E. Hasselmo,[1] and Howard Eichenbaum[1]

[1]Center for Memory and Brain and Department of Psychological and Brain Sciences, Boston University, Boston, Massachusetts 02215, and [2]Department of Neuroscience, State University of New York, Upstate Medical University, Syracuse, New York 13210

The medial temporal lobe (MTL) is believed to support episodic memory, vivid recollection of a specific event situated in a particular place at a particular time. There is ample neurophysiological evidence that the MTL computes location in allocentric space and more recent evidence that the MTL also codes for time. Space and time represent a similar computational challenge; both are variables that cannot be simply calculated from the immediately available sensory information. We introduce a simple mathematical framework that computes functions of both spatial location and time as special cases of a more general computation. In this framework, experience unfolding in time is encoded via a set of leaky integrators. These leaky integrators encode the Laplace transform of their input. The information contained in the transform can be recovered using an approximation to the inverse Laplace transform. In the temporal domain, the resulting representation reconstructs the temporal history. By integrating movements, the equations give rise to a representation of the path taken to arrive at the present location. By modulating the transform with information about allocentric velocity, the equations code for position of a landmark. Simulated cells show a close correspondence to neurons observed in various regions for all three cases. In the temporal domain, novel secondary analyses of hippocampal time cells verified several qualitative predictions of the model. An integrated representation of spatiotemporal context can be computed by taking conjunctions of these elemental inputs, leading to a correspondence with conjunctive neural representations observed in dorsal CA1.

Key words: **entorhinal cortex; hippocampus; place cells; spatiotemporal context; subiculum; time cells**

Introduction

Space and time are both essential aspects of episodic memory, which integrates information about what, where, and when an event is experienced (Tulving, 1983). It is clear that the medial temporal lobe (MTL) is essential for episodic memory. Rodent neurophysiology studies have accumulated a tremendous amount of data describing the representation of "where" information (O'Keefe and Dostrovsky, 1971; Muller and Kubie, 1987; Wilson and McNaughton, 1993). In recent years, growing evidence has shown that neurons in the hippocampus and MTL also have robust temporal correlates (Pastalkova et al., 2008; MacDonald et al., 2011; Naya and Suzuki, 2011). Behavioral models of episodic memory and learning have focused on the representation of time (Estes, 1955; Brown et al., 2007; Balsam and Gallistel, 2009).

Received Dec. 19, 2012; revised Feb. 13, 2014; accepted Feb. 18, 2014.

Author contributions: M.W.H., M.E.H., and H.E. designed research; M.W.H., C.J.M., Z.T., K.H.S., and Q.D. performed research; M.W.H., Z.T., K.H.S., and Q.D. contributed unpublished reagents/analytic tools; M.W.H. and C.J.M. analyzed data; M.W.H., M.E.H., and H.E. wrote the paper.

This work was supported by AFOSR FA9550-12-1-0369. Colin Lever, Loren Frank, and Paul Lipton generously provided path data. Ehren Newman helped import some of these data. We thank Ehren Newman and Sam McKenzie for reading the paper.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr Marc Howard, Boston University, 2 Cummington Street, Boston, MA 02215. E-mail: marcwhoward777@gmail.com.

DOI:10.1523/JNEUROSCI.5808-12.2014

 0270-6474/14/344692-16$15.00/0

The brain constructs functions of variables that can be directly extracted from the current state of the physical world. This ability depends on specialized receptors that react to physical inputs in a particular range. For instance, the visual system can represent luminosity as a function of retinal position because there are photoreceptors at different locations; the auditory system can represent sound intensity as a function of frequency because different hair cells are stimulated by different frequency bands. Because there are not specialized receptors for past events or distant objects that are not visible, constructing functions of space and time presents a computational challenge. The current state of the world does not provide explicit information about the past. Similarly, the place code cannot rely on dedicated receptors that directly detect the location of distant objects in the dark (Gothard et al., 2001). To code space and time we need to support "hidden" dimensions that are not physically present in the input.

In this paper, we introduce a method for representing functions of hidden variables for which the time derivative is available. Using the same equations with different inputs we generate simulated cells that resemble time cells (MacDonald et al., 2011), history-dependent place cells (Frank et al., 2000; Wood et al., 2000), and boundary vector cells that code for the distance to a boundary of the enclosure (Lever et al., 2009). Novel secondary analyses of time cells confirm several qualitative predictions of the equations. Moreover, we show that these cell types can be combined conjunctively to account for a variety of place cell find-

ings from dorsal CA1 including changes in place fields over time (Mankin et al., 2012) and in response to spatial deformations of the environment (O'Keefe and Burgess, 1996).

Materials and Methods

In many cases, the nervous system constructs a representation of a function of a sensory variable, say the amplitude of a sound as a function of frequency. Consider the representation of a chord; two notes presented simultaneously. It is not sufficient to represent the chord as the average of the two tones; rather we require the amplitude of the sound for each frequency, a function. In sensory representations, functions of physical variables are possible because of specialized receptors. For instance, each hair cell has a tuning curve centered around some central frequency. We say that each hair cell supports part of the dimension of frequency. The activity of all of the hair cells ordered by their central frequencies gives a representation of the amplitude of the sound as a function of frequency.

Unfortunately, not all physical dimensions have specialized receptors. Consider the problem of representing the temporal history leading up to the present moment. For instance, we might want to represent the history of when one particular tone was presented as part of a melody. To construct the history we would want to have a set of cells that support the temporal dimension. The basic strategy is that we first construct an intermediate representation and then extract an estimate of the history from this intermediate representation. The intermediate representation is a set of cells that respond to the present input, each with a different time constant. The pattern of activity of the exponentially decaying cells in the intermediate representation does not resemble the history. However, we can extract an approximation of the history from the intermediate representation using a straightforward set of synaptic weights from the intermediate representation to the cells that represent the history. The cells in the estimate of the history behave as if they were specialized receptors tuned to various points in the past. The width of these cells' "tuning curves" grows for cells representing points further in the past. The result is a fuzzy representation of history, a function of past time, that grows more fuzzy for time points further in the past.

There are other physical dimensions other than time for which we lack specialized receptors that are extremely important behaviorally. For instance, the location of a landmark, such as a nest or a food source, as an animal forages around an environment requires a representation of spatial position. This paper describes a mechanism that enables a set of cells to function like specialized receptors for any physical dimension as long as the rate of change of that dimension is available at each moment. For instance, to represent position we need access to the velocity, the change in position per unit time. In the temporal case described above, the cells in the intermediate representation change as time passes. The end result of the calculation is a representation of time. The key insight is that access to the rate of change of the hidden variable per unit time enables the intermediate representation to change in a given moment only insofar as the hidden variable changes in that given moment. Rather than a representation of a function of time, the approximation yields a representation of a function of the hidden variable via path integration. We consider three variables that appear to be of particular relevance for neurons in the rodent hippocampus and related structures; time, position within a sequence, and allocentric position.

Mathematical framework. Figure 1 formalizes the case in which there are specialized receptors and lets us introduce some notation. Let us refer to a physical dimension in the world as x and a function of that dimension as $f(x)$ (Fig. 1). To be concrete, x might be sound frequency and $f(x)$ might be amplitude of a chord consisting of a low tone and a high tone. In the bottom of the figure, we see (schematically) the response of a set of receptors, each with a tuning curve with some width. Sorting the activity of the receptors by the center of their receptive fields gives a representation of the function in the world; we refer to the internal representation of $f(x)$ as $\tilde{f}(\overset{*}{x})$ Note that this representation is not necessarily a precise match to the physical function out in the world. For instance, there will likely be a non-zero width to the tuning curve of the receptors. There is also no reason to expect that the cells are equally spaced in their values of $\overset{*}{x}$. We refer to $\tilde{f}$ as a representation of f and $\overset{*}{x}$ as an internal dimension.

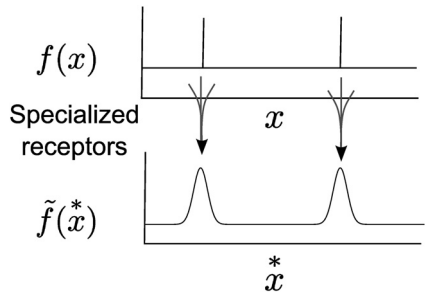

Figure 1. The brain constructs functions of physical variables. A function of some physical variable x, $f(x)$ is present in the environment. Here x could be a variable such as position along the retina or the frequency of tones in audition. Specialized receptors respond to the value of the function in the neighborhood of some particular value of x. Each receptor has its own tuning curve; together the set of receptors support the dimension x over some range. We refer to the internal estimate of the function $f(x)$ as $\tilde{f}(\overset{*}{x})$. Note that whereas x is a variable in the physical world, $\overset{*}{x}$ is a mapping from an ordering of cells onto that physical variable based on their firing correlates.

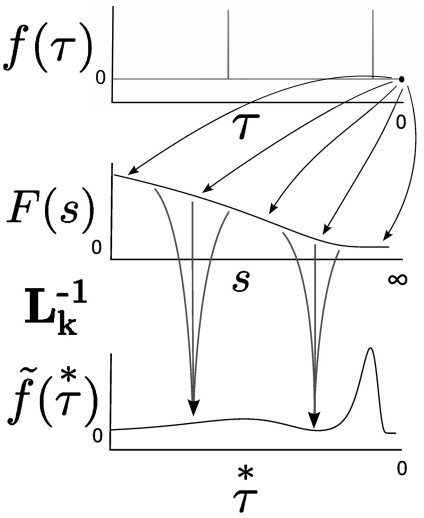

Figure 2. Constructing a function of past time. The top row shows the history leading up to the present moment ($\tau = 0$, right). There are no receptors that can directly detect events that happened in the past. To solve this problem, we assume that at each moment the present value of the stimulus (black dot) provides input to all of the cells in an intermediate representation $F(s)$. In the figure, the input is zero at $\tau = 0$ but was non-zero at two instants in the past (vertical lines). Those past values of $f(\tau)$ provided input to all of the cells in $F(s)$. The middle row shows $F(s)$ as a function of the rate constant s at $\tau = 0$. The time constant of each cell in $F(s)$ is just $1/s$. The value of ∞ for s corresponds to a time constant of zero and is simply meant to align the axis of this schematic figure correctly. The cells in $F(s)$ with short time constants (right) have decayed almost back to zero; the cells with longer time constants (left) have not decayed nearly as much. In principle, the pattern of activation across cells in $F(s)$ contains complete information about the history of the inputs from prior moments. The bottom row shows the estimate of the reconstruction $\tilde{f}(\overset{*}{\tau})$. At each moment $\tilde{f}$ is constructed from $F(s)$ via a matrix of feedforward connections $\mathbf{L}_k^{-1}$. These connections from F to $\tilde{f}$ can be understood as several projections with lateral inhibition in series, analogous to sensory processing (see text for details). Each cells' value of $\overset{*}{\tau}$ is aligned with the time constant of the corresponding cell in $F(s)$. We can see that at time $\tau = 0$, $\tilde{f}(\overset{*}{\tau})$ estimates the past values of $f(\tau < 0)$; across cells $\tilde{f}(\overset{*}{\tau})$ contains a smeared estimate of $f(\tau)$.

Figure 2 introduces notation and provides a schematic overview for how to approximate a function of time, $f(\tau)$. Our goal at time τ is to be able to construct the function for all time points τ' leading up to the present moment, $f(\tau' < \tau)$. At each moment, the present value of $f(\tau)$ drives an intermediate representation $F(s)$. The cells in the intermediate representation are leaky integrators indexed by their rate constant s. The time constant of each cell in the intermediate representation is just $1/s$.

We use a linear operator $\mathbf{L}_k^{-1}$ to construct an estimate of $f(\tau)$ from $F(s)$; the approximation is referred to as $\tilde{f}(\overset{*}{\tau})$. The pattern of activity $\tilde{f}(\overset{*}{\tau})$ is an imprecise estimate of $f(\tau)$, with a temporal precision that decreases for times further in the past. We generalize the framework by modulating the leaky integrators by a function $\alpha(\tau)$ (not shown). If $\alpha(\tau)$ is the rate of change of some physical variable x and if $f(\tau)$ is only a function of $x(\tau)$, then $\tilde{f}$ estimates a function of x rather than a function of time and we write $\tilde{f}(\overset{*}{x})$.

The remainder of this subsection describes the equations in detail and explains why they work. The next subsections describe the qualitative properties of the representation in the temporal and spatial case.

Let $F(s)$ describe the firing rate of a set of cells each with a different positive real value of s. At each moment, each of the cells in $F(s)$ receives input from the current value of $f(\tau)$ and updates its current value by obeying the following differential equation:

$$\frac{dF(s,\tau)}{d\tau} = \alpha(\tau)[-sF(s,\tau) + f(\tau)]. \quad (1)$$

F is a function of both s, which indexes the cells, and of time τ. We will sometimes write $F(s, \tau)$ when considering the time dependence of $F(s)$; sometimes we will simply write F to keep the expressions simple. The function $\alpha(\tau)$ is the same for all of the cells in $F(s)$ but can change over time. To code for time, we set $\alpha(\tau) = 1$. In this case, we find:

$$\frac{dF(s,\tau)}{d\tau} = -sF(s,\tau) + f(\tau). \quad (2)$$

Now suppose that the environment contains some hidden variable x that changes as a function of time $x(\tau)$. Further suppose that the current value of x controls the current value of f so that $f(\tau) = f[x(\tau)]$. To code for a hidden variable $x(\tau)$ we set $\alpha(\tau) = \frac{dx}{d\tau}$. Now, dividing both sides of Equation 1 by α and applying the chain rule we find:

$$\frac{dF(s,x)}{dx} = -sF(s,x) + f(x). \quad (3)$$

Note that because F is only a function of f and under these circumstances f is only a function of x, F is also a function of x. Note that Equation 3 is identical to Equation 2 with time τ replaced by the variable x. Also note that when $\alpha(\tau)$ is velocity in the x direction and f is a function of x, $F(s,\tau)$ meets the conditions for a path integrator in the x direction, as described by Issa and Zhang (2012).

Now, consider the temporal case with $\alpha(\tau)$. Solving Equation 2 yields

$$F(s,\tau) = \int_{-\infty}^{\tau} e^{-s(\tau-\tau')} f(\tau')\, d\tau'. \quad (4)$$

That is, at time τ the activity of each cell in $F(s)$ is an exponentially weighted sum over the prior states of the input at times $\tau' < \tau$; each cell's time constant is $1/s$. Solving Equation 3 would yield a precisely analogous expression giving $F(s)$ as a function of the hidden variable x instead of time τ.

Comparing Equation 4 to the standard definition of the Laplace transform, we see that $F(s)$ is just the Laplace transform of $f(\tau' < \tau)$ with real coefficients. Because Equation 3 yields a precisely analogous expression, when the conditions that yield Equation 3 hold, $F(s)$ is the Laplace transform of $F(x)$. In operator notation we can write:

$$F = \mathcal{L}f. \quad (5)$$

The knowledge that F is the Laplace transform of f is very powerful. The Laplace transform of a particular function f causes a unique F; no information about f is lost in transforming it into F. Moreover, like the Fourier transform, the Laplace transform is invertible. That means that applying the inverse Laplace transform to F would recover f. If f is a function of time, the inverse Laplace transform of F would recover the original function $f(\tau)$. If Equation 3 holds, then the inverse Laplace transform of F would recover a function of x rather than time. Recovering the function f amounts to finding a neurally plausible way to invert the Laplace transform.

Let us consider what it would mean to successfully invert the Laplace transform and recover the function f in the firing rate across a set of cells. In much the same way that the sensory receptors enabled us to support a dimension x (Fig. 1), each cell in the reconstruction would support some part of the domain of f. In the case of time, a cell in the reconstruction would fire in response not to the current value of the input, but in response to the value the input had a certain time in the past. Each cell would represent a different delay and the temporal history would be written into a pattern across cells. A perfect inversion of the transform would require zero-width tuning curves. Perfect reconstruction would also require an infinite number of cells to represent a function. An imperfect reconstruction, with broad tuning curves, is clearly preferable.

Following Shankar and Howard (2012), we define a linear operator $\mathbf{L}_k^{-1}$ that approximates the inverse Laplace transform. The operator describes a set of connections from the cells in F to the cells in the reconstruction, $\tilde{f}$:

$$\tilde{f} \equiv \mathbf{L}_k^{-1} F(s) \equiv C_k\ s^{k+1} \frac{d^k}{ds^k}\ F(s). \quad (6)$$

This equation requires significant elaboration; the remainder of this subsection is devoted to unpacking it. First, we note that k is an integer that controls the accuracy of the reconstruction. It can be shown (Post, 1930) that in the limit as k grows without bound, $\mathbf{L}_k^{-1}$ is exactly the inverse Laplace transform. When k is finite, the reconstruction is imperfect. We set $k = 4$ in this paper, except where noted.

Next, let us consider the left hand side of Equation 6. Each cell in $\tilde{f}$ is aligned to a particular value of s. The right hand side of Equation 6 describes how the activity of each cell in $\tilde{f}$ is determined from the activity of cells in $F(s)$ around that particular value of s. Depending on the properties of f, we can assign a mapping between the value of s for each cell in $\tilde{f}$ and a physical value in the world. Note that the choice of mapping does not affect the firing correlates of a cell in $\tilde{f}$; the firing correlates of a cell in $\tilde{f}$ are completely controlled by its value of s, the input $f(\tau)$ driving it, and the properties of $\alpha(\tau)$ modulating $F(s)$. Put another way, s describes a physical property of each neuron in $\tilde{f}$, whereas the mapping is a choice we make to describe its firing correlates with variables out in the world. When $\alpha(\tau) = 1$, it is convenient to consider $\tilde{f}$ as a function of $\overset{*}{\tau}$, where $\overset{*}{\tau}$ is defined as:

$$\overset{*}{\tau} \equiv -k/s, \quad (7)$$

as this mapping between $\overset{*}{\tau}$ and s aligns each cell in $\tilde{f}$ with the physical time at which it peaks firing (Fig. 2; see also below). With this choice, the units of $\overset{*}{\tau}$ are the same as the units of $1/s$. With this mapping, $\overset{*}{\tau}$ can be understood as the time in the past (with the present set to 0) that each cell in $\tilde{f}$ supports. If $F(s)$ corresponds to the Laplace transform of some variable s, we can consider $\tilde{f}$ a function of a variable $\overset{*}{x}$ chosen as:

$$\overset{*}{x} \propto k/s. \quad (8)$$

The proportionality allows us to take into account units from $\alpha(\tau)$ and choose the sign of $\overset{*}{x}$ and ensure that $\overset{*}{x}$ aligns with physical position x.

Now, let us step through the right-hand side of Equation 6 to understand how it approximates the reconstruction of the original function f. There are three factors on the right hand side of Equation 6. The first factor, C_k is a constant that depends on the value of k (Shankar and Howard, 2012). This value does not change across cells in $\tilde{f}$ nor across time. The first factor simply sets the scale of the overall firing rates and ensures that the sign of the reconstruction corresponds to the sign of the input function. The second factor in Equation 6, s^{k+1}, says that the overall value of the firing rate of each cell in $\tilde{f}$ depends on its value of s. This term does not change over time for a particular cell. But it does affect the relative firing of one cell in $\tilde{f}(\overset{*}{\tau})$ compared with other cells in $\tilde{f}(\overset{*}{\tau})$. Cells in $\tilde{f}$ that are aligned to a smaller value of s, corresponding to a longer time

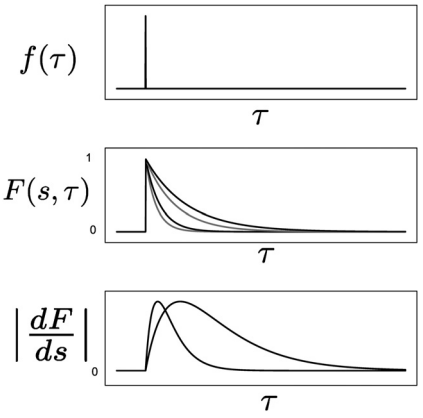

Figure 3. How lateral inhibition causes a peak in firing after the stimulus. Top, $f(\tau)$ is non-zero at one point in time. Middle, Time course of $F(s, \tau)$ for several values of s. Each of the black lines gives $F(s, \tau)$ for a particular value of s; the gray lines give $F(s, \tau)$ for a nearby value of s. Note that the difference between each black line and its accompanying gray lines is zero immediately after the stimulus but also zero for very long times. The difference is maximal at an intermediate time that depends on the value of s. Bottom, The magnitude of the difference between pairs of adjacent values of $F(s, \tau)$ is shown, i.e., the distance between each pair of lines. This difference is proportional to the derivative with respect to s. Note that the difference peaks after the stimulus; the time of the peak depends on the value of s.

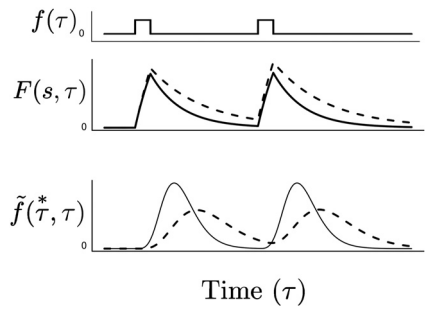

Figure 4. Dynamics of cells when $\alpha(\tau) = 1$. Top, The stimulus function f providing input to a set of cells is non-zero for two periods of time. Cells in $F(s)$ respond to non-zero f and then decay exponentially after the input is turned off. Different cells respond with different decay constants s. Cells in $\tilde{f}$ do not respond to the stimulus immediately, but after some characteristic delay. Different cells have different delays determined by their value of $\overset{*}{\tau}$. The spread of the firing of each cell in $\tilde{f}(\overset{*}{\tau})$ depends on its value of $\overset{*}{\tau}$. The value of $\overset{*}{\tau}$ for a cell in $\tilde{f}$ is determined by the value of s of the cells in $F(s)$ that project to it. Note also that the activity of cells in $\tilde{f}$ show an asymmetric profile in time.

constant, require more input to reach the same level of activity as a cell aligned to a higher value of s.

The last factor in Equation 6 is responsible for the changes in the state of $\tilde{f}(\overset{\star}{\tau})$ over time. The notation $\frac{d^k F}{ds^k}$ refers to the k'th derivative of $F(s)$ with respect to s. For each cell in $\tilde{f}$, the derivative is calculated at the value $s = -k/\overset{\star}{\tau}$. The derivative with respect to s is just the difference between neighboring values of $F(s)$. It has long been appreciated that many receptive fields in the visual system approximate second spatial derivatives (Marr and Hildreth, 1980). The synaptic weights needed to approximate the k'th derivative are well understood (Shankar and Howard, 2012, 2013; see simulation methods below). A simple feedforward circuit with lateral inhibition is sufficient to approximate the second derivative. Higher-order derivatives can be calculated by placing circuits that calculate lower-order derivatives in series. For instance, the fourth derivative is just the second derivative of the second derivative.

To gain some intuition into how taking derivatives of $F(s)$ can lead to an estimate of past history, let us consider a simple thought experiment. Suppose that $\alpha(\tau) = 1$, and that the input function is zero except for a very high value for one brief moment of time. In Figure 3, middle, we follow the activity of selected cells in F with various values of s in time after the stimulus was presented. Under these circumstances, we can readily write down the firing of each cell in $F(s)$. If the non-zero input was presented a time τ_0 in the past, the firing of each cell in $F(s)$ is given by

$$F(s) = e^{-s\tau_o}.$$

That is, each cell in F shows exponential decay with a time constant controlled by its value of s. Now, at each moment in time the derivative with respect to s is the difference between the activity of cells with adjacent values of s. Note that immediately after the stimulus is presented $\tau_0 = 0$ and all of the cells in F have activation 1; when $\tau_0 = 0$, the derivative with respect to s is zero because all of the cells have the same activity. At very long times, all of the cells in F have decayed back to zero and the derivative again vanishes. The time at which the derivative peaks depends on s and is later for smaller values of s corresponding to longer time constants. We can formalize this a bit by computing the k'th derivative with respect to s. When s values are equally spaced, we find:

$$\frac{d^k F}{ds^k} = (-1)^k \tau_o{}^k e^{-s\tau_o}.$$

Ignoring the change in sign (the C_k term in Eq. 6 changes sign to cancel out this part of the derivative), the magnitude of the derivative is the product of a power law function that goes to zero as τ_0 decreases to zero (τ_o^k) and an exponential function that decays to zero as τ_0 increases ($e^{-s\tau_o}$). By setting the derivative equal to zero we find that the k'th derivative of $F(s)$ around s is maximal a time $\tau_{\max} = k/s$ after the stimulus was presented. Here we see that the choice of indexing the cells by $\overset{\star}{\tau} \equiv -k/s$, results in a cell in $\tilde{f}$ aligned with the value s peaking in its activity $\overset{\star}{\tau}$ seconds after presentation of a brief stimulus. This aligns $\tilde{f}(\overset{\star}{\tau},\tau)$ with $f(\tau + \overset{\star}{\tau})$ (recall that $\overset{\star}{\tau}$ is negative). This mapping makes $\tilde{f}(\overset{\star}{\tau},\tau)$ a good reconstruction of the value of f had at a time $\overset{\star}{\tau}$ in the past. An analogous calculation also shows that the width of each cell's firing is also controlled by s (Shankar and Howard, 2012).

Properties of cells in the temporal case. Figure 4 illustrates the important properties of cells in $F(s)$ and $\tilde{f}(\overset{\star}{\tau})$ in the temporal case when $\alpha(\tau) = 1$. Here we assume that the input is zero except for two square wave pulses. Before receiving input, $F(s)$ is zero. After non-zero values of f are encountered, the firing rates of cells in $F(s)$ grow exponentially over time, like a charging capacitor. After the input becomes zero, the firing rates of cells in $F(s)$ decay exponentially over time. The time constant of each cell in $F(s)$ is controlled by its value of s (the time constant is $1/s$). In contrast, cells in $\tilde{f}(\overset{\star}{\tau})$ do not respond immediately to the stimulus. Rather, cells in $\tilde{f}(\overset{\star}{\tau})$ fire a characteristic time after a non-zero value of f is presented.

The delay in the firing of cells in $\tilde{f}(\overset{\star}{\tau})$ can be understood as a consequence of their roles in reconstructing $f(\tau)$. A cell in $\tilde{f}(\overset{\star}{\tau},\tau)$ estimates f not at the current moment τ, but the value of f a time $\overset{\star}{\tau}$ in the past (Fig. 2). For instance, a cell in $\tilde{f}$ supporting $\overset{\star}{\tau}$ 5 s in the past does not respond immediately to a transient stimulus because that stimulus is not yet part of the history 5 s in the past. After some time has passed, the stimulus is now in that part of the history. A cell in $\tilde{f}$ supporting $\overset{\star}{\tau}$ 10 s in the past would fire later than the cell supporting $\overset{\star}{\tau}$ 5 s in the past. A transient stimulus causes a sequence of cells to fire after it was presented.

Figure 4 illustrates the property that cells coding more distant parts of the history show firing that is more spread out in time. This is a natural outcome of the time derivative of an exponential (Fig. 3). The spread implies that events experienced further in the past are represented with less temporal accuracy than events experienced more recently. This is an extremely useful property for describing behavioral aspects of memory and timing (Gallistel and Gibbon, 2000; Brown et al., 2007; Chater and Brown, 2008), and it holds for all choices of k (Shankar and Howard, 2012). Note that this spread is not a consequence of "noise" *per se*. The spread is deterministic and is a commitment to representing the history with decreasing accuracy. Shankar and Howard (2012, 2013) provide detailed treatments of the effect of noise on the temporal representation.

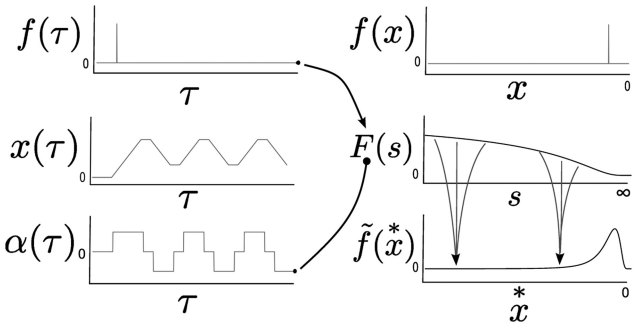

Figure 5. Schematic figure illustrating the encoding and extraction of a hidden stimulus dimension. We assume that the input f is a function of some hidden time-varying variable $x(\tau)$; that is, the value of f at each time only depends on the value of x at that time. In addition we assume we have access at each moment to the time derivative of $x(\tau)$, which enters Equation 1 as $\alpha(\tau)$. At each moment, the current value of f (top) and the current value of the time derivative of x are used to update $F(s,\tau)$. To be concrete, let us assume that x is physical location and $f(\tau)$ is contact with a fixed landmark as the animal moves around the environment. That is, at times when the position of the animal is such that it touches the landmark, $f(\tau)$ is non-zero. On the left, f is non-zero early in time as the animal encounters the landmark. After encountering the landmark, the animal moves away, changing direction several times, and ending up close to the landmark. The plots on the right of the figure shows $f(x)$, the true position of the landmark relative to the current position of the animal at the last moment in time, along with the state of the intermediate representation $F(s)$ and the reconstruction $\tilde{f}$. In this figure we have chosen the mapping from s onto $\tilde{x}^*$ to make $\tilde{x}^*$ negative to emphasize the analogy to the temporal case in Figure 2. At each moment $\tilde{f}(\tilde{\tau}^*)$ provides an estimate of the distance to the landmark relative to the animal's current location.

Estimating hidden dimensions by varying $\alpha(\tau)$. Here we consider the more general case in which we represent functions of a hidden variable x rather than time τ. To represent functions of x, we require that $\alpha(\tau) = \frac{dx}{d\tau}$ and that the input f at time τ depends only on the value of x at that time τ, $f(\tau) = f[x(\tau)]$ (Fig. 5). Here, $\tilde{f}(\tilde{x}^*,\tau) = \mathbf{L}_k^{-1}F(s,\tau)$ estimates the current value of $f(x)$ at time τ.

Let us work through a simple concrete example. Suppose that x is physical position along a linear track. A rat starts out touching the landmark at one end and then runs back and forth along the track, not quite returning to the location of the landmark. Let the input $f(\tau)$ be one at time 0 when the animal is in contact with the landmark but zero at all other times. Let $\alpha(\tau)$ be the animal's signed velocity in the x direction, $\alpha(\tau) = \frac{dx}{d\tau}$. Let the landmark be at position $x = 0$, so that $f(x)$ is 1 when the animal is at $x = 0$ at one end of the track but zero elsewhere. As the animal moves away from the landmark, velocity is positive causing $\alpha(\tau)$ to be positive and thus the cells in $F(s,\tau)$ decay exponentially in time (Eq. 1). Because time is perfectly correlated with position on this outbound journey, the activity of the cells also decays exponentially as a function of position x (Eq. 3). Now, after the animal reaches the end of the linear track and turns back toward the landmark with the same constant speed, $\frac{dx}{d\tau}$ changes sign. On the return journey, $\alpha(\tau)$ is thus negative and the values of $F(s,\tau)$ grow exponentially until the animal reaches the start of the track, recapitulating their initial state. As the animal repeats this journey many times (Fig. 6), although the states of $F(s,\tau)$ change over time, they are always the same for a particular position x. Like a bead on a wire, the values of $F(s,\tau)$ slide down an exponential function as the animal moves away from the origin and then slide back up as the animal returns. Because at each moment $\tilde{f}$ only depends on F, and because F is a function of position, $\tilde{f}(\tilde{x}^*,\tau)$ is also a function of position and is an internal estimate of the animal's location relative to the landmark. For convenience, in Figure 6, we define $\tilde{x}^*$ to be positive, $\tilde{x}^* = k/s$, so that it codes the distance of the landmark along the x direction. In the same way that in the temporal case cells in $\tilde{f}(\tilde{\tau}^*)$ became active a certain time after the start of the delay, in the spatial case cells in $\tilde{f}(\tilde{\tau}^*)$ become activated when the animal's position is a characteristic distance from the starting point. Analogous to the temporal case, the cells in $\tilde{f}(\tilde{x}^*,x)$ show an asymmetric spatial profile; cells with larger values of $\tilde{x}^*$ show a wider spatial spread of their firing field.

Spatial location is a concrete example of a variable x that could be represented in this way. However, it should be noted that mathematically at least this approach could be applied for any variable for which the time derivative is available. In principle at least, one could imagine using a similar approach to represent numerosity or the location of moving targets that cease to be visible, or to integrate views of an object across multiple eye movements or from different perspectives. In this paper, we restrict our attention to variables that appear to be represented in the rodent hippocampus and parahippocampal cortex.

Overview of secondary analyses and simulations. The goal of the simulations is to compare predictions from the mathematical framework to the firing correlates of actual neurons. In any particular experiment, the firing correlates of a cell encoding F or $\tilde{f}$ depends on its value of s, on the properties of the input, $f(\tau)$ and the properties of $\alpha(\tau)$. We will consider three cases (Table 1).

In Case I we assume that $\alpha(\tau)$ is constant and the stimulus function $f(\tau)$ corresponds to the time-varying presence of a nonspatial stimulus. With these settings, the mathematical framework gives rise to a representation of recent history. With many cells receiving input corresponding to many nonspatial stimuli, the cells in the reconstruction would contain both what and when information about the history leading up to the present moment. Previous work has shown that this representation can be used to describe behavioral data from diverse fields, including conditioning, interval timing, episodic memory, and working memory (Shankar and Howard, 2012; Howard and Eichenbaum, 2013). Our interest here is in determining whether cells in the brain exhibit qualitative properties similar to those predicted for $\tilde{f}(\tilde{\tau}^*)$. To this end, we conduct secondary analyses on results from time cells observed in the rodent hippocampus (MacDonald et al., 2011).

Case II constructs a representation of sequences of stimuli not as a function of time but as a function of the path taken leading up to the present position. In Case II, we set $f(\tau)$ to correspond to a simple simulation of a head direction cell with preferred direction θ, $f_\theta(\tau)$, and set $\alpha(\tau)$ to be equal to the animal's speed $|\mathbf{v}(\tau)|$ (Hasselmo, 2007). The notation $\mathbf{v}$ denotes the two-dimensional velocity. Note that speed does not depend on the direction the animal is traveling, merely how fast it is going. If we define $p(\tau)$ as a variable that parameterizes distance traveled along the path, then under these circumstances $\tilde{f}(\tilde{p}^*)$ codes for the movements taken along the path leading up to the present, where $\tilde{p}^*$ is an internal estimate of distance along the path. We will see that simulated cells in $F(s,\tau)$ and $\tilde{f}(\tilde{p}^*,\tau)$ exhibit history-dependent firing correlates similar to those of neurons in the hippocampus and entorhinal cortex during navigation along linear mazes (Frank et al., 2000; Wood et al., 2000; Lipton et al., 2007).

Case III constructs an allocentric representation of the position of a landmark relative to the current position of the animal. In Case III $\alpha(\tau)$ is equal to the component of the animal's allocentric velocity at time τ along a particular direction, $v_\theta(\tau)$. Unlike Case II, $\alpha(\tau)$ can take both positive and negative values in Case III. In Case III, the input $f(\tau)$ is solely determined by the animals location at time τ, $f[x(\tau)]$, corresponding to physical contact with a fixed landmark. In the one-dimensional case (or with a point landmark), this satisfies the conditions for Equation 3 to hold and the inverse Laplace transform constructs an approximation of the spatial location of the landmark along the direction in which velocity is defined. In two-dimensional environments, we set $f(x)$ to proximity with boundaries of the enclosure with a specific orientation. In this case, simulated cells comprising $F(s,\tau)$ have spatial correlates that resemble border cells observed in the entorhinal cortex and subiculum (Solstad et al., 2006; Savelli et al., 2008; Lever et al., 2009); simulated cells in $\tilde{f}(\tilde{x}^*)$ resemble boundary vector cells that fire a certain allocentric distance from a landmark (Lever et al., 2009).

The three "pure" cases in Table 1 show a correspondence to firing correlates of cells from a variety of regions in the hippocampal formation and parahippocampal cortices. Different simulated cells reflect a variety

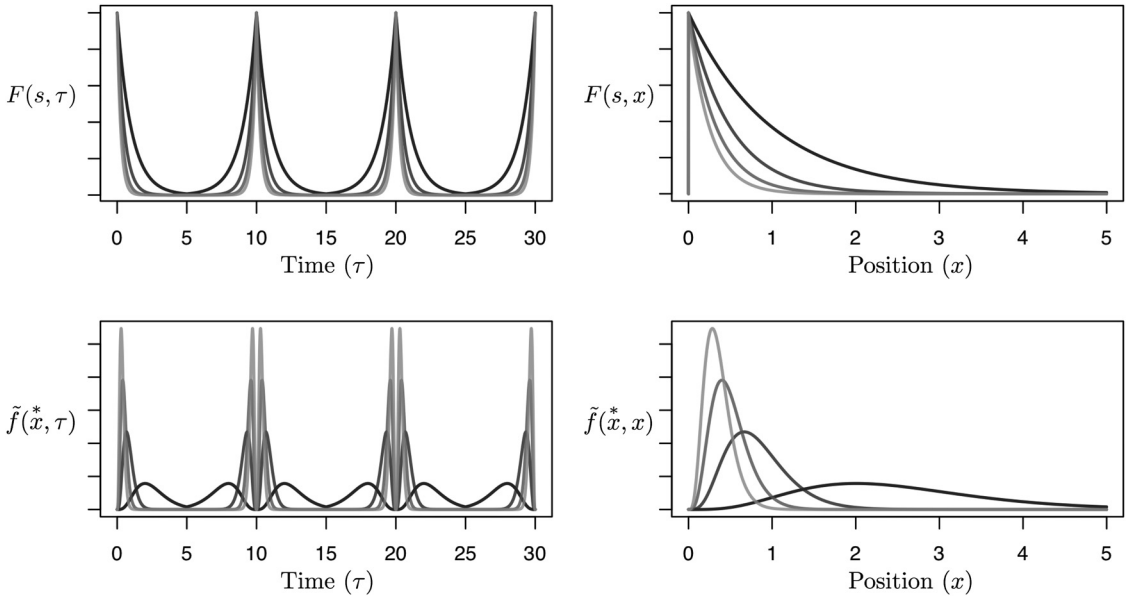

Figure 6. A worked example of coding for a hidden dimension, here spatial location of a landmark. The animal starts at a position on the left-hand side of a one-dimensional track where a stimulus is encoded at position zero. The input f is non-zero at this one moment. The animal then runs left-to-right and back repeatedly at constant speed. The left column plots firing rates as a function of time. The sign of velocity, and hence $\alpha(\tau)$ depends on the direction of motion. The right column replots these values as a function of position, x. The top row plots gives values for $F(s, \tau)$ (left) and $F(s, x)$ (right). The bottom row shows $\tilde{f}(\overset{*}{x},\tau)$ (left) and $\tilde{f}(\overset{*}{x},x)$ (right). For this figure, we have chosen the mapping $\overset{*}{x} = k/s$ so that $\overset{*}{x}$ is positive for positive values of position. Differently shaded lines correspond to different cells with different values of s (top) and the corresponding values of x (bottom). Thus, this set of cells codes for different positions of the animal relative to the landmark. See text for details.

Table 1. Firing correlates for F and $\tilde{f}$ with different input functions and different modulatory functions

	$\alpha(\tau)$	$f(\tau)$	$F(s)$	$\tilde{f}$
Purely temporal	1	$f(\tau)$	Exponentially-decaying	Time cells
History-dependent spatial	$\lvert\mathbf{v}(\tau)\rvert$	$f_\theta(\tau)$	Trajectory-coding cells	Retrospective splitter cells
Purely spatial	$v_\theta(\tau)$	$f[(x(\tau))]$	Border cells	Boundary vector cells

of different inputs $f(\tau)$ and a variety of sources of multiplicative modulation $\alpha(\tau)$. A cell that receives input from several "pure" cells would show a mixture of these firing correlates. We consider the firing correlates of cells taking conjunctive inputs from "pure" cells. The firing correlates of these conjunctive cells correspond well to the firing correlates of principal cells in dorsal CA1, enabling us to account for both canonical place fields and the change in place fields with deformations in the environment and changes over time. Following previous work (O'Keefe and Burgess, 1996; Hartley et al, 2000; Barry et al., 2006), conjunctions of boundary vector cells (Case III) result in canonical place cells that respond appropriately to changes in the dimensions of the environment (O'Keefe and Burgess, 1996). Moreover, conjunctions of boundary vector cells (Case III) and time cells (Case I) yield cells with place fields that drift over long periods of time (Mankin et al., 2012). Although not explicitly simulated, this strategy is broadly consistent with extensive evidence supporting conjunctive coding in the hippocampus (Shapiro et al., 1997; Wood et al., 1999; Anderson and Jeffery, 2003; Komorowski et al., 2009, 2013).

Secondary analysis methods. To compare detailed predictions of the mathematical framework to firing of time cells in Case I, we conducted secondary analyses on multiunit recordings from dorsal CA1 that have previously been reported (MacDonald et al., 2011). The MacDonald et al. (2011) dataset contained 245 putative pyramidal cells recorded across a total of four sessions each recorded from a different rat. For each cell, we estimated the peak of firing during the delay period from a smoothed firing rate distribution (Fig. 7*a*). There was variation in the duration of the delay period across trials and across rats. Across animals, the shortest delay period was 5.3 s; we used this as the end of the delay interval for all sessions. To avoid edge effects we allowed spikes slightly before and slightly after the end of the delay period to enter the analyses, but only included cells with a well defined peak that was at least 300 ms from the beginning and end of the delay period. The spikes each cell fired during the interval from 300 ms before the start of each delay period up to 6 s after the start of each delay period were taken before smoothing. Smoothing was implemented with the R function *density*(), which uses a kernel size proportional to the SD of the data, to estimate the density in 512 bins between −300 ms and 6 s. We only included cells with peak times of >300 ms and <5 s in the analyses. We also required that the leading edge of the half-height range to be greater than zero and the tailing edge of the half-height range to be <6 s. In all, 63 cells were included in the analyses reported here.

To generate model predictions for the cosine of the angle between two population vectors, we numerically evaluated the following expression:

$$\cos\theta(\tau_1, \tau_2) = \frac{I(\tau_1, \tau_2)}{\sqrt{I(\tau_1, \tau_1)}\sqrt{I(\tau_2, \tau_2)}}, \tag{9}$$

where $I(\tau_1,\tau_2)$ is the inner product between $\tilde{f}(\overset{*}{\tau},\tau_1)$ and $\tilde{f}(\overset{*}{\tau},\tau_2)$ evaluated for a given range of values of $\overset{*}{\tau}$. More explicitly, $I(\tau_1,\tau_2)$ is given by

$$I(\tau_1,\tau_2) = \int_{\overset{*}{\tau}_{\min}}^{\overset{*}{\tau}_{\max}} \left(\frac{\tau_1\tau_2}{\overset{*}{\tau}^2}\right)^k e^{k(\tau_1+\tau_2)/\overset{*}{\tau}}\, d\overset{*}{\tau}. \tag{10}$$

Equation 10 is based on expressions for the similarity of $\tilde{f}$ with itself derived in Shankar and Howard (2012) with the number density of nodes set to be a constant in $\overset{*}{\tau}$, consistent with the uniform distribution of modes we observed across cells in the experiment. Note that this does not

correspond to equal spacing of nodes in s. Consistent with the experimental values we set $\overset{*}{\tau}_{\min} = 0.3$ and $\overset{*}{\tau}_{\max} = 5.3$.

Simulation methods. In Cases II and III, we provided input to the model taken from paths through an environment of the appropriate shape. A single session was taken for each environment. In Case II, paths for the W-maze were taken from Frank et al. (2000), and paths for the alternating T-maze were taken from Lipton et al. (2007). In Case III, paths for open field simulations were taken from Lever et al. (2002) (circular environment) and Lever et al. (2009) (all other environments). Paths were smoothed with the R function *kernapply*() with a Daniell kernel and a width of six bins. Velocity for each time point was computed as the average of the forward and backward difference in position. For the conjunctive simulation of place fields (see Fig. 12), we simulated paths using an algorithm used by Brunel and Trullier (1998).

In all simulations, we used an exponential estimate of the solution of Equation 1 with $\delta\tau$ set appropriate to the experiment. Values of s were chosen to be equally spaced. At each time step, $\tilde{f}(\overset{*}{\tau},\tau)$ was computed from $F(s,\tau)$ with a matrix approximation to $\mathbf{L}_k^{-1}$ (Shankar and Howard, 2012, 2013). This is a simple feedforward connection from $F(s)$ to $\tilde{f}$. As mentioned earlier, each node in $\tilde{f}$ corresponds to a particular value of s. The feedforward connection to that node is given by alternating excitatory and inhibitory weights to approximate the k'^{th} derivative. For instance, to estimate the first derivative around a particular value of s, one would take the firing rate of the next cell with a higher value of s and subtract from that the firing rate of the next cell with a smaller value of s. To compute the second derivative one would take the derivative of the derivative. Ultimately, this can be expressed as a set of weights multiplying the original firing rates. To estimate the fourth derivative of $F(s)$ around s_o, we simply took:

$$\frac{d^4F(s_o)}{ds^4} \simeq \frac{1}{24}F(s_{-2}) - \frac{1}{6}F(s_{-1}) + \frac{1}{4}F(s_o) - \frac{1}{6}F(s_{+1}) + \frac{1}{24}F(s_{-2}). \quad (11)$$

To compute $\tilde{f}$ we multiplied this value by the corresponding value of s raised to the $k + 1^{\text{st}}$ power. However, because the results of the simulations all report relative firing rate of one cell over time, this scaling factor does not make a difference when comparing one cell to itself.

In the open field simulations used to generate border cells and boundary vector cells (Case III) we allowed the input f corresponding to contact with a border to be non-zero when the distance from the boundary was <3 cm. In the open field simulations, we estimated $F(s)$ from pairs of monotonically decaying cells to avoid numerical errors due to exponential growth. Given a particular direction θ, we constructed a pair of cells with the appropriate value of s but with velocity inputs in opposite directions. For instance, if cell 1 took the velocity along an angle θ, cell 2 took the velocity in the direction $\theta + \pi$. The α for each of these cells was set to be zero if the signed velocity was negative. So, when the animal moves in the direction θ, cell 1 decays and cell 2 remains constant with $\alpha = 0$. When the animal moves in the direction $\theta + \pi$, cell 1 remains constant, whereas cell 2 decays. To estimate $F(s)$ with $\alpha(\tau)$ corresponding to velocity in the direction θ, we took $F(s)$ equal to the ratio of cell 1 to cell 2 when cell 2's firing rate was greater than cell 1's firing rate and zero otherwise. Note that for positive values of x, the ratio is exactly equal to $F(s)$. For negative values of x, the inverse Laplace transform cannot be constructed (this would be analogous to constructing future values of history) so setting $F(s)$ to zero does not affect the reconstruction.

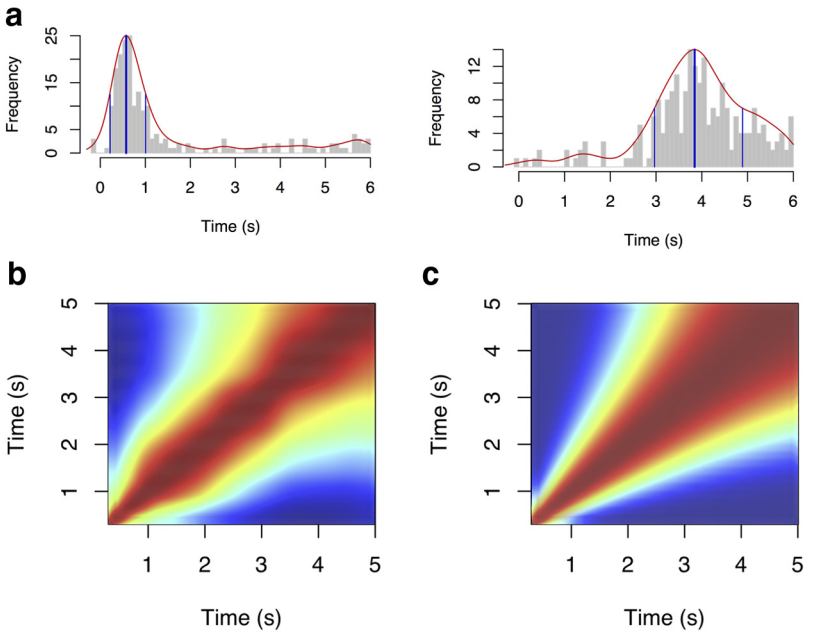

Figure 7. Comparison of experimentally observed hippocampal time cells to $\tilde{f}(\overset{*}{\tau})$. ***a***, Experimentally observed time cells exhibit increasing spread with their delay and asymmetry. The firing of two representative cells is shown as a function of time. In each plot, the total number of spikes fired in each bin is shown as a function of time synchronized to the onset of the delay period. The smooth red line gives an estimate of the density of spikes as function of time. The thick vertical blue line gives the estimate of the cell's mode; the two thin vertical blue lines give the estimate of the half-height region. Note that the cell that fires later in the delay also shows a wider spread. Note that both cells are asymmetric, with the distance between the mode and the end of the half-height region greater than the distance between the mode and the beginning of the half-height region. ***b***, Empirical ensemble similarity during the delay period. Color scale gives the ensemble similarity (cosine of the angle between the vectors) of the smoothed population vectors for each pairs of times. ***c***, Analogous plot to ***b*** computed for $\tilde{f}$ with $k = 8$ and a uniform distribution of characteristic delays. Data are from a secondary analysis of MacDonald et al. (2011).

For the conjunctive simulations (see Fig. 12), we took sums of products of "pure" inputs. In the simulations of canonical place cells responding to changing the dimensions of an open field environment we took sums of products of boundary vector cells ($\tilde{f}$ from Case III) computed as described above. For the simulation showing a place field that changes over time, we took the product of a place field, i.e., a conjunction of boundary vector cells, and a time cell ($\tilde{f}$ from Case I).

Results

Case I: purely temporal representation

In MacDonald et al. (2011), rats had to wait in a small enclosure during the delay period of a memory task. The rat had to remember which of two objects was presented to choose appropriately among a pair of odors presented after the delay and obtain reward. In analogy to place cells that fire when the animal is in a circumscribed part of the environment, "time cells" in the hippocampus fired during circumscribed periods of the delay period (Pastalkova et al., 2008; Gill et al., 2011; Kraus et al., 2013). Each cell in $\tilde{f}(\overset{*}{\tau})$ should respond to conjunctive information about what stimulus occurred a particular time in the past. MacDonald et al. (2011) observed some neurons that responded at a particular delay and that distinguished the object that was experienced before the delay, demonstrating that this stimulus-specificity is possible. However, most neurons they recorded from responded on both types of trials, consistent with the hypothesis that they were coding for the stimulus (opening of the door to the waiting enclosure) that initiated the delay. To examine in detail the temporal profile of hippocampal time cells and compare them to

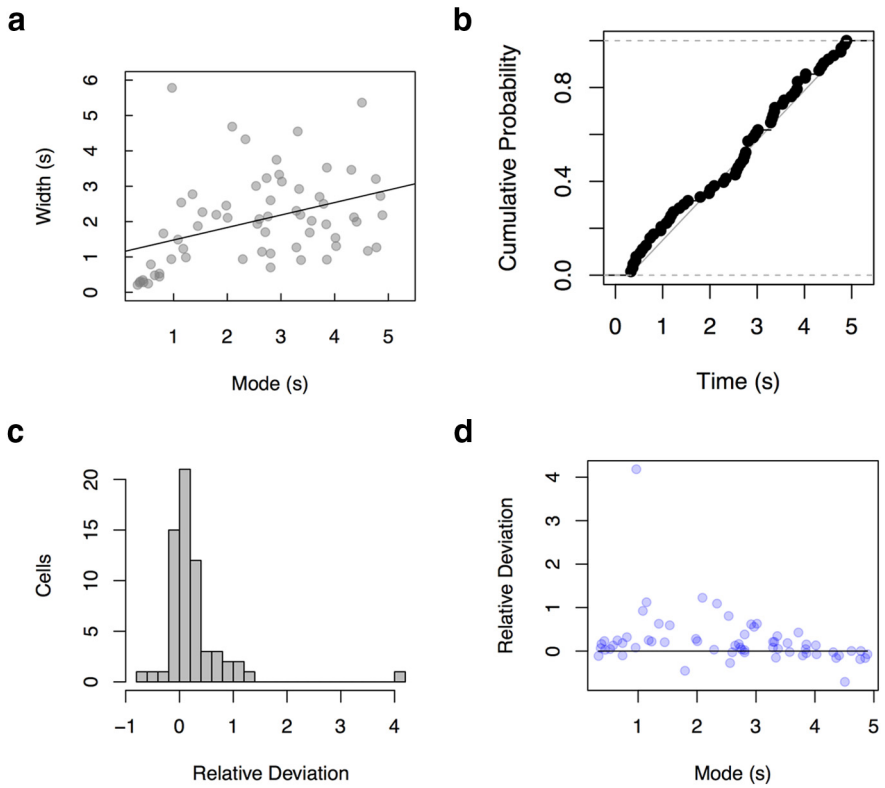

Figure 8. Detailed properties of experimentally observed time fields. ***a***, Width of time fields grow with their mode. See text for details. ***b***, The distribution of modes was not different from uniform. Only cells with modes at least 300 ms from both the start of the delay interval and the shortest end of the delay interval across rats were included. The straight line gives the cumulative probability of the uniform distribution. ***c***, ***d***, Asymmetry in time fields. ***c***, An asymmetry index was calculated for each cell. The distribution is reliably different from zero. ***d***, The asymmetry index was not driven by edge effects, with a residue of skewed time fields for cells with a wide range of modes >0.3 s and <5 s. Data are from secondary analyses of MacDonald et al. (2011).

predictions from $\tilde{f}(\tilde{\tau}^{\star})$ we collapsed the results across both trials types.

Inspection of Figure 4 reveals two properties of $\tilde{f}$ that can be compared with these empirical data. Consider the temporal profile of two cells in $\tilde{f}(\tilde{\tau}^{\star},\tau)$ with different values of $\tilde{\tau}^{\star}$. The cell that codes for the presence of the stimulus a short time in the past peaks a shorter time after presentation of the stimulus than does the cell that codes for the stimulus a longer time in the past. The first cell also shows less spread in its firing field than the second cell. This increasing spread is an extremely important property for $\tilde{f}(\tilde{\tau}^{\star},\tau)$ in behavioral applications, accounting for the decrease in temporal accuracy with delay, an extremely general phenomenon (Yntema and Trask, 1963; Hinrichs and Buschke, 1968; Hacker, 1980). Second, note that the shape of the response for $\tilde{f}(\tilde{\tau}^{\star},\tau)$ is asymmetric in time, showing a positive tail. Secondary analyses of cells reported in MacDonald et al. (2011) show that hippocampal time cells exhibit both of these qualitative properties. Figure 7*a* shows empirical data giving the firing profile of two representative time cells that exhibit both the increased spread and asymmetry. The first cell shown peaks a shorter time into the delay interval than the second. The width of the first cell's firing rate profile is less than the width of the second cell's firing rate profile. Each cell also shows an asymmetric firing rate profile, with a longer tail than a leading edge. Both of these informal observations were also supported by quantitative analyses across the population.

Analyses showed that the cells' spread in firing increased with the delay to their peak firing. A linear regression of each cell's spread to its peak showed a significant regression coefficient: $0.35 \pm 0.1, p < 0.002, R^2 = 0.15$ (Fig. 8*a*). The temporal profile of time cells was also positively skewed. We took the distance between the first time bin where the smoothed firing rate exceeded half of the maximum value and the cell's mode as a measure of the leading edge. We took the distance between the mode and the last time bin where the smoothed density exceeded half of the maximum as a measure of the trailing edge. The trailing edge was greater than the leading edge for 45 of 63 cells included in the analysis, a proportion that was significantly different from chance; $\chi^2(1) = 10.7, p < 0.002$ (Fig. 8*c*). This effect was not attributable to an edge effect that resulted from cutting off the delay interval: even restricting attention to the cells with modes 1000 ms from the beginning and end of the interval, we found that 31/39 were positively skewed, a proportion significantly greater than chance; $\chi^2(1) = 12.4, p < 0.001$ (Fig. 8*d*).

To further characterize the temporal profile of time cells, we computed ensemble similarity for the set of experimentally observed time cells with well defined peaks between all pairs of times during the delay period. Before calculating the cosine of the population vector, each cell's smoothed profile was normalized by its maximum. We then computed the cosine between the population vectors corresponding to each of the 512 smoothed time bins and reported the results over the interval from 300 ms to 5 s. Figure 7*b* shows the cosine of the angle between smoothed population vectors excluding 300 ms from the beginning and end of the delay interval. The bins along the diagonal are necessarily one. Notably, the similarity spreads out such that the representation changes more slowly later in the delay period than it does earlier in the delay period. The distribution of modes across cells was uniform (Fig. 8*b*), suggesting that this finding is attributable to the shape of time fields.

Figure 7*c* shows an analogous plot of the ensemble similarity constructed for $\tilde{f}(\tilde{x}^{\star})$. $\tilde{f}(\tilde{x}^{\star})$ Also changes more slowly as the stimulus, here the start of the delay interval, recedes into the past. This spread with the delay is potentially very important. It parallels results from human and animal studies that show the precision of temporal judgments decreases with the length of time being judged (for review, see Wearden and Lejeune, 2008).

To summarize, we compared the properties of time cells recorded from dorsal CA1 to those predicted by the representation $\tilde{f}(\tilde{x}^{\star})$. We were able to confirm two qualitative predictions. First, time cells in CA1 showed firing that spread out in time for cells that fired later in the delay period. Second, the temporal profile of time cells was asymmetric.

Case II: history-dependent spatial representation

It is not always important to remember the times between events. In many circumstances, an ordinal representation that retains

sequential information about the stimuli leading up to the present is preferable to a representation that includes temporal information. In many cases, cells in the hippocampus and entorhinal cortex respond as if they were coding for an ordinal representation of the sequence of movements taken leading up to the present moment in time. In particular, cells in the hippocampus and entorhinal cortex show history-dependent firing in mazes where the paths the animal can take are constrained. Such cells fire in different locations that correspond to similar trajectories, a phenomenon referred to as path-equivalent firing (Frank et al., 2000; Singer et al., 2010). Some neurons also fire differently in the same location depending on the path the animal has taken to arrive at that point, a phenomenon referred to as retrospective coding (Frank et al., 2000; Wood et al., 2000).

The mathematical framework developed here constructs an ordinal representation of the sequence of movements leading to the present location if $\alpha(\tau)$ is set to the animal's speed $|\mathbf{v}(\tau)|$ and $f(\tau)$ is given by the degree to which the current head direction overlaps with a specific angle θ, $f(\tau) = f_\theta(\tau)$. We restricted $f_\theta(\tau)$ to only have positive values, approximating the response of head direction cells that are observed in various brain regions associated with the hippocampus (Taube, 1998; Solstad et al., 2008). Note that, like the speedometer reading on a car, $\alpha(\tau)$ is always greater than or equal to zero but is not constant over time. We can introduce a variable $p(\tau)$ to describe the total distance traveled along the path up to time τ. Then, the speed $|\mathbf{v}(\tau)| = \frac{dp}{d\tau}$. Because $\alpha(\tau) = \frac{dp}{d\tau}$, $F(s, \tau)$ is the Laplace transform of $f_\theta(\tau)$ with respect to p. As a consequence, we write $\tilde{f}(\overset{\star}{p})$ for the approximate inverse Laplace transform. In much the same way that $\tilde{f}(\overset{\star}{\tau})$ provided a representation of the times at which the stimulus was experienced leading up to the present moment in Case I, here the inverse Laplace transform $\tilde{f}(\overset{\star}{p})$ describes the degree to which the path leading up to the present includes movements in the direction θ; with many cells taking input from many different head directions all of the cells in $\tilde{f}(\overset{\star}{p})$ reconstruct the entire sequence of movements.

Figure 9 illustrates how path-equivalent firing works on the W-maze. The trajectories are taken from Frank et al. (2000), where the animal navigates along a W-maze visiting each of the three arms in sequence (L-C-R-C-L …). There are four types of journeys along the maze (Fig. 9, top). On eastward journeys, L-C and C-R, the animal starts heading north, then heads east until it reaches the next arm, then heads south (left column). On westward journeys, R-C and C-L, the animal starts heading north, then heads *west* until it reaches the next arm, then heads south (right column). Consider an input given by a head direction cell with a preferred direction pointing toward the east (Fig. 9*a*). Here, $f_\theta(\tau)$ would tend to be non-zero a certain distance along the L-C and C-R journeys as the animal heads east (thin lines, left column), but would tend to be close to zero (except for some side-to-side fluctuation) on R-C and C-L journeys (right column). Now consider the activity of cells in $F(s)$ driven by such an input. If the rate constant is extremely fast, a cell in $F(s)$ would be very hard to distinguish from a speed-modulated head direction cell. However, with relatively slow rate constants, the activity of a cell in $F(s)$ would gradually grow exponentially as it is presented with non-zero input, then gradually decay exponentially after the input is no longer present (Fig. 9*a*, medium lines). Now consider the activity of the corresponding cells in $\tilde{f}(\overset{\star}{p})$. In just the same way that the cells in Figure 4 responded a certain time after the input in the temporal domain, here the cells in $\tilde{f}(\overset{\star}{p})$ respond to the input after the animal has traveled some characteristic distance past the point at which the input was received (thick lines in Fig. 9). Naturally, if the input is similar on two journeys, the cells in $F(s)$ and $\tilde{f}(\overset{\star}{p})$ will show similar firing along those to journeys.

Consider the history leading up to a position on the central arm on L-C paths and R-C paths. Under these circumstances, although the animal is in the same location heading in the same direction, the history of movements leading to the present location differs dramatically. This property enables cells in $F(s)$ and $\tilde{f}(\overset{\star}{p})$ to also exhibit retrospective coding on the W-maze. Figure 10*a*,*b* shows results from simulated cells coding for $F(s)$ and $\tilde{f}(\overset{\star}{p})$ on L-C and R-C journeys. The simulated cells in Figure 10*a*,*b* both receive input from the velocity component in the eastward direction. The cell in $F(s)$ (Fig. 10*a*) was chosen to have a slow rate constant. As a consequence, the elevated firing from the eastern movement decays exponentially over a long distance such that it spreads down the central arm. The cell in $\tilde{f}(\overset{\star}{p})$ (Fig. 10*b*) also distinguishes these journeys. However, it does not begin firing at an elevated rate until the animal has already made the turn to enter the center arm. This cell peaks in its firing well along the center arm. Both of these patterns of activity have been observed in retrospective coding cells (Frank et al., 2000).

Figure 10*c*,*d* shows retrospective coding on the continuous alternating T-maze (Wood et al., 2000; Lee et al., 2006; Lipton et al., 2007). On this maze, the animal is rewarded for turning left and then right after successive journeys down the central arm of the maze. That is, the animal travels repeatedly in a figure-eight pattern. On the alternating T-maze, journeys down the central arm can be part of two different segments of the larger path. Figure 10*c* shows a simulated cell in $F(s)$ receiving input from the velocity component toward the east; Figure 10*d* shows the corresponding cell in $\tilde{f}(\overset{\star}{p})$. The cell in $F(s)$ fires just before entering the central arm on journeys from the western return arm but not on journeys from the eastern return arm. As with the W-maze, cells in $F(s)$ with slower rate constants s would show retrospective coding toward the bottom of the central arm (data not shown for clarity). The cell in $\tilde{f}(\overset{\star}{p})$ shows firing that peaks in the middle of the central arm. Neurons in the hippocampus and entorhinal cortex have been observed to show differential activation on the center arm in the continuous alternating T-maze (Wood et al., 2000; Lee et al., 2006; Lipton et al., 2007). On the T-maze, journeys down the center arm differ in both the movements that precede and follow the segment on the center arm. As a consequence, differential activity on the center arm of the T-maze could be attributed to either retrospective coding or prospective coding. There is good evidence that purely prospective coding exists, at least under some circumstances (Ferbinteanu and Shapiro, 2003; Bower et al., 2005; Ainge et al., 2007a,b). Recent evidence, however, suggests that prospective activity enters the hippocampus from another region (Catanese et al., 2012), perhaps the prefrontal cortex (Rich and Shapiro, 2009). Prospective coding could be incorporated into the current framework if information about goals provides part of the input for some cells.

To summarize, we showed that in Case II Equation 1 computes the Laplace transform of movements with respect to distance traveled along a path. We found that cells participating in $F(s)$ and $\tilde{f}(\overset{\star}{p})$ exhibited trajectory coding, or path-equivalent firing, and retrospective coding. These properties allowed the

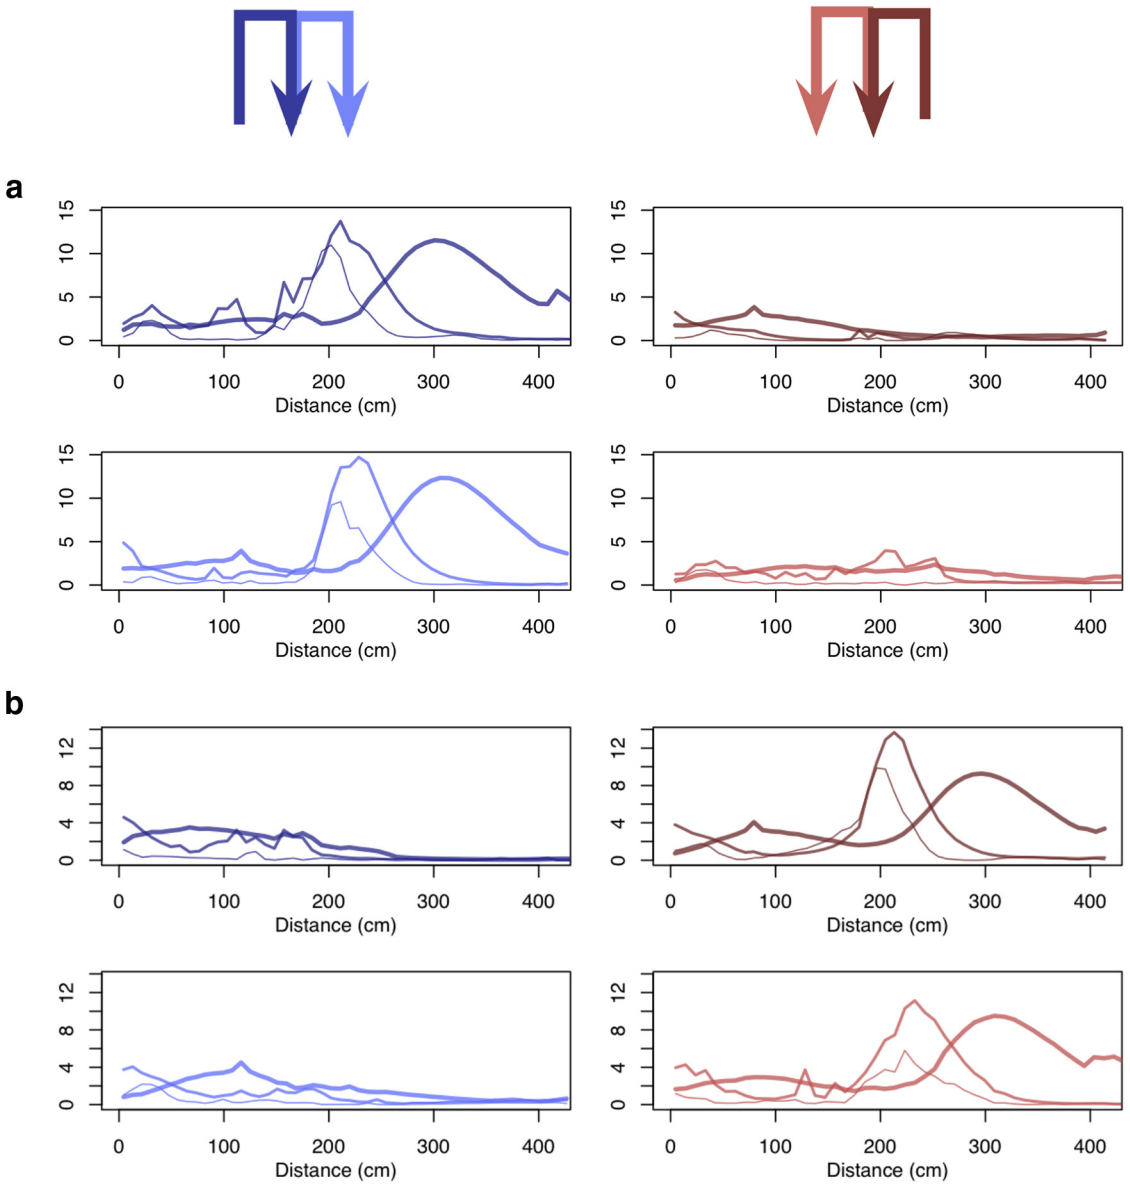

Figure 9. Simulation of path equivalent firing on the W-maze. The animal traverses the W-maze by sequentially visiting the left arm, and then the center arm, then the right arm, then back to the center arm and then back to the left arm, L-C-R-C-L. Each circuit is separated into four paths, indicated by the schematic at the top. ***a***, ***b***, The activity of simulated cells for each of the four paths. The color of each of these four plots correspond to the four paths in the schematic figures; L-C and C-R paths are in the left column, R-C and C-L paths are in the right column. For each set of cells the plots in the bottom row shows the second of the two journeys. That is, the bottom left plot gives C-R and the bottom right plot gives C-L. In each plot, the *x*-axis gives the total distance traveled since the beginning of each journey. Note that this is not the animal's location projected onto an average path, but the total distance traveled; there is some trial-to-trial variability in the location corresponding to a particular distance. In each plot, thin lines are the average value of *f* across trials; medium lines are the average value of a cell in $F(s)$; thick lines are the average value of a cell in $\tilde{f}(\tilde{p}^{*})$. In each plot, the cell in F and in $\tilde{f}$ correspond to the same value of s. ***a***, Cells that take input from a head direction cell with a preferred direction pointing toward the east. Note that the input is activated on journeys where the animal moves to the east (left column) but not on journeys where the animal moves to the west (right column). When the input is non-zero, the cell participating in $F(s)$ rises to a peak and then persists, decaying exponentially. The cell in $\tilde{f}(\tilde{p}^{*})$ follows the input by a characteristic distance along the path. On the western journeys (right), there is little or no input so neither of the simulated cells are activated. ***b***, A set of cells that take input from a head direction cell with preferred direction that points west.

model to account for firing correlates of neurons that have been observed in the hippocampus and entorhinal cortex during navigation along the W-maze and the continuous alternating T-maze. Notably, there are cells in the rodent hippocampal region that exhibit qualitative properties consistent with both $F(s)$ and $\tilde{f}(\tilde{p}^{*})$.

Case III: purely spatial representation

Observing that place fields are systematically distorted by changes in the shape of the environment, O'Keefe and Burgess (1996) proposed that the boundaries of the environment set the reference frames for the hippocampal place code. They hypothesized the existence of boundary vector cells that fire a certain distance and heading from a wall. Conjunctive input from several boundary vector cells anchored to different walls in an environment would result in selective firing at a specific location within the environment; i.e., a place cell. Boundary vector cells predict systematic changes in place field location across changes in the shape of an environment (Hartley et al., 2000; Barry et al., 2006)) and have since been directly observed

in the subiculum (Barry et al., 2006; Lever et al., 2009), a subregion of the hippocampus. In the temporal domain $\tilde{f}(s,\tau)$ keeps track of the time since a stimulus was encountered; when the flow of time is modulated by velocity along some heading, $\tilde{f}(s,x)$ keeps track of the net distance traveled along a heading since a stimulus, here a fixed environmental landmark, was encountered.

Let the animal's trajectory through the two-dimensional environment be given by $x(\tau)$. Let the input function $f(\tau)$ be a function of position, $f[x(\tau)]$. We assume that $f(x)$ is non-zero only when the animal is within some small distance of a particular spatially fixed stimulus. In the open field simulations described here, we set $f[x(\tau)]$ to be non-zero when the animal is within a certain distance of the boundary of the enclosure and zero elsewhere (see Materials and Methods for details). As in the one-dimensional case, in the two-dimensional case each cell is modulated by signed velocity in a single direction x. Generalizing from the simple one-dimensional situation described above (Fig. 6) to the more general case of a two-dimensional open field introduces a few complexities. First, the walls of an enclosure are extended in space. If the landmark is perpendicular to the direction along which velocity is coded, then the two-dimensional case reduces to the one-dimensional case. However, if the landmark is not perpendicular to the direction of motion that controls $\alpha(\tau)$, then the input is not a simple function of position in the x direction. Second, in two-dimensional environments, it is possible to have x take on negative values. To see how this is possible, suppose that there is a point landmark in the center of an open environment and a cell in $F(s)$ is modulated by velocity along the west-east axis. After encountering the landmark, the animal moves to the east. During this movement, the velocity is positive. After traveling some distance to the east, the animal moves a short distance to the north. During this movement, the velocity in the east-west direction is zero and the cell's firing does not change. Now, when returning to the west, the velocity is negative. When the animal reaches the same east-west position as the landmark, the firing of the cell has returned to the value right after encountering the landmark. Now, because the environment is two-dimensional, the animal can continue to the west. Velocity remains negative and the firing of the cell continues to grow exponentially. This is a problem from both a neural perspective, firing rates in the brain are bounded, but also from a mathematical perspective. Under these circumstances, the reconstruction of the Laplace transform no longer makes sense with positive values of s.

To address this second limitation we took the intermediate representation $F(s)$ to be computed from pairs of precursor cells. As describe in detail in the methods, for each direction, we took a pair of cells obeying Equation 1. Each of the cells in the pair took its $\alpha(\tau)$ to be given by a speed-modulated head direction cell with a preferred direction; the preferred directions of the two cells pointed in opposite directions. Neither of the intermediate cells ever received a negative $\alpha(\tau)$; when one of the cells is experiencing a positive $\alpha(\tau)$ the other cell is necessarily experiencing $\alpha(\tau) = 0$. $F(s)$ was computed by taking the ratio of cell 1 over cell 2 if cell 2's firing is greater than cell 1; $F(s)$ was set to zero when cell 2's firing was less than cell 1's. In this way, $F(s)$ enables reconstruction of positive values of position along an axis oriented along cell 1's preferred direction. $\tilde{f}(s,x)$ was computed exactly as in the one-dimensional case as $\mathbf{L}_k^{-1}$ operating on $F(s)$.

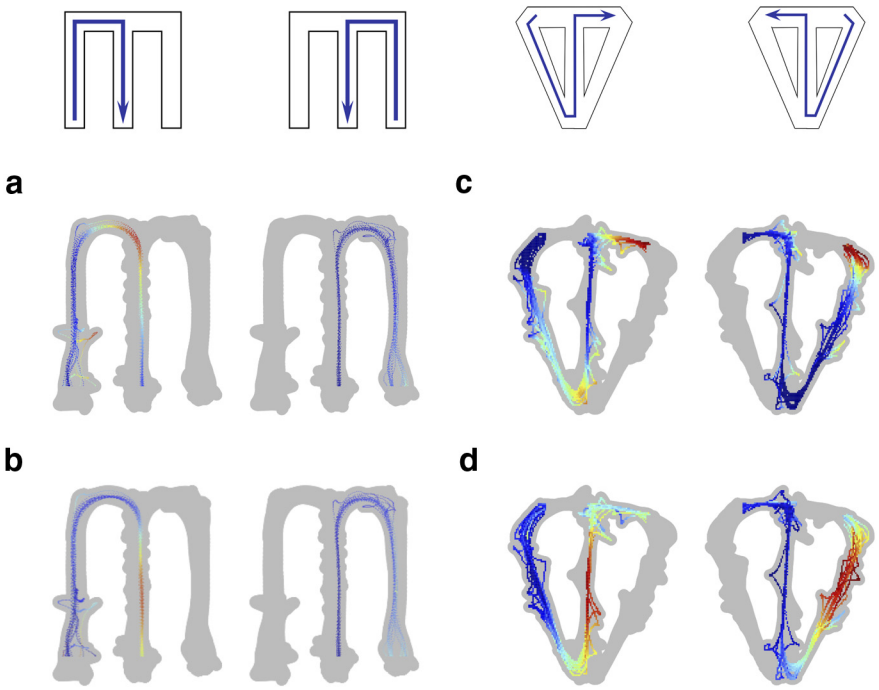

Figure 10. Path-dependent firing on the W-maze and continuous T-maze. Top, Schematic of two journeys that can be taken on each of the two mazes. Left, Retrospective coding on the W-maze (Frank et al., 2000). Right, History-dependent spatial firing on the continuous alternating T-maze (Wood et al., 2000). All cells take their input $f(\tau)$ to be a head direction cell with preferred direction toward the east; all cells have $\alpha(\tau)$ set to be the animal's speed. ***a***, ***c***, Firing rate map of a simulated cell participating in $F(s)$. ***b***, ***d***, Simulated cell coding for $\tilde{f}(p^*)$ corresponding to the cell in $F(s)$. ***a***, Note that elevated firing persists into the central arm on journeys from the west but not from the east. ***b***, This simulated cell also shows retrospective coding that distinguishes the history on the central arm, but now the firing peaks at the center of the arm. ***c***, Note that this simulated cell fires before entering the central stem when coming from the west return arm, but not when coming from the east return arm. ***d***, This simulated cell fires on the central arm on journeys coming from the west but not on journeys from the east. Paths in ***a*** and ***b*** were taken from Frank et al. (2000). Paths in ***c*** and ***d*** were taken from Lipton et al. (2007).

Figure 11 summarizes the results of the two-dimensional simulations. We chose the preferred direction for cell 1 to be due south and for cell 2 to be due north. $f(x)$ was non-zero when the animal was within 3 cm of wall whose perpendicular vector had a component pointing south. For walls that were not exactly east–west, the input was taken to fall off with the cosine between the perpendicular vector and the southerly direction. The left-most column of Figure 11 gives the results for $F(s)$ (Fig. 11, top) and for $\tilde{f}(s,x)$ (Fig. 11, bottom) in a square environment aligned perfectly along the preferred direction. In this case, the model estimates the Laplace transform of distance from the Northern wall. The firing rate of cells participating in $F(s)$ is a decreasing exponential function of distance from the wall. This property is similar to border cells observed in subiculum and other regions (Solstad et al., 2006; Savelli et al., 2008; Lever et al., 2009; Boccara et al., 2010). Simulated cells participating in $\tilde{f}(s,x)$ fire in a strip aligned to the wall. Note that there is some variability in the spatial firing even

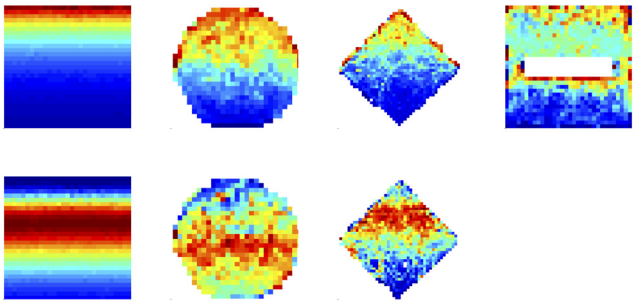

Figure 11. Simulated boundary vector cells in two-dimensional environments. For the simulated cells, velocity was taken as the component of velocity in the southern direction. The input function $f(x)$ was set to be non-zero when the animal was within 3 cm from a northern wall of the enclosure. Spatial trajectories were actual paths taken by rats for the various environments (see methods for details). Top, The Laplace transform $F(s)$ mimics border cells. Bottom, The inverse Laplace transform $\tilde{f}(\overset{*}{x})$ mimics boundary vector cells. The simulations correspond well with experimental results from Lever et al. (2009).

though the simulation is in principle deterministic. That is because this simulation, like all path integration methods, is subject to cumulative error in the estimate of velocity.

In the more general two-dimensional case, $F(s)$ is not precisely the Laplace transform of north-south position. Nonetheless, cells in F and $\tilde{f}(\overset{\star}{x})$ (we will continue to write $\tilde{f}(\overset{\star}{x})$ under these circumstances) correspond reasonably well to firing correlates of subicular neurons under these circumstances. The remaining plots in Figure 11 show firing rate maps for the same simulated cells in environments of different geometries. Across environments, the firing of the simulated cell in $F(s)$ (Fig. 11, top) and the simulated cell in $\tilde{f}(\overset{\star}{x})$ (Fig. 11, bottom) resemble the qualitative pattern observed for border cells and boundary vector cells, respectively (Barry et al., 2006; Lever et al., 2009).

To summarize, we demonstrated that the mathematical framework used here can provide a reconstruction of position along one spatial dimension rather than time. Under these circumstances, cells in $F(s)$ resemble border cells whereas cells in $\tilde{f}(\overset{\star}{x})$ resemble boundary vector cells. Even under circumstances where the algorithm cannot implement the Laplace transform with respect to position, the cells generated by the model nonetheless provide a reasonable approximation to firing correlates of actual neurons.

Conjunctive representation of spatiotemporal context

It has been hypothesized that one of the functions of the hippocampus is to maintain a representation of spatiotemporal context that incorporates all of the aspects of an episode (Cohen and Eichenbaum, 1993; Eichenbaum et al., 2007). According to this account, the hippocampus receives input from a variety of regions that code for distinct classes of stimuli and integrates them into a common representational framework that can capture the relationships between them. Consistent with this view, conjunctive cells that respond preferentially to the simultaneous presence of multiple features within an environment are extremely common (Shapiro et al., 1997; Wood et al., 1999; Anderson and Jeffery, 2003; Komorowski et al., 2009, 2013). Up to now, we have considered three cases in which cells receive input $f(\tau)$ and modulation $\alpha(\tau)$ with varying properties. These "pure" inputs can be combined into conjunctive representations that capture the relationships between the various stimuli. By taking combinations of unimodal functions for a wide variety of stimuli, the hippocampus can build a holistic representation that captures the relationships between stimuli separated in time and space (Staresina and Davachi, 2009). Conjunctive representations are a high-dimensional representation of the different aspects of an episode (Rigotti et al., 2013).

We consider two simple examples of conjunctive representation and compare the results to the firing correlates of cells in dorsal CA1. In Case III, each set of cells coded for the distance along some heading to a particular landmark. In most real world situations, there are multiple landmarks available at any moment. In the first demonstration, we consider the conjunction of boundary vector cells coding for the different walls of an enclosure. We find that the resulting simulated cells correspond well to canonical place fields and respond appropriately to changing the dimensions of an enclosure (O'Keefe and Burgess, 1996; Hartley et al., 2000). In the second example, we take a conjunction of a place cell, itself constructed as the conjunction of boundary vector cells from Case III and a time cell from Case I. The result is a cell with a place field that changes slowly over time (Hyman et al., 2012; Mankin et al., 2012).

It has long been hypothesized that the existence of well defined place cells in dorsal CA1 is a consequence of conjunctions of inputs from boundary vector cells that each code for distance from an environmental boundary (O'Keefe and Burgess, 1996; Hartley et al., 2000). If place cells in dorsal CA1 are caused by boundary vector cells that take as input contact with the boundaries of an enclosure, then changing the shape of the should result in a change in some place fields. Figure 12*a*,*b* shows two cases in which a simulated place cell was constructed by taking combinations of inputs from boundary vector cells like those shown in Figure 11. The firing rate of the simulated cell in Figure 12*a* is given by the product of the firing rate of two BVCs, one "attached" to the northern wall, the other attached to the western wall. The firing rate of the simulated cell in Figure 12*b* is caused by a slightly more complicated combination of the input of three simulated boundary vector cells. As described by previous authors (O'Keefe and Burgess, 1996; Hartley et al., 2000; Barry et al., 2006), conjunctive inputs from a variety of boundary vector cells can account for a range of findings from the place cell literature.

The present mathematical framework can provide for purely spatial cells (Case III) as well as purely temporal cells (Case I). By taking conjunctions of cells from these two classes we can account for place cells with place fields that change gradually over time (Manns et al., 2007; MacDonald et al., 2011; Hyman et al., 2012; Mankin et al., 2012). For instance, many of the time cells observed by MacDonald et al. (2011) that fired for circumscribed time periods during the delay interval of a memory task also showed place fields during the delay. These cells fired only when the animal was in a circumscribed location during a circumscribed time period within the delay. Similarly, Mankin et al. (2012) recorded from the same arena over several sessions. They observed place cells that showed firing fields that changed gradually via rate remapping over hours. Figure 12*c* shows a simulated cell with firing rate computed as the product of two boundary vector cells (Case III) and a time cell (Case I) that takes as input a stimulus available only at the beginning of the session. The conjunctive cell does not show any firing early in the period because the time cell has not yet become active. As the time cell becomes active, the conjunctive cell exhibits a place field in the location where both of the boundary vector cells are active. As time continues onward, the time cell ceases firing and the conjunctive cell again becomes inactive. Combining time cells coding for different aspects of stimulus history in a gradually changing environment, could account for the widespread finding that the place

code of dorsal CA1 ensembles changes gradually with the passage of time (Manns et al., 2007; MacDonald et al., 2011; Hyman et al., 2012; Mankin et al., 2012).

Discussion

Equations 1 and 6 provide a unified framework for constructing functions of time, position in a sequence, or spatial position. The method starts by constructing an intermediate representation of leaky integrators with different time constants. From this intermediate representation, we extract a fuzzy representation of a function of time, or position in a sequence, or spatial position. This framework organizes superficially different neurophysiological findings in the rodent hippocampus and parahippocampal structures. Different pathways receive a variety of inputs $f(\tau)$ and are modulated by different functions $\alpha(\tau)$, but all of the simulated cells considered are generated by the same two equations. With appropriate settings, we obtain a fuzzy representation of what happened when (Case I), or of what happened in which part of the sequence (Case II), or what is located where (Case III). All of these sources of information combine in the hippocampus to provide a conjunctive description of the current spatiotemporal context (Fig. 12).

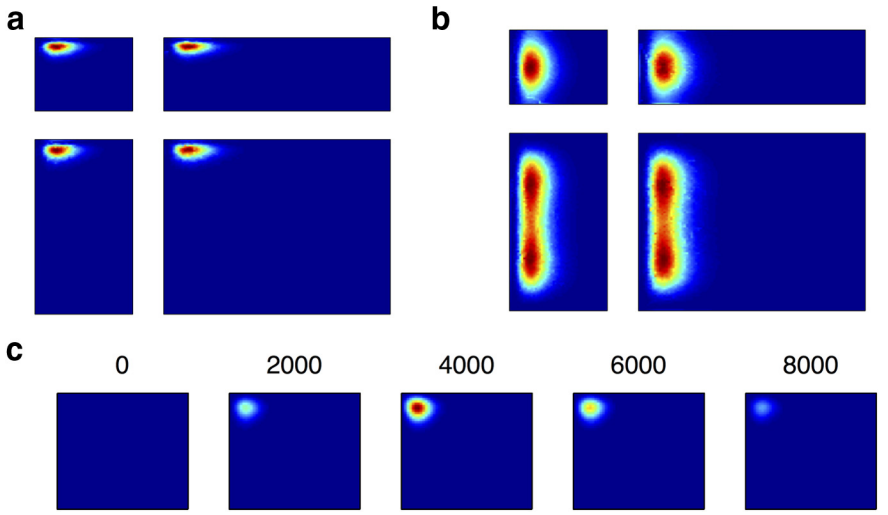

Figure 12. Cells that take conjunctions of simple functions of space and time generate a rich set of firing correlates. ***a***, ***b***, Conjunctions of boundary vector cells generate canonical place cells with fields that move when the dimensions of the environment are changed. Each of the four panels gives a firing rate map for four different environments of different sizes. ***a***, A cell taking the product of input from one boundary vector cell coding for distance to the northern wall and a cell coding for distance from the western wall. ***b***, Input from three boundary vector cells was combined. The first cell peaks in its firing 2 units from the northern wall, the second cell peaks firing 4 units from the western wall, and the third peaks 6 units from the eastern wall. The firing rate of the unit shown is the firing rate of cell 1 times the sum of cells 2 and 3. ***c***, Firing rate map for a simulated conjunctive time cell/place cell for different periods of time. The simulated cell's firing rate is given by the product of two boundary vector cells (Case III) and a time cell (Case I) coding for an event at the beginning of the session. The first firing rate map is for 0 –2000 units of time; the second map is for 2000 – 4000 units of time. Because the temporal history is scale-invariant, the units of time are arbitrary except insofar as the time scale is long enough to allow the path to thoroughly sample positions.

Comparison to other work

In the temporal domain, the present approach extends several existing hypotheses for sequential activation of cells. Although cells are sequentially activated in $\tilde{f}$ following presentation of a stimulus, cells in $\tilde{f}$ do not have direct connections between them, in contrast to chaining models that exploit direct links between sequentially activated cells (Goldman, 2009; Itskov et al., 2011). In the temporal domain, the present mathematical framework is a special case of a liquid state machine (Maass et al., 2002; Buonomano and Maass, 2009), although an extremely simple one that is trivially decodable. The present approach has perhaps the most in common with spectral timing theory (Grossberg and Merrill, 1992). However, the present approach is much more simple facilitating insight into its workings and extension to other variables.

Other work has generalized from temporal to spatial variables. For instance, Hasselmo (2009) presented a framework for learning spatial and temporal trajectories using changes in synaptic weights between grid cells and cells coding velocity. In this paper, sequential firing does not depend on changes in synaptic weights; $\mathbf{L}_k^{-1}$ is fixed throughout an experiment.

Biological requirements for the computation

As discussed above, given a set of leaky integrators $F(s)$, it is straightforward to approximate the inversion of the Laplace transform and construct $\tilde{f}$ using feedforward connections with lateral inhibition. There are several computational requirements necessary to implement Equation 1, but these are at least plausible given current knowledge.

First, we require exponentially decaying cells with a variety of time constants; to describe behavioral effects over long time scales, time constants of perhaps a few thousand seconds would be necessary. Although network effects are a possible mechanism for long time constants, firing that lasts over long periods of time could also result from intrinsic currents. Graded persistent firing (Egorov et al., 2002; Fransén et al., 2006) is similar to an integrator with an infinite time constant. Long-lasting persistent firing has been observed *in vitro* in a variety of regions throughout the hippocampal region and some cells terminate persistent firing after variable time periods (Egorov et al., 2006; Yoshida and Hasselmo, 2009; Navaroli et al., 2011; Sheffield et al., 2011; Hyde and Strowbridge, 2012; Knauer et al., 2013; Jochems and Yoshida, 2013). We recently developed a relatively detailed computational model that adapts cellular properties believed to support persistent firing to implement Equation 1 with long time constants (Z. Tiganj, M. E. Hasselmo, M. W. Howard, unpublished observations).

Second, the time constants of these integrators should be distributed along some anatomical axis to facilitate taking the derivative with respect to s. The dorsoventral axis is a good candidate for this gradient. Place field size changes systematically along the dorsoventral axis (Jung et al., 1994; Kjelstrup et al., 2008). The time constant of the integrators controls the spatial resolution of the representation that results in $\tilde{f}(x^*)$. Cells in MEC exhibit border cell activity consistent with $F(s)$ in Case III (Solstad et al., 2008). Cellular properties in MEC change systematically along the dorsoventral axis. There are systematic differences in resonance properties (Giocomo et al., 2007; Heys et al., 2010), synaptic summation properties (Garden et al., 2008), and the radius of inhibitory synaptic interactions (Beed et al., 2013) along the dorsoventral axis *in vitro*.

Third, Equation 1 requires that $\alpha(\tau)$ multiplicatively modulate the inputs to a leaky integrator and its ongoing activity. On a cellular level, multiplicative interactions could be caused by the

properties of NMDA receptors, in which postsynaptic depolarization multiplicatively enhances the synaptic currents induced by presynaptic glutamate release (Poirazi et al., 2003). Multiplicative interactions have been demonstrated in both experimental data (Otmakhova et al., 2002; Jarsky et al., 2005) and in detailed biophysical simulations (Poirazi et al., 2003). Gain control via noisy synaptic inputs is another potential mechanism to rapidly change the time constant of cells (Chance et al., 2002). Head direction cells (Taube, 1998; Solstad et al., 2008; Brandon et al., 2011) coupled with neural correlates of speed are sufficient to generate a velocity signal (O'Keefe et al., 1998; Maurer et al., 2005; Wills et al., 2012).

Specific experimental predictions

The equations described here make several qualitative and quantitative predictions that could be tested with existing technology. In practice, it is difficult assign a specific function to a specific region. Presumably, neurons in different pathways receive different inputs and different modulators; predicting the firing correlates of a population of cells computing $F(s)$ or $\tilde{f}$ requires knowing their inputs and modulators. If dorsal CA1 contains a conjunctive representation of many inputs then CA1 is a good location to test qualitative predictions of the model but a poor location to test quantitative predictions. Quantitative predictions would be better tested in "pure" populations that provide input to the hippocampus.

Behavioral considerations suggest that the temporal information in $\tilde{f}(\tilde{\tau}^{\star})$ should extend much longer than the time scale observed thus far; perhaps as long as a few thousand seconds (Shankar and Howard, 2012; Howard and Eichenbaum, 2013). Stimulus-specific time cells should exist over similarly long periods of time. Stimulus-specific time cells have been observed (MacDonald et al., 2011, 2013) and time cells have been observed over the longest scales that have been considered (~16 s; Kraus et al., 2013). There is ample evidence for changes in ensemble representations over time scales much longer than 10 s (Manns et al., 2007; Hyman et al., 2012; Howard et al., 2012; Mankin et al., 2012), but it is not known whether those changes carry information about past stimuli. Addressing the question of whether stimulus-specific time cells exist over long time scales requires establishing experimental control over gradually changing firing.

Whereas neurons that resemble the behavior of cells constructing the intermediate representation have been observed in Case II and Case III (Table 1), the major gap in the functional correlates predicted by the framework for these three cases thus far is the absence of the intermediate representation $F(s)$ for Case I. In the temporal domain, cells participating in $F(s)$ would exhibit stimulus specific firing with exponential decay. Different cells would exhibit a variety of time constants. Given the strong spatial correlates in medial entorhinal cortex (MEC) and the absence of clear spatial firing correlates in lateral entorhinal cortex (LEC; Hargreaves et al., 2005), LEC seems like an excellent candidate to search for cells participating in $F(s)$ in Case I. The perirhinal cortex, which provides cortical input to LEC, is another good candidate.

The analogy between temporal and spatial functions proposed here makes several qualitative predictions about the properties of spatial firing correlates. By analogy with time cells, this framework predicts that the spatial extent of cells in $\tilde{f}(\tilde{x}^{\star})$ coding for the distance from a landmark should spread out with the distance to the landmark, as predicted by descriptive models of boundary vector cells (Barry et al., 2006; Hartley et al., 2000). To test this prediction, it is necessary to establish that a landmark controls firing of a particular cell. Experimental control over the firing of place cells by a landmark has been established on a linear track (Gothard et al., 1996, 2001).

The most theoretically important quantitative prediction of the model is mathematical scale-invariance (Shankar and Howard, 2012). Scale-invariance is a central aspect of many models of a variety of memory and timing tasks (Gallistel and Gibbon, 2000; Brown et al., 2007). Scale-invariance manifests neurally as a constant ratio between the width of a cell's time field and the latency with which it responds to the stimulus. This should hold on individual trials and not be due to an averaging artifact across trials. In practice, there are many methodological challenges to evaluating scale-invariance. For instance, uncontrolled variables, including behavior, could inadvertently introduce a scale to the input function.

References

Ainge JA, Tamosiunaite M, Woergoetter F, Dudchenko PA (2007a) Hippocampal CA1 place cells encode intended destination on a maze with multiple choice points. J Neurosci 27:9769–9779. CrossRef Medline

Ainge JA, van der Meer MA, Langston RF, Wood ER (2007b) Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. Hippocampus 17:988–1002. CrossRef Medline

Anderson MI, Jeffery KJ (2003) Heterogeneous modulation of place cell firing by changes in context. J Neurosci 23:8827–8835. Medline

Balsam PD, Gallistel CR (2009) Temporal maps and informativeness in associative learning. Trends Neurosci 32:73–78. CrossRef Medline

Barry C, Lever C, Hayman R, Hartley T, Burton S, O'Keefe J, Jeffery K, Burgess N (2006) The boundary vector cell model of place cell firing and spatial memory. Rev Neurosci 17:71–97. Medline

Beed P, Gundlfinger A, Schneiderbauer S, Song J, Böhm C, Burgalossi A, Brecht M, Vida I, Schmitz D (2013) Inhibitory gradient along the dorsoventral axis in the medial entorhinal cortex. Neuron 79:1197–1207. CrossRef Medline

Boccara CN, Sargolini F, Thoresen VH, Solstad T, Witter MP, Moser EI, Moser MB (2010) Grid cells in pre- and parasubiculum. Nat Neurosci 13:987–994. CrossRef Medline

Bower MR, Euston DR, McNaughton BL (2005) Sequential-context-dependent hippocampal activity is not necessary to learn sequences with repeated elements. J Neurosci 25:1313–1323. CrossRef Medline

Brandon MP, Bogaard AR, Libby CP, Connerney MA, Gupta K, Hasselmo ME (2011) Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. Science 332:595–599. CrossRef Medline

Brown GD, Neath I, Chater N (2007) A temporal ratio model of memory. Psychol Rev 114:539–576. CrossRef Medline

Brunel N, Trullier O (1998) Plasticity of directional place fields in a model of rodent CA3. Hippocampus 8:651–665. CrossRef Medline

Buonomano DV, Maass W (2009) State-dependent computations: spatiotemporal processing in cortical networks. Nat Rev Neurosci 10:113–125. CrossRef Medline

Catanese J, Cerasti E, Zugaro M, Viggiano A, Wiener SI (2012) Dynamics of decision-related activity in hippocampus. Hippocampus 22:1901–1911. CrossRef Medline

Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. Neuron 35:773–782. CrossRef Medline

Chater N, Brown GD (2008) From universal laws of cognition to specific cognitive models. Cogn Sci 32:36–67. CrossRef Medline

Cohen NJ, Eichenbaum H (1993) Memory, amnesia, and the hippocampal system. Cambridge, MA: MIT.

Egorov AV, Hamam BN, Fransén E, Hasselmo ME, Alonso AA (2002) Graded persistent activity in entorhinal cortex neurons. Nature 420:173–178. CrossRef Medline

Egorov AV, Unsicker K, von Bohlen und Halbach O (2006) Muscarinic control of graded persistent activity in lateral amygdala neurons. Eur J Neurosci 24:3183–3194. CrossRef Medline

Eichenbaum H, Yonelinas AP, Ranganath C (2007) The medial temporal lobe and recognition memory. Annu Rev Neurosci 30:123–152. CrossRef Medline

Estes WK (1955) Statistical theory of spontaneous recovery and regression. Psychol Rev 62:145–154. CrossRef Medline
Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. Neuron 40:1227–1239. CrossRef Medline
Frank LM, Brown EN, Wilson M (2000) Trajectory encoding in the hippocampus and entorhinal cortex. Neuron 27:169–178. CrossRef Medline
Fransén E, Tahvildari B, Egorov AV, Hasselmo ME, Alonso AA (2006) Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. Neuron 49:735–746. CrossRef Medline
Gallistel CR, Gibbon J (2000) Time, rate, and conditioning. Psychol Rev 107:289–344. CrossRef Medline
Garden DL, Dodson PD, O'Donnell C, White MD, Nolan MF (2008) Tuning of synaptic integration in the medial entorhinal cortex to the organization of grid cell firing fields. Neuron 60:875–889. CrossRef Medline
Gill PR, Mizumori SJ, Smith DM (2011) Hippocampal episode fields develop with learning. Hippocampus 21:1240–1249. CrossRef Medline
Giocomo LM, Zilli EA, Fransén E, Hasselmo ME (2007) Temporal frequency of subthreshold oscillations scales with entorhinal grid cell field spacing. Science 315:1719–1722. CrossRef Medline
Goldman MS (2009) Memory without feedback in a neural network. Neuron 61:621–634. CrossRef Medline
Gothard KM, Skaggs WE, McNaughton BL (1996) Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. J Neurosci 16:8027–8040. Medline
Gothard KM, Hoffman KL, Battaglia FP, McNaughton BL (2001) Dentate gyrus and ca1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. J Neurosci 21:7284–7292. Medline
Grossberg S, Merrill JW (1992) A neural network model of adaptively timed reinforcement learning and hippocampal dynamics. Brain Res Cogn Brain Res 1:3–38. CrossRef Medline
Hacker MJ (1980) Speed and accuracy of recency judgments for events in short-term memory. J Exp Psychol Hum Learn Mem 6:651–675. CrossRef
Hargreaves EL, Rao G, Lee I, Knierim JJ (2005) Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. Science 308:1792–1794. CrossRef Medline
Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J (2000) Modeling place fields in terms of the cortical inputs to the hippocampus. Hippocampus 10:369–379. CrossRef Medline
Hasselmo ME (2007) Arc length coding by interference of theta frequency oscillations may underlie context-dependent hippocampal unit data and episodic memory function. Learn Mem 14:782–794. CrossRef Medline
Hasselmo ME (2009) A model of episodic memory: mental time travel along encoded trajectories using grid cells. Neurobiol Learn Mem 92:559–573. CrossRef Medline
Heys JG, Giocomo LM, Hasselmo ME (2010) Cholinergic modulation of the resonance properties of stellate cells in layer ii of medial entorhinal cortex. J Neurophysiol 104:258–270. CrossRef Medline
Hinrichs JV, Buschke H (1968) Judgment of recency under steady-state conditions. J Exp Psychol 78:574–579. CrossRef Medline
Howard MW, Eichenbaum H (2013) The hippocampus, time, and memory across scales. J Exp Psychol Gen 142:1211–1230. CrossRef Medline
Howard MW, Viskontas IV, Shankar KH, Fried I (2012) A neural signature of mental time travel in the human MTL. Hippocampus 22:1833–1847. CrossRef Medline
Hyde RA, Strowbridge BW (2012) Mnemonic representations of transient stimuli and temporal sequences in the rodent hippocampus in vitro. Nat Neurosci 15:1430–1438. CrossRef Medline
Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK (2012) Contextual encoding by ensembles of medial prefrontal cortex neurons. Proc Natl Acad Sci U S A 109:5086–5091. CrossRef Medline
Issa JB, Zhang K (2012) Universal conditions for exact path integration in neural systems. Proc Natl Acad Sci U S A 109:6716–6720. CrossRef Medline
Itskov V, Curto C, Pastalkova E, Buzsáki G (2011) Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus. J Neurosci 31:2828–2834. CrossRef Medline
Jarsky T, Roxin A, Kath WL, Spruston N (2005) Conditional dendritic spike propagation following distal synaptic activation of hippocampal CA1 pyramidal neurons. Nat Neurosci 8:1667–1676. CrossRef Medline
Jochems A, Yoshida M (2013) Persistent firing supported by an intrinsic cellular mechanism in hippocampal ca3 pyramidal cells. Eur J Neurosci 38:2250–2259. CrossRef Medline
Jung MW, Wiener SI, McNaughton BL (1994) Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. J Neurosci 14:7347–7356. Medline
Kjelstrup KB, Solstad T, Brun VH, Hafting T, Leutgeb S, Witter MP, Moser EI, Moser MB (2008) Finite scale of spatial representation in the hippocampus. Science 321:140–143. CrossRef Medline
Knauer B, Jochems A, Valero-Aracama MJ, Yoshida M (2013) Long-lasting intrinsic persistent firing in rat ca1 pyramidal cells: a possible mechanism for active maintenance of memory. Hippocampus 23:820–831. CrossRef Medline
Komorowski RW, Manns JR, Eichenbaum H (2009) Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. J Neurosci 29:9918–9929. CrossRef Medline
Komorowski RW, Garcia CG, Wilson A, Hattori S, Howard MW, Eichenbaum H (2013) Ventral hippocampal neurons are shaped by experience to represent behaviorally relevant contexts. J Neurosci 33:8079–8087. CrossRef Medline
Kraus BJ, Robinson RJ 2nd, White JA, Eichenbaum H, Hasselmo ME (2013) Hippocampal "time cells": time versus path integration. Neuron 78:1090–1101. CrossRef Medline
Lee I, Griffin AL, Zilli EA, Eichenbaum H, Hasselmo ME (2006) Gradual translocation of spatial correlates of neuronal firing in the hippocampus toward prospective reward locations. Neuron 51:639–650. CrossRef Medline
Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J (2002) Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature 416:90–94. CrossRef Medline
Lever C, Burton S, Jeewajee A, O'Keefe J, Burgess N (2009) Boundary vector cells in the subiculum of the hippocampal formation. J Neurosci 29:9771–9777. CrossRef Medline
Lipton PA, White JA, Eichenbaum H (2007) Disambiguation of overlapping experiences by neurons in the medial entorhinal cortex. J Neurosci 27:5787–5795. CrossRef Medline
Maass W, Natschläger T, Markram H (2002) Real-time computing without stable states: a new framework for neural computation based on perturbations. Neural Comput 14:2531–2560. CrossRef Medline
MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H (2011) Hippocampal "time cells" bridge the gap in memory for discontiguous events. Neuron 71:737–749. CrossRef Medline
MacDonald CJ, Carrow S, Place R, Eichenbaum H (2013) Distinct hippocampal time cell sequences represent odor memories immobilized rats. J Neurosci 33:14607–14616. CrossRef Medline
Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK (2012) Neuronal code for extended time in the hippocampus. Proc Natl Acad Sci U S A 109:19462–19467. CrossRef Medline
Manns JR, Howard MW, Eichenbaum H (2007) Gradual changes in hippocampal activity support remembering the order of events. Neuron 56:530–540. CrossRef Medline
Marr D, Hildreth E (1980) Theory of edge detection. Proc R Soc Lond B Biol Sci 207:187–217. CrossRef Medline
Maurer AP, Vanrhoads SR, Sutherland GR, Lipa P, McNaughton BL (2005) Self-motion and the origin of differential spatial scaling along the septo-temporal axis of the hippocampus. Hippocampus 15:841–852. CrossRef Medline
Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J Neurosci 7:1951–1968. Medline
Navaroli VL, Zhao Y, Boguszewski P, Brown TH (2012) Muscarinic receptor activation enables persistent firing in pyramidal neurons from superficial layers of dorsal perirhinal cortex. Hippocampus 22:1392–1404. CrossRef Medline
Naya Y, Suzuki WA (2011) Integrating what and when across the primate medial temporal lobe. Science 333:773–776. CrossRef Medline
O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. Nature 381:425–428. CrossRef Medline
O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. Brain Res 34:171–175. CrossRef Medline
O'Keefe J, Burgess N, Donnett JG, Jeffery KJ, Maguire EA (1998) Place cells,

navigational accuracy, and the human hippocampus. Philos Trans R Soc Lond B Biol Sci 353:1333–1340. CrossRef Medline

Otmakhova NA, Otmakhov N, Lisman JE (2002) Pathway-specific properties of AMPA and NMDA-mediated transmission in ca1 hippocampal pyramidal cells. J Neurosci 22:1199–1207. Medline

Pastalkova E, Itskov V, Amarasingham A, Buzsáki G (2008) Internally generated cell assembly sequences in the rat hippocampus. Science 321:1322–1327. CrossRef Medline

Poirazi P, Brannon T, Mel BW (2003) Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. Neuron 37:977–987. CrossRef Medline

Post E (1930) Generalized differentiation. Trans Am Math Soc 32:723–781. CrossRef

Rich EL, Shapiro M (2009) Rat prefrontal cortical neurons selectively code strategy switches. J Neurosci 29:7208–7219. CrossRef Medline

Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. Nature 497:585–590. CrossRef Medline

Savelli F, Yoganarasimha D, Knierim JJ (2008) Influence of boundary removal on the spatial representations of the medial entorhinal cortex. Hippocampus 18:1270–1282. CrossRef Medline

Shankar KH, Howard MW (2012) A scale-invariant representation of time. Neural Comput 24:134–193. CrossRef Medline

Shankar KH, Howard MW (2013) Optimally fuzzy scale-free memory. J Mach Learn Res 14:3753–3780.

Shapiro ML, Tanila H, Eichenbaum H (1997) Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. Hippocampus 7:624–642. CrossRef Medline

Sheffield ME, Best TK, Mensh BD, Kath WL, Spruston N (2011) Slow integration leads to persistent action potential firing in distal axons of coupled interneurons. Nat Neurosci 14:200–207. CrossRef Medline

Singer AC, Karlsson MP, Nathe AR, Carr MF, Frank LM (2010) Experience-dependent development of coordinated hippocampal spatial activity representing the similarity of related locations. J Neurosci 30:11586–11604. CrossRef Medline

Solstad T, Moser EI, Einevoll GT (2006) From grid cells to place cells: a mathematical model. Hippocampus 16:1026–1031. CrossRef Medline

Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI (2008) Representation of geometric borders in the entorhinal cortex. Science 322:1865–1868. CrossRef Medline

Staresina BP, Davachi L (2009) Mind the gap: binding experiences across space and time in the human hippocampus. Neuron 63:267–276. CrossRef Medline

Taube JS (1998) Head direction cells and the neurophysiological basis for a sense of direction. Prog Neurobiol 55:225–256. CrossRef Medline

Tulving E (1983) Elements of episodic memory. New York: Oxford UP.

Wearden JH, Lejeune H (2008) Scalar properties in human timing: conformity and violations. Q J Exp Psychol (Hove) 61:569–587. CrossRef Medline

Wills TJ, Barry C, Cacucci F (2012) The abrupt development of adult-like grid cell firing in the medial entorhinal cortex. Front Neural Circuits 6:21. CrossRef Medline

Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. Science 261:1055–1058. CrossRef Medline

Wood ER, Dudchenko PA, Eichenbaum H (1999) The global record of memory in hippocampal neuronal activity. Nature 397:613–616. CrossRef Medline

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27:623–633. CrossRef Medline

Yntema DB, Trask FP (1963) Recall as a search process. J Verb Learn Verb Behav 2:65–74. CrossRef

Yoshida M, Hasselmo ME (2009) Persistent firing supported by an intrinsic cellular mechanism in a component of the head direction system. J Neurosci 29:4945–4952. CrossRef Medline