Place from time: Reconstructing position from the temporal context model

Marc W. Howard and **Vaidehi S. Natu**
Department of Psychology Syracuse University

Abstract

The temporal context model (TCM Howard & Kahana, 2002) was proposed to describe recency and associative effects observed in episodic recall. Episodic recall depends on an intact medial temporal lobe, a region of the brain that also supports a place code. Howard, Fotedar, Datey, and Hasselmo (2005) demonstrated that the leaky integrator that supports a gradually-changing representation of temporal context in TCM is sufficient to describe properties of cells observed in ventromedial entorhinal cortex during spatial navigation if it is provided with input about the animal's current velocity. This representation of temporal context generates noisy place cells in the open field, unlike the clearly-defined place cells observed in the hippocampus. Here we demonstrate that a reasonably accurate spatial representation can be extracted from temporal context with as few as eight cells, suggesting that the spatial precision observed in the place code in the hippocampus is not inconsistent with input from a representation of temporal-spatial context in entorhinal cortex.

It has long been believed that the medial temporal lobe (MTL), including the hippocampus and parahippocampal cortical regions, is involved in episodic recall—memory for specific instances from ones' life (e.g. Nadel & Moscovitch, 1997). It has also been clear from neurophysiological results that the MTL is involved in maintaining a representation of position within an open environment (e.g. O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978; Quirk, Muller, Kubie, & Ranck, 1992; Fyhn, Molden, Witter, Moser, & Moser, 2004; Wilson & McNaughton, 1993). It is of considerable theoretical interest to know if these two functions —episodic memory and a current representation of position—correspond to a single computational function (e.g. O'Keefe & Nadel, 1978; Eichenbaum, Dudchenko, Wood, Shapiro, & Tanila, 1999). Recently Howard et al. (2005) proposed that the temporal context model (TCM, Howard & Kahana, 2002) could provide a framework in which to describe aspects of both episodic recall and the construction and maintenance of a place code. We will review this work briefly here before presenting new results on the reconstruction of spatial position from a representation of temporal context.

TCM in episodic recall

TCM (Howard & Kahana, 2002; Howard, 2004; Howard, Wingfield, & Kahana, Revised; Howard et al., 2005) provides a set of equations that describe a distributed representation of temporal context as a vector in a high-dimensional space. This model was developed to describe performance in situations in which a list of words are presented for a memory test. In TCM, the state of context at time step i, $\mathbf{t}_i$ is updated according to

$$\mathbf{t}_i = \rho_i \mathbf{t}_{i-1} + \beta \mathbf{t}_i^{IN}, \quad \rho_i : \left|\left| \mathbf{t}_i \right|\right| = 1, \tag{1}$$

where β is a free parameter that controls the rate of change of temporal context. The value of the scalar ρ_i is chosen at each time step to keep the context vector of unit length. The context

vector is formed from the context vector from the previous time step, $\mathbf{t}_{i-1}$ and an input vector $\boldsymbol{t}_i^{IN}$. We assume throughout that the length of $\boldsymbol{t}_i^{IN}$ is less than or equal to one. The input vector is caused by the currently presented stimulus. The properties of $\boldsymbol{t}_i^{IN}$ will be described extensively below.

It may help to attempt to visualize the process of contextual evolution described by Eq. 1. The pattern of activity $\mathbf{t}_i$ may be thought of as a pattern of sustained firing across cells in extrahippocampal MTL regions (Howard et al., 2005;Egorov, Hamam, Fransén, Hasselmo, & Alonso, 2002). The constraint that the length of $\mathbf{t}_i$ is always equal to 1 means that $\mathbf{t}_i$ lives on the surface of a hypersphere in a high-dimensional space. At each time step, the input pattern pushes $\mathbf{t}_{i-1}$ in a particular direction. Then ρ_i is chosen to bring the resulting vector sum back to the surface of the sphere, resulting in $\mathbf{t}_i$. If we assume, as we typically have in modeling random lists of words each presented once, that the input patterns $\boldsymbol{t}_i^{IN}$ are orthonormal to each other, then Eq. 1 describes a random walk and the similarity of context vectors falls off exponentially with the number of presentations between them:

$$\boldsymbol{t}_i \cdot \boldsymbol{t}_j = \rho^{|i-j|}, \tag{2}$$

where ρ without the subscript is the asymptotic value of ρ_i when presented with an infinitely long list of orthonormal inputs, $\rho = \sqrt{1 - \beta^2}$.

In TCM, the current state of the context vector is used to probe recall of a set of vectors corresponding to the items that are to be recalled, typically words from episodic recall tasks. This is accomplished by means of an outer product matrix $\mathbf{M}^{TF}$ which is updated by

$$\boldsymbol{M}_i^{TF} = \boldsymbol{M}_{i-1}^{TF} + \boldsymbol{f}_i \boldsymbol{t}_i', \tag{3}$$

where $\mathbf{f}_i$ is the pattern corresponding to the item vector presented at time step i and the prime denotes the transpose. Although a mapping hypothesis between TCM as a formal model and the structure of the MTL will be developed more fully later (see also Howard et al., 2005), we note here that the entries in the matrix $\mathbf{M}^{TF}$ may be roughly interpreted as the strength of the set of synapses connecting extrahippocampal MTL regions, especially entorhinal cortex, to neocortical association areas. When $\mathbf{M}^{TF}$ is multiplied from the right with a particular state of context $\mathbf{t}_j$, this yields a superposition of item vectors, each weighted by the inner product of the probe context to that item's encoding context(s):

$$\boldsymbol{M}^{FT} \boldsymbol{t}_j = \sum_i \boldsymbol{f}_i (\boldsymbol{t}_i \cdot \boldsymbol{t}_j). \tag{4}$$

Each item is cued by a particular context cue to the extent that context cue overlaps with the study context of that item. In modeling episodic recall tasks, we have used a non-linear competitive recall rule to map this superposition of item vectors onto a probability of recall for each item.

The recency effect

In the free recall task, the subject is presented with a list of items one at a time. At recall the subjects' task is to recall as many items as possible without regard to order. Subjects can initiate recall after a delay (e.g. Glanzer & Cunitz, 1966), or even from specific prior lists (Shiffrin, 1970; Ward & Tan, 2004), so that recall cannot be initiated by using item representations available in short-term memory. Accordingly, models of free recall have typically assumed

that subjects maintain some representation of the list context to use as the cue in free recall (e.g. Raaijmakers & Shiffrin, 1980; Anderson & Bower, 1972; Hasselmo & Wyble, 1997). TCM builds on this tradition by assuming that the current state of temporal context is the cue used to initiate recall in the free recall task. Because $\mathbf{t}_i$ changes gradually over time (Eqs. 1, 2), and items are cued to the extent that their study context overlaps with the probe context, this naturally leads to a recency effect; items presented at the end of the list should be more strongly activated, and hence more likely to be recalled than items from earlier in the list.

This explanation of the recency effect contrasts to some extent with accounts of recency that depend on the presence of items in a limited capacity short-term store (Atkinson & Shiffrin, 1968; Raaijmakers & Shiffrin, 1980; Davelaar, Goshen-Gottstein, Ashkenazi, & Usher, 2005). On the one hand $\mathbf{t}_i$ is a record of recent experience, like notions of short-term store. On the other, $\mathbf{t}_i$ changes gradually over time, rather than in a discrete fashion like buffer models (Atkinson & Shiffrin, 1968; Raaijmakers & Shiffrin, 1980). This gradual decay, coupled with a competitive retrieval rule enabled TCM to provide an explanation of the long-term recency effect (Bjork & Whitten, 1974; Glenberg et al., 1980; Howard & Kahana, 1999) observed in continuous-distractor free recall.

Associative effects

Traditional accounts of interitem associations describe associations between, say, members of a pair as some form of direct connection between item representations (e.g. Raaijmakers & Shiffrin, 1980; Murdock, 1982). In TCM, associations between list items are mediated by the effect those items have on context. When A is used as a probe item, there is no direct connection between A and B to support an association. Rather, A causes a change in the state of context, which then overlaps with the context that B was presented in, resulting in a behaviorally-observed association. Items have their effect on context by determining the input patterns $\mathbf{t}_i^{IN}$ in Eq. 1. Suppose we present an item A at some time A_i and then repeat it at some later time A_{i+1}. The input pattern that A evokes when it is repeated is given by:

$$\mathbf{t}_{A_{i+1}}^{IN} = a_O \mathbf{t}_{A_i}^{IN} + a_N \mathbf{t}_{A_i}.$$ (5)

The pattern $\mathbf{t}_{A_i}$ is the state of temporal context when A was presented previously. The pattern $\mathbf{t}_{A_i}^{IN}$ is the input pattern caused by A when it was presented previously. This input pattern $\mathbf{t}_{A_i}^{IN}$, is also governed by Eq. 5, meaning that it in turn, will also include prior states of context, $\mathbf{t}_{A_{i-1}}^{IN}$ in this case. The coefficients α_O and α_N are chosen such that the input pattern when A is repeated, $\mathbf{t}_{A_{i+1}}^{IN}$, is unity. A free parameter γ is used to control the relative contribution of $\mathbf{t}_{A_i}^{IN}$ and $\mathbf{t}_{A_i}$ in contstructing the retrieved context pattern.

The repetition of an item results in an input to the current state of context that includes prior components of contextual states. Because these two components overlap with the encoding contexts of other items, the current state of context after an item repetition serves as an effective retrieval cue for items that were presented near in time to the previous presentation of the repeated item. This provides a basis for explaining temporally-defined associations (Kahana, 1996; Howard & Kahana, 1999). In particular, $\mathbf{t}_{A_i}^{IN}$ is part of contextual states that followed the previous presentation of A, but not those that preceded it (see Eq. 1), so that the combination

of $t_{A_i}^{IN}$ and $\mathbf{t}_{A_i}$ together provide an asymmetric retrieval cue that mimics the form of temporally-defined associations measured from episodic recall tasks (Howard & Kahana, 2002; Howard et al., 2005, Revised).

A mapping hypothesis onto the MTL

Performance in episodic recall tasks depends on an intact MTL (e.g. Graf, Squire, & Mandler, 1984). If TCM is an realistic model of episodic recall tasks, then it should be possible to map the components of TCM onto MTL structures and use this mapping to explain neurophysiological and neuropsychological results. One possiblity would be to put $\mathbf{t}_i$ in the hippocampus. Hasselmo and Wyble (1997) placed context in the hippocampus proper in a previous model of free recall. In that model, free recall proceeded by using hippocampal context as a probe for cortical items; in item recognition cortical items were used as a probe to try and recover hippocampal context (see also Dennis & Humphreys, 2001; Schwartz, Howard, Jing, & Kahana, in press). Similarly, a model of sequence learning that has been applied to a number of learning tasks (Levy, 1996; Shon, Wu, Sullivan, & Levy, 2002; Wu & Levy, 1998, 2001) argues that one function of the hippocampus is to support local context states that depend on the item presented and the temporal context it is presented in. These local context neurons bridge across temporally proximate item presentations, providing a possible explanation of the temporally-defined associations described by TCM (Kahana, 1996; Howard & Kahana, 1999). In contrast to these prior approaches, various considerations led Howard et al. (2005) to hypothesize that $\mathbf{t}_i$ resides in the entorhinal cortex, and perhaps other parahippocampal cortical areas as well, rather than in the hippocampus proper. The hippocampus proper was hypothesized to support new item-to-context learning. That is, Howard et al. (2005) hypothesized that the hippocampus was responsible for maintaining a non-zero value of α_N in Eq. 5. This mapping has been shown to provide an explanation of neuropsychological and neurophysiological findings from relational memory tasks (Bunsey & Eichenbaum, 1996; Higuchi & Miyashita, 1996) and neurophysiological findings from spatial navigation tasks (Frank, Brown, & Wilson, 2000; Quirk et al., 1992). We briefly review these findings and their explanations within TCM.

The hippocampus and transitive association

In TCM, repetition of an item A results in two components of contextual input (Eq. 5). One component, $\mathbf{t}_{A_i}$, is the state of context present when the item was last presented. The other component, $t_{A_i}^{IN}$, is the input that the item caused when it was last presented. Consider for a moment what would happen if only one of those components was available. If $\alpha_N = 1$ and $\alpha_O = 0$, meaning that only $\mathbf{t}_{A_i}$ contributed, then this would constitute a symmetric retrieval cue for items presented near in time to A (see Eq. 2). In contrast, if $\alpha_N = 0$ and $\alpha_O = 1$, meaning that only $t_{A_i}^{IN}$ contributed, then this would constitute an asymmetric retrieval cue. In particular, if $\alpha_N = 0$, then this would result in robust forward associations, but no temporally-defined backward associations. Bunsey and Eichenbaum (1996) observed that hippocampal lesions impaired the development of backward associations, while leaving forward associations intact. This finding led Howard et al. (2005) to hypothesize that the hippocampus functions to enable items to reconstruct the temporal contexts in which they are presented, resulting in a non-zero value of α_N.

This hypothesis about the function of the hippocampus in reconstructing temporal context states also provides an explanation of transitive associations. If a subject is trained on a pair of items A–B, and then learns the pair B–C at some other time, a transitive association results if A is associated to C, despite the fact that A and C were never presented together. Bunsey and

Eichenbaum (1996) observed that although hippocampal lesions had no effect on learning the forward pairwise associations A–B and B–C, they selectively disrupted transitive associations between A and C. Howard et al. (2005) showed that setting $\alpha_N = 0$, simulating a disruption in the ability for items to reconstruct the temporal contexts in which they were presented, selectively disrupted transitive associations. In TCM, reconstruction of temporal contexts give rise to transitive associations because these temporal context states are caused by inputs from the items themselves. That is, when B is presented with C, B recovers the temporal context associated with its presentation with A. The state of context in which C is encoded is therefore similar to the state of contextual input that will be caused by A as a probe. Howard et al. (2005) showed that reconstruction of temporal context states enables the development of a cortical stimulus representation that reflects higher-order temporal relationships among the stimuli; TCM can be seen as a quantitative implementation of ideas about the role of the hippocampus and entorhinal cortex in sequence encoding and relational memory (e.g. Eichenbaum, 2004).

A pseudo-integrator in entorhinal cortex

What computations might give rise to a representation of position like that observed in the hippocampus? One possibility is a representation derived from landmark information. The hippocampal place code persists unchanged in the dark (Quirk, Muller, & Kubie, 1990) and even in blind rats (Save, Cressant, Thinus-Blanc, & Poucet, 1998), suggesting that the hippocampal place code does not depend critically on visual input. One way to construct a representation of position is to use a dead reckoning strategy. The first integral of velocity is position. If provided with information about the current velocity, a computation that integrates its inputs would result in a representation of position. The temporal context vector $\mathbf{t}_i$ resembles an integrator. Equation 1 states that $\mathbf{t}_i$ includes $\mathbf{t}_{i-1}$. In turn, $\mathbf{t}_{i-1}$ includes $\mathbf{t}_{i-1}^{IN}$ and $\mathbf{t}_{i-2}$:

$$
\begin{aligned}
\mathbf{t}_i &= \rho_i \mathbf{t}_{i-1} + \beta \mathbf{t}_i^{IN} \\
&= \rho_i \rho_{i-1} \mathbf{t}_{i-2} + \beta (\mathbf{t}_i^{IN} + \rho_i \mathbf{t}_{i-1}^{IN}) \\
&= \beta (\mathbf{t}_i^{IN} + \rho_i \mathbf{t}_{i-1}^{IN} + \rho_i \rho_{i-1} \mathbf{t}_{i-2}^{IN} + \; \dots \;)
\end{aligned}
\tag{6}
$$

By "unwinding" in this way we see that $\mathbf{t}_i$ is the result of something like an integration of the $\mathbf{t}^{IN}$s. This expression also shows that $\mathbf{t}_i$ is not the result of a perfect integration; rather the factors of ρ_i cause a decay of information over time. The temporal context vector is the result of a leaky integration of a series of input patterns.

Howard et al. (2005) examined the properties of the place code that would result from combining the leaky integrator implemented by Eq. 1 with inputs that consist of velocity vectors:

$$
\mathbf{t}_i^{IN} = \mathbf{v}_i
\tag{7}
$$

First, we note that this assumption could be implemented with known properties of the entorhinal cortex. Most notably, cells in layer V of the entorhinal cortex show stable graded firing when stimulated (Egorov et al., 2002). The presence of persistent stable firing is essential to implement Eq. 1. To see this clearly, set $\mathbf{t}_i^{IN} = \mathbf{0}$ in Eq. 1. This results in $\rho_i = 1$ and we find that $\mathbf{t}_i = \mathbf{t}_{i-1}$; to implement Eq. 1 cells must persist in their firing rates in the absence of input. Remarkably, this property is met by principal cells in layer V *in vitro* (Egorov et al., 2002). Moreover, to implement Eq. 1, cells must be able to assume a new stable firing rate when presented with a depolarizing input—they must be able to integrate. This property is also met

by layer V cells, even in the absence of recurrent synaptic inputs (Egorov et al., 2002). In summary, it appears that cells in layer V have precisely the cellular properties necessary to implement Eq. 1. The mechanisms that give rise to these properties at the cellular level are of considerable interest. Several proposals have been made (e.g. Fransen, Egorov, Hasselmo, & Alonso, 2003; Loewenstein & Sompolinsky, 2003)

Having cells that are capable of implementing the integration property of Eq. 1 is necessary but not sufficient to implement a place code using a leaky integrator. Two other mechanisms are necessary—a way to normalize the length of $\mathbf{t}_i$ and a way to provide velocity vectors as inputs to the set of integrator cells. Chance, Abbott, and Reyes (2002) showed that cultured cortical neurons showed a gain that decreased as the simulated background activity of the network increased. If the intrinsic currents that give rise to persistent firing are subject to this type of gain control, then a set of integrator cells should eventually find a level of stable overall firing rate, implementing something like the normalization represented by ρ_i in Eq. 1. To provide velocity vectors as input (Eq. 1) we need information about heading and speed. Heading is provided by head direction cells (Taube, 1998); it is known that cells in layer V of the entorhinal cortex receive input from regions containing head direction cells (Haeften, Wouterlood, & Witter, 2000). The only other requirement is a way to represent speed on these inputs. There are a couple of ways to implement this. Here we note that average firing rate in hippocampal pyramidal cells is strictly proportional to running speed (Figure 8, Zhang, Ginzburg, McNaughton, & Sejnowski, 1998), suggesting that the MTL manages to find a solution to this problem.

Having motivated the proposal that the place code in entorhinal cortex implements $\mathbf{t}_i$ with velocity vectors as inputs, Howard et al. (2005) examined the correlates of cells that compose $\mathbf{t}_i$ when presented with paths from actual rats during navigation tasks. We first examined place fields generated by simulated neurons during navigation around an open field. Figure 1*a* and *b* show representative place fields from the open simulation. The cells showed reliable but noisy modulation by location. The location of the place fields was consistent across different environments. Both of these properties have been observed for entorhinal cells during exploration of the open field (Quirk et al., 1992; but see Fyhn et al., 2004). The simulated cells showed place fields located towards the preferred direction of the head direction cell that projected to that place cell. For instance, the cell in Figure 1*a* had a preferred direction that pointed toward the east.

Although the simulated cells showed positional modulation in the open field similar to that observed for cells in the entorhinal cortex, it is something of a misnomer to refer to these as place cells. The integrator described in Eq. 6 is leaky, meaning that more recent movements are more strongly represented than movements that took place a longer time ago. Rather than a place code, $\mathbf{t}_i$ is more accurately described as a weighted sum over recent movements. Rather than being a shortcoming, this property turns out to be essential for describing properties of the enorhinal "place code" when the animal is navigating around the W-maze. Frank et al. (2000) had animals move around a three-armed maze (Figure 1*c*) from arm to arm. The animals were trained to visit a food well at the bottom of the center arm, then one at the bottom of the left arm, then return to the center arm, then the right arm, back to the center and so forth. Frank et al. (2000) observed cells in entorhinal cortex and the hippocampus proper that responded to specific sequences of movements rather than locations. A *trajectory coding* cell would, for instance, respond on the bar of the maze when the animal was making left-center or center-right journeys, but not when in the same location on other trips. The simulation of Eq. 1 on the W-maze showed huge numbers of trajectory coding cells with properties like this. Frank et al. (2000) also observed cells that showed *retrospective coding*. A unique property of the W-maze is that it allows one to compare trips down the center arm according to which arm the animal is coming from (Figure 1*c*). On these trips, the animals position, heading and behavioral goal

(the chocolate milk at the end of the arm) are the same. However, a group of cells that code for the sequence of recent movements (Eq. 6) would be expected to show differential firing according to where the animal is coming from. Frank et al. (2000) observed retrospective coding cells in the entorhinal cortex, including both deep and superficial layers, as well as the CA1 field of the hippocampus. The simulation showed cells with this property (Figure 1*d*). These cells were activated when the animal assumed the cell's preferred direction on the bar of the maze. Because firing persists over a macroscopic period of time, the same property that enabled us to describe the recency effect in free recall, the cell does not turn off right away, but shows firing that persists along the center arm of the maze. One aspect of the Frank et al. (2000) results that the model did not capture robustly was the phenomenon of *prospective coding*, in which cells fire differentially according to movements the animal is about to make (see also Ferbinteanu & Shapiro, 2003;Wood, Dudchenko, Robitsek, & Eichenbaum, 2000, for more evidence of retrospective and prospective coding in the hippocampus).

Could the hippocampal place code come from t_i?

Howard et al. (2005) demonstrated that $\mathbf{t}_i$, when provided with velocity vectors as inputs, can describe spatial correlates of cells from ventromedial entorhinal cortex (Frank et al., 2000; Quirk et al., 1992, see also Fyhn et al., 2004 for recent findings on the dorsolateral entorhinal cortex). Does the entorhinal cortex function like $\mathbf{t}_i$ during spatial navigation? The fact that $\mathbf{t}_i$ is not really a spatial code at all seems at first glance like a serious obstacle to accepting this point of view. As discussed above, $\mathbf{t}_i$, when provided with velocity vectors as inputs, is better interpreted as a weighted sum over recent movements than as a positional code. This seems inconsistent with what we know about the hippocampus in the open field.

Hippocampal place fields in the open field are omnidirectional (Muller, Bostock, Taube, & Kubie, 1994), suggesting an accurate representation of position. Recording from over a hundred cells, Wilson and McNaughton (1993) observed reconstruction errors of as small as 7 cm in a familiar 62×62 cm open field environment using a template matching method. This discrepancy between actual and reconstructed position was on the order of the 5 cm tracking error intrinsic to their set up. It is possible to obtain even more precise reconstruction when the motion is constrained, as in a linear track (Jensen & Lisman, 2000) or a figure eight maze (Zhang et al., 1998). If the hippocampal place code results from processing $\mathbf{t}_i$ driven by velocity vectors, then there must be a way to extract accurate positional information from a weighted sum over recent movements. If it proves impossible to reconstruct position to the precision possible with hippocampal cells (Wilson & McNaughton, 1993; Zhang et al., 1998; Jensen & Lisman, 2000) using $\mathbf{t}_i$, then this rules out $\mathbf{t}_i$ as an explanation of entorhinal activity, at least in spatial applications. To address these questions, we will attempt to perform spatial reconstruction on the output of the cellular implementation of $\mathbf{t}_i$ driven by velocity vectors in the open field.

Simulation

Methods

We studied the ability to reconstruct position from the output of a cellular simulation implementing $\mathbf{t}_i$ with velocity vectors as inputs. The methods of the cellular implemention of $\mathbf{t}_i$ follow those used in Howard et al. (2005). The cellular implementation of $\mathbf{t}_i$ was driven by velocity vectors from simulated movements. Reconstruction was performed on the output of the cellular implementation. To the extent that we can reconstruct position, we can conclude that a brain region, presumably the hippocampus, receiving input from $\mathbf{t}_i$ driven by velocity vectors could in principle construct an accurate place code in the open field.

Cellular simulation—Following Howard et al. (2005) we implemented $\mathbf{t}_i$ with a set of simulated cells. This is illustrated schematically in Figure 2. The activity of cell i at time step s was computed using

$$t_i(s) = \rho(s)[t_i(s-1) + \beta t_i^{IN}(s)], \tag{8}$$

where β is a free parameter, $\rho(s)$ implements cortical gain control and $t_i^{IN}(s)$ is cell i's input from the head direction system implementing Eq. 7. Notice that the notation used here differs from that used previously for the vectors. Whereas the subscript in the vector $\mathbf{t}_i$ refers to the time step, the subscripts in Eq. 8 refer to the cells with the time step s provided as the argument. Equation 8 differs from Eq. 1 substantively in that the factor of ρ multiplies the input term as well as the previous state of context term. For sufficiently small time steps, this is unlikely to be an important difference.

Cortical gain control (Chance et al., 2002; Chance & Abbott, 2000) was implemented according to

$$\rho(s) = \left\{ \sum_i [t_i(s-1)]^2 \right\}^{-1/2}. \tag{9}$$

This form of $\rho(s)$ has the effect of normalizing $\mathbf{t}$ after each time step (see Eq. 8). To construct the velocity input (Eq. 7) we first simulated the output of a head direction system (e.g. Taube, 1998). Each cell had a preferred direction φ_i. These were spaced evenly around the circle starting with $\varphi_1 = 0$. At each time step s, we compute $\varphi_i^{\text{diff}}(s)$, the absolute difference between each cells preferred direction and the animal's actual head direction $\varphi(s)$:

$$\varphi_i^{\text{diff}}(s) := \begin{cases} |\varphi(s) - \varphi_i|, & |\varphi(s) - \varphi_i| < \pi \\ 2\pi - |\varphi(s) - \varphi_i| & |\varphi(s) - \varphi_i| \geq \pi \end{cases}. \tag{10}$$

Then the input to each cell at time step s is computed using

$$t_i^{IN}(s) = \left| \left| \mathbf{p}(s+1) - \mathbf{p}(s) \right| \right| \frac{1}{\sigma\sqrt{2\pi}} \exp\frac{-[\varphi_i^{\text{diff}}(s)]^2}{2\sigma^2}, \tag{11}$$

where $\mathbf{p}(s)$ is the position at time step s and $\|\mathbf{p}(s+1)-\mathbf{p}(s)\|$ is just the animal's speed at time s. The parameter σ is the width of the tuning curve of the head direction cells. To be roughly consistent with observed values (Taube, 1998), we fixed σ to $\pi/6$ for the simulations in this paper.

Simulated paths: Simulated paths were constructed using an algorithm introduced by Brunel and Trullier (1998). At each time step the animal moved one unit of distance. The direction of movement θ changed from time-step to time-step using Equation 12:

$$\tau_\theta \dot{\theta} = -\theta + \hat{\theta} + \sigma_\theta \sqrt{\tau_\theta} \eta(t), \tag{12}$$

where the dot denotes differentiation with respect to time, $\hat{\theta}$ is the direction to the target, τ_θ is a time constant that controls the animal's turning radius, $\eta(t)$ is a white noise source and σ_θ controls the magnitude of the noise.

In the open field simulations (Figure 4), the animal moved around an 80×80 cm simulated environment, with $\tau_\theta = 2$ and $\sigma_\theta = 0.5$. To simulate random foraging, a set of ten "food locations"

was chosen from a uniform distribution. The closest location to the animal's current position was chosen as the first goal location. After the animal navigated to within a radius of one cm to the goal location the location was removed from the list of available food locations and the closest remaining food location to the animal's current position was chosen as the new goal location. When all ten food locations had been visited, a new set of ten locations was chosen.

Reconstruction—Several existing techniques for estimating position from the activity of cells in the hippocampus would be ineffective for estimating position from the cellular simulation described here. Methods, like the template matching methods used by Wilson and McNaughton (1993), that depend on estimating the mean firing rate of each cell as a function of position are unlikely to succeed given the highly variable nature of the simulated cells in the open field (see Figure 1). This is because t_i is really responding to the set of recent movements and a position may be reached from numerous different trajectories. This means that a particular position may be associated with a wide variety of states of context. The activity of any individual cell is a poor predictor of position taken alone. Bayesian methods of reconstruction (Zhang et al., 1998; Jensen & Lisman, 2000) are poorly suited for this model because the firing rate of cells are highly dependent on each other. This means that the probability of a particular vector state must be calculated from the observed joint probability at each positional bin. Typically joint probabilities are estimated by treating the firing of the cells as independent variables (e.g. Jensen & Lisman, 2000). The need to directly estimate joint probability becomes combinatorically very costly. In order to reconstruct position, we need to have some insight into how the high-dimensional t_i vectors project onto two-dimensional space.

The following thought experiment will motivate the reconstruction method we used. Consider the situation in which we have four cells fed by head direction cells with preferred directions pointing in the cardinal directions (ENWS). Now assume that the animal moves around a one-dimensional square track clockwise with constant speed. After a sufficient number of circuits, it should be possible to reconstruct position precisely because each position along the linear track is reached by only one trajectory. As the animal moves to the East, along the southern edge of the track, the firing rate of "easterly" cell, t_E will gradually increase, much like the voltage of a charging capacitor. As the animal turns left and begins to head North, t_E will decay exponentially and t_N will begin to increase. Each cell's firing rate will increase as the animal moves along the corresponding edge of the track, then decay exponentially as it starts moving in a different direction. As long as the time constant of the increase is not too short, we should be able to read off the animal's position simply by looking at the appropriate cell and the appropriate neighbor. For instance, t_E assumes a particular intermediate value while the animal is moving to the east and the cell is "charging" and when the animal is moving to the north while t_E is decaying. Examination of t_S and t_N is necessary to disambiguate these two possibilities. Now consider the case in which the animal moves along an east-west line. In this case, t_E should increase its firing as the animal moves to the east and decay as the animal moves to the west. It is necessary to examine both t_E and t_W, the firing rate of the westerly cell, to determine position along the east-west axis, but each east-west position corresponds to one (t_E, t_W) pair. What happens if the animal suddenly stops moving east-west and starts moving north-south? t_E and t_W will both start decaying exponentially. Although both t_E and t_W are changing, their ratio remains unchanged as the animal moves north-south. We can apply similar logic to construct and maintain position information using arbitrary sequences of movements along a grid.

Figure 3 illustrates the validity of this idea. We constructed paths between corners of a box. This sequence of paths is more complex than a one-dimensional path but captures the essential complexities of the open field problem. We simulated t_i with four cells given the sequence of movements.[1] We then looked at the two-dimensional "shadow" of the four-dimensional t_i

when we projected it onto ($\log t_E/t_W$, $\log t_N/t_S$) space. As can be seen by comparing corresponding panels in Figure 3, the projection of the four-dimensional $\mathbf{t}_i$ onto this two-dimensional space corresponds roughly to the topology of the actual paths.

The correspondence shown in Figure 3 motivated us to undertake a more elaborate reconstruction. We generated simulated paths with 100,000 movements, then generated a series of $\mathbf{t}_i$s for each movement with a variety of values of β. We estimated reconstructed x and y coordinates using

$$x_r = a \sum_{i=1}^{N} \cos\varphi_i \log(t_i) \tag{13}$$

$$y_r = a \sum_{i=1}^{N} \sin\varphi_i \log(t_i), \tag{14}$$

where N is the number of cells (set to 8 here to cover the circle) and φ_i is the preferred direction of cell i. The preferred direction of the first cell $\varphi_1 =$. The preferred directions were evenly spaced over the circle. The slope a was estimated from the data. This was done by picking 10,000 time steps randomly (excluding the first 1000 points to avoid edge effects), calculating the sums on the rhs of Eqs. 13 and 14, concatenating these values, and doing a simple linear regression to the observed x and y coordinates from the sampled points. Our estimate of a was simply the slope of this regression. The intercept of the regression was typically approximately zero. Because by symmetry the intercept should be precisely zero we fixed it at zero and omitted it from the reconstruction.

This reconstruction method is just population vector decoding. It differs slightly from traditional population decoding approaches (e.g., Abbott, 1994) in that the basis vectors are not derived from a tuning curve relating firing to position, but rather the preferred direction of the head direction signal each cell receives. Also, taking the natural log of the firing rate linearizes the exponential character of the firing rates.

Results

After finding a from a sample of the points, the reconstructed position was calculated for all points in the sample. Figure 4 displays the results of the reconstruction. On the left is the average distance between the reconstructed position and the actual position as a function of β. With β as low as .01, the reconstruction error is about 7 cm. With β at .001, the reconstruction error decreased to about 2.2 cm.[2] The right side of the figure shows representative paths of actual and reconstructed positions taken from successive 500 time-step segments of the reconstructed path. Qualitatively, the reconstructed paths show an excellent correspondence with the observed data.

General Discussion

We started with the question of whether an entorhinal place code derived from $\mathbf{t}_i$, a representation of temporal context from a model of episodic recall, retains sufficient information about position to support the hippocampal place code. We used a simple population vector reconstruction scheme in which each cell voted for its preferred direction with the natural

---

[1]Because the diagonal movements corresponded to a region where the tuning curves for both of the adjacent cells were low, we renormalized the input vector, such that the length of $\boldsymbol{t}_i^{IN}$ was one at each time step for the purposes of this illustration.

[2]With even lower values of β, *e.g.* 10^{-4} the observed error began to increase again. This was probably due to either rounding errors and/or a very slow change from the initial state leading to an edge effect.

log of its firing rate. Using this method we showed that $\mathbf{t}_i$ implemented with as few as eight cells was able to provide excellent positional reconstruction. Reconstruction error was a function of β, which controls the rate of forgetting in the network. Good reconstruction was observed with small values of β. This makes sense in that small values of β correspond to a slower rate of forgetting, which enables movements from farther in the past to contribute to the estimate of position. These findings suggests that it is possible that the excellent spatial precision observed in the hippocampus in the open field (Wilson & McNaughton, 1993) could be the consequence of an appropriate calculation performed on an entorhinal representation described by $\mathbf{t}_i$ provided with velocity vectors as input.

There are a number of potential limitations of the applicability of the present theoretical excercise. The simple linear reconstruction method used here will not work well for larger numbers of cells; more sophisticated regression methods that take into account the correlational structure among the firing rates of the cells would be necessary to provide adequate reconstruction if more cells were used. It is also possible that the simulated movements did not capture some important property of the actual movements a rat might take in the open field. For instance, actual paths could foil attempts at reconstruction if they are sufficiently pathological. If for some reason the rat decided to run in one of several small circular paths for extremely long periods of time, this method would be able to reconstruct position along each path, but not distinguish the location of the circular paths themselves. This is because the information that discriminates a circular path in, say, the northwest quadrant of an environment from one in the southeast would be the movements prior to initiation of the circular path. As the animal runs on the circle this information would recede into the distance. This property of the model is perhaps not disadvantageous. When an animal in the open field starts a sequence of stereotyped movements, the firing of hippocampal place fields changes (Markus, Qin, Leonard, et al., 1995). Further, if the movement of a rat is constrained to quadrants of an environment by means of barriers, the firing of hippocampal place cells is similar across the different quadrants (Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002). Finally we note that the population vector method used here relies on an accurate measurement of low firing rates. Although integrator cells (Egorov et al., 2002) can maintain stable firing rates over a wide variety of values, their capacity is finite. With sufficient hyperpolarization, they eventually shut off. This thresholding behavior would constrain the positional accuracy that can be maintained with $\mathbf{t}_i$.

Variation in the speed of movement was not included in the simulations used here. To illustrate the dependence of reconstruction on running speed, we will briefly derive an expression for the change in our log measure in a simplified case. Consider a situation in which the animal moves to the east with speed $v(s)$ units per time step at time step s. Let's assume that we have four cells with preferred directions in the cardinal directions and sufficiently small tuning curves that we can neglect the input activation of the cells other than the one with the preferred direction pointing East. Let's consider the change in activity of the "East" cell, t_E and the "West" cell, t_W as a result of a movement of $v(s)$ units to the east. Using Eq. 8, the activities of these cells at time step $s+1$ is given by:

$$t_E(s+1) = \rho(s)[t_E(s) + \beta v(s)]$$

$$t_W(s+1) = \rho(s)t_W(s).$$

We're interested in the dependence of the change in $\log(t_E/t_W)$ on the speed of the movement at time step s, $v(s)$. We start by deriving the change in $\log(t_E)$.

$$\Delta \, \log \, t_E(s) = \log\{\rho(s)[t_E(s) + \beta v(s)]\} - \log[t_E(s)]$$

$$= \log\left[\rho(s) + \frac{\rho(s)\beta v(s)}{t_E(s)}\right]$$

$$= \log \rho(s) + \log\left(1 + \frac{\beta v(s)}{t_E(s)}\right).$$

Following similar logic, the change in $\log(t_W)$ is given by:

$$\Delta \, \log \, t_W(s) = \log[\rho(s) t_W(s)] - \log[t_W(s)]$$

$$= \log \rho(s).$$

Using these two expressions, we find

$$\Delta \, \log \, [t_E(s) \,/\, t_W(s)] = \log\left[1 + \frac{\beta v(s)}{t_E(s)}\right] \tag{15}$$

Of course $\log(1 + x) \simeq x$ when x is sufficiently small. When $\beta/t_E(s)$ is sufficiently small, log (t_E/t_W) responds almost linearly to changes in running speed. Insofar as this is the case, log (t_E/t_W) transforms exactly as one would desire from a measure of position. The requirement that $\beta/t_E(s)$ be small for this reconstruction scheme to work also provides insight into the regime over which this reconstruction scheme will work. Setting β to a small value not only enables $\mathbf{t}_i$ to retain information from a longer time, but also facilitates the constraint that $\beta/t_E(s)$ be small. The dependence of this factor on $t_E(s)$ is somewhat unsettling, as this value will change over time as, say, the animal continues to move to the east. This variability undoubtedly accounts for some if not all of the residual error observed in the simulations reported here. However, when β is small, the values of the components of $\mathbf{t}_i$ are restricted to a relatively narrow dynamic range. Under those conditions, if the variability in $v(s)$ is large with respect to the variability in $t_E(s)$, then an appropriate sensitivity to speed of movement will be preserved.

The diversity of spatial firing properties in entorhinal cortex

The model of the entorhinal place code examined here (Howard et al., 2005) predicts the properties of trajectory and retrospective coding on the W-maze, consistent with observations from the lateral entorhinal cortex (Frank et al., 2000). The present model makes a number of predictions for the activity of entorhinal cells in the open field. First, these cells should show some directionality. If the preferred direction can be accurately estimated from a population of principal cells (probably a difficult problem in practice), the reconstruction method used here could in principle reconstruct position in the open field from the activity of noisy EC cells like those observed by Quirk et al. (1992). A stronger prediction of the model is that the ensemble firing at a given time should depend on the sequence of events that led up to that time, even in the open field.

In the open field, the model predicts noisy place fields that are consistent across environments, consistent with observations from the entorhinal cortex (Quirk et al., 1992). Recently, Fyhn et al. (2004) showed that the entorhinal cortex contains cells with a broad range of spatial firing properties in the open field. They recorded from a range of locations in entorhinal cortex along a dorsolateral-to-ventromedial axis, corresponding to regions of EC that project to the hippocampus proper along its dorsal to ventral extent respectively. They found that cells in the most ventromedial parts of EC showed no little or no spatial modulation. Cells in the intermediate range showed weak spatial modulation, like those observed previously (Quirk et al., 1992), and like the model discussed here predicts (Figure 1, see also Howard et al., 2005). Surprisingly, in the most dorsolateral parts of the entorhinal cortex Fyhn et al. (2004)

observed cells with well-defined consistent omnidirectional place fields in the open field. These cells typically showed multiple fields within an environment. Remarkably, this allocentric spatial representation was maintained even in animals with lesions to the hippocampus, suggesting that the locus of the precise allocentric representation in the MTL is upstream from the hippocampus proper.

The findings of Fyhn et al. (2004) are extremely important for a number of reasons, not least of which is that the dorsolateral entorhinal cortex provides input to the dorsal hippocampus, which has provided that vast majority of studies of hippocampal place cells (but see Jung, Wiener, & McNaughton, 1994). These findings should provoke a thorough reworking of our views on the role of the hippocampus in constructing a spatial representation. They also point out the importance of systematically mapping out the firing properties of hippocampal cells along the septotemporal axis. The finding of entorhinal cells with well-defined omnidirectional place fields in dorsolateral entorhinal cortex present a theoretical challenge for the view that the entorhinal cortex supports a general representation of temporal-spatial context. Although the present model of the place code apparently does a reasonable job at describing the activity of cells in the intermediate zone, more dorsolateral regions of entorhinal cortex behave differently. It is possible that $\mathbf{t}_i$ is one of several computations that the entorhinal cortex computes. Perhaps ventromedial entorhinal regions, where Fyhn et al. (2004) did not observe spatial modulation of firing, implement $\mathbf{t}_i$, but do not receive velocity inputs derived from self-motion information, instead receiving non-spatial inputs. The omnidirectional place fields in dorsolateral entorhinal cortex could result from inputs from the intermediate zone of entorhinal cortex, or some other upstream region. In any event, it is clear that a systematic exploration of the spatial, and especially non-spatial, firing correlates of cells in entorhinal cortex is long overdue and would shed considerable light on current models of hippocampal function. The computational function of the hippocampus cannot be understood without a systematic a study of its inputs and outputs.

Constructing an accurate spatial representation

The finding that we could reconstruct position from $\mathbf{t}_i$ reasonably well using a population vector method suggests that $\mathbf{t}_i$ contains sufficient amounts of information about allocentric position in the open field to construct a place code like that observed in the hippocampus. The existence of spatial information does not directly answer the question of how such a place code could be constructed. It does, however, provide a couple of hints. First, the population reconstruction method presented here takes the natural log of the firing rate. This suggests that downstream neurons should have an input-output relationship that implements this property. Further, to get an accurate reading of position in one direction, it is necessary to not only observe one cell with a preferred direction pointing along the axis one is interested in, but also observe cells with a preferred direction pointing in the opposite direction. This suggests that downstream cells that compute an omnidirectional place code in the open field might take input from a non-random set of cells that participate in $\mathbf{t}_i$, preferentially sampling matching pairs. This does not necessarily require tuning to pairs of cells that differ by precisely π in their preferred directions. Perhaps tuning to anticorrelated pairs of cells would be sufficient.

A set of downstream cells that compute omnidirectional place fields from input consisting of $\mathbf{t}_i$ would need to receive input from a number of cells and construct a conjunctive representation from them. Existing models of omnidirectional place fields in the hippocampus (e.g. Brunel & Trullier, 1998; Kali & Dayan, 2000; Hartley, Burgess, Lever, Cacucci, & O'Keefe, 2000) have also proposed that the hippocampus computes a conjunctive representation of its entorhinal inputs. A conjunctive function for the hippocampus has also been proposed on the basis of considerations from the performance of non-spatial memory tasks (e.g., O'Reilly & Rudy, 2001).

Regardless of the nature of the coding in hippocampus, an assertion of the mapping hypothesis between TCM and the MTL is that a primary function of the hippocampus in spatial navigation is to allow reconstruction of states in entorhinal cortex in response to discrete objects. This view is very similar to that proposed by Burgess and colleagues (e.g. Burgess, 2002; Burgess, Maguire, & O'Keefe, 2002). This position is supported by the suggestion that the hippocampus apparently plays a time-delimited role in the establishment of cognitive maps (Winocur, Moscovitch, Fogel, Rosenbaum, & Sekeres, 2005; Rosenbaum, Ziegler, Winocur, Grady, & Moscovitch, 2004 but see Clark, Broadbent, & Squire, 2005), as well as the observed deficit in stimulus-place associations found in hippocampallesioned animals (Gilbert & Kesner, 2002).

Conclusions

We examined the properties of the place code that result from driving $\mathbf{t}_i$ from the temporal context model, a formal model of episodic recall performance, with velocity vectors during spatial navigation. In particular, we attempted to reconstruct a veridical representation of position from the noisy place code that results from driving $\mathbf{t}_i$ with velocity vectors during random exploration in an open field. We found that a simple population coding scheme in which each cell votes in the preferred direction of the head direction cell that drives it with the natural logarithm of its firing rate was successful in providing good spatial reconstruction with as few as eight cells. We conclude that it is possible that the accurate representation of place observed in the hippocampus could, in principle at least, be extracted from $\mathbf{t}_i$. Perhaps episodic memory performance and the hippocampal place code both result from a representation of temporal-spatial context.

References

Abbott LF. Decoding neuronal firing and modelling neural networks. Quarterly reviews of biophysics 1994;27(3):291–331. [PubMed: 7899551]

Anderson JR, Bower GH. Recognition and retrieval processes in free recall. Psychological Review 1972;79(2):97–123.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, p. 89–105). New York: Academic Press.

Bjork RA, Whitten WB. Recency-sensitive retrieval processes in long-term free recall. Cognitive Psychology 1974;6:173–189.

Brunel N, Trullier O. Plasticity of directional place fields in a model of rodent CA3. Hippocampus 1998;8 (6):651–665. [PubMed: 9882023]

Bunsey M, Eichenbaum HB. Conservation of hippocampal memory function in rats and humans. Nature 1996;379(6562):255–257. [PubMed: 8538790]

Burgess N. The hippocampus, space, and viewpoints in episodic memory. Quarterly Journal of Experimental Psychology 2002;55(4):1057–80. [PubMed: 12420984]

Burgess N, Maguire EA, O'Keefe J. The human hippocampus and spatial and episodic memory. Neuron 2002;35(4):625–41. [PubMed: 12194864]

Chance FS, Abbott LF. Divisive inhibition in recurrent networks. Network 2000;11(2):119–29. [PubMed: 10880002]

Chance FS, Abbott LF, Reyes AD. Gain modulation from background synaptic input. Neuron 2002;35 (4):773–82. [PubMed: 12194875]

Clark RE, Broadbent NJ, Squire LR. Hippocampus and remote spatial memory in rats. Hippocampus 2005;15(2):260–72. [PubMed: 15523608]

Davelaar EJ, Goshen-Gottstein Y, Ashkenazi A, Usher M. A context activation model of list memory: Dissociating short-term from long-term recency effects. Psychological Review 2005;112(1):3–42. [PubMed: 15631586]

Dennis S, Humphreys MS. A context noise model of episodic word recognition. Psychological Review 2001;108(2):452–78. [PubMed: 11381837]

Egorov AV, Hamam BN, Fransén E, Hasselmo ME, Alonso AA. Graded persistent activity in entorhinal cortex neurons. Nature 2002;420(6912):173–8. [PubMed: 12432392]

Eichenbaum H. Hippocampus: cognitive processes and neural representations that underlie declarative memory. Neuron 2004;44(1):109–20. [PubMed: 15450164]

Eichenbaum H, Dudchenko P, Wood E, Shapiro M, Tanila H. The hippocampus, memory, and place cells: is it spatial memory or a memory space? Neuron 1999;23(2):209–226. [PubMed: 10399928]

Ferbinteanu J, Shapiro ML. Prospective and retrospective memory coding in the hippocampus. Neuron 2003;40(6):1227–39. [PubMed: 14687555]

Frank LM, Brown EN, Wilson M. Trajectory encoding in the hippocampus and entorhinal cortex. Neuron 2000;27 (1):169–178. [PubMed: 10939340]

Fransen E, Egorov A, Hasselmo M, Alonso A. Model of graded persistent activity in entorhinal cortex neurons. Society for Neuroscience Abstracts 2003;557.6

Fyhn M, Molden S, Witter MP, Moser EI, Moser MB. Spatial representation in the entorhinal cortex. Science 2004;305(5688):1258–64. [PubMed: 15333832]

Gilbert PE, Kesner RP. Role of the rodent hippocampus in paired-associate learning involving associations between a stimulus and a spatial location. Behavioral Neuroscience 2002;116(1):63–71. [PubMed: 11895184]

Glanzer M, Cunitz AR. Two storage mechanisms in free recall. Journal of Verbal Learning and Verbal Behavior 1966;5:351–360.

Glenberg AM, Bradley MM, Stevenson JA, Kraus TA, Tkachuk MJ, Gretz AL. A two-process account of long-term serial position effects. Journal of Experimental Psychology: Human Learning and Memory 1980;6:355–369.

Graf P, Squire LR, Mandler G. The information that amnesic patients do not forget. Journal of Experimental Psychology : Learning, Memory, and Cognition 1984;10(1):164–78.

van Haeften T, Wouterlood FG, Witter MP. Pre-subicular input to the dendrites of layer-v entorhinal neurons in the rat. Annals of the New York Academy of Science 2000;911:471–3.

Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J. Modeling place fields in terms of the cortical inputs to the hippocampus. Hippocampus 2000;10(4):369–79. [PubMed: 10985276]

Hasselmo ME, Wyble BP. Free recall and recognition in a network model of the hippocampus: simulating effects of scopolamine on human memory function. Behavioural Brain Research 1997;89(1–2):1–34. [PubMed: 9475612]

Higuchi S, Miyashita Y. Formation of mnemonic neuronal responses to visual paired associates in inferotemporal cortex is impaired by perirhinal and entorhinal lesions. Proceedings of the National Academy of Science, USA 1996;93(2):739–743.

Howard MW. Scaling behavior in the temporal context model. Journal of Mathematical Psychology 2004;48:230–238.

Howard MW, Fotedar MS, Datey AV, Hasselmo ME. The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. Psychological Review 2005;112(1):75–116. [PubMed: 15631589]

Howard MW, Kahana MJ. Contextual variability and serial position effects in free recall. Journal of Experimental Psychology : Learning, Memory, and Cognition 1999;25:923–941.

Howard MW, Kahana MJ. A distributed representation of temporal context. Journal of Mathematical Psychology 2002;46 (3):269–299.

Howard, M. W., Wingfield, A., & Kahana, M. J. (Revised). Dissociations between recency and association in the temporal context model: A description of the mnemonic deficit observed in cognitive aging.

Jensen O, Lisman JE. Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. Journal of Neurophysiology 2000;83:2602–2609. [PubMed: 10805660]

Jung MW, Wiener SI, McNaughton BL. Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. Journal of Neuroscience 1994;14(12):7347–56. [PubMed: 7996180]

Kahana MJ. Associative retrieval processes in free recall. Memory & Cognition 1996;24:103–109.

Kali S, Dayan P. The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. Journal of Neuroscience 2000;20(19):7463–7477. [PubMed: 11007906]

Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J. Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature 2002;416(6876):90–4. [PubMed: 11882899]

Levy WB. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. Hippocampus 1996;6:579–590. [PubMed: 9034847]

Loewenstein Y, Sompolinsky H. Temporal integration by calcium dynamics in a model neuron. Nature Neuroscience 2003;6 (9):961–7.

Markus EJ, Qin YL, Leonard B, et al. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. Journal of Neuroscience 1995;15(11):7079–7094. [PubMed: 7472463]

Muller RU, Bostock E, Taube JS, Kubie JL. On the directional firing properties of hippocampal place cells. Journal of Neuroscience 1994;14(12):7235–51. [PubMed: 7996172]

Murdock BB. A theory for the storage and retrieval of item and associative information. Psychological Review 1982;89:609–626.

Nadel L, Moscovitch M. Memory consolidation, retrograde amnesia and the hippocampal complex. Current Opinion in Neurobiology 1997;7(2):217–27. [PubMed: 9142752]

O'Keefe J, Dostrovsky J. The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. Brain Research 1971;34:171–175. [PubMed: 5124915]

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map.* New York: Oxford University Press.

O'Reilly RC, Rudy JW. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. Psychological Review 2001;108(2):311–345. [PubMed: 11381832]

Quirk GJ, Muller RU, Kubie JL. The firing of hippocampal place cells in the dark depends on the rat's recent experience. Journal of Neuroscience 1990;10(6):2008–17. [PubMed: 2355262]

Quirk GJ, Muller RU, Kubie JL, Ranck JB Jr. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. Journal of Neuroscience 1992;12 (5):1945–1963. [PubMed: 1578279]

Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 14, p. 207–262). New York: Academic Press.

Rosenbaum RS, Ziegler M, Winocur G, Grady CL, Moscovitch M. "i have often walked down this street before": fMRI studies on the hippocampus and other structures during mental navigation of an old environment". Hippocampus 2004;14 (7):826–35. [PubMed: 15382253]

Save E, Cressant A, Thinus-Blanc C, Poucet B. Spatial firing of hippocampal place cells in blind rats. Journal of Neuroscience 1998;18(5):1818–1826. [PubMed: 9465006]

Schwartz, G., Howard, M. W., Jing, B., & Kahana, M. J. (in press). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychonomic Science*

Shiffrin RM. Forgetting: Trace erosion or retrieval failure? Science 1970;168:1601–1603. [PubMed: 5420546]

Shon AP, Wu XB, Sullivan DW, Levy WB. Initial state randomness improves sequence learning in a model hippocampal network. Physical review E, Statistical, nonlinear, and soft matter physics 2002;65(3 Pt 1):031914.

Taube JS. Head direction cells and the neurophysiological basis for a sense of direction. Progress in Neurobiology 1998;55(3):225–256. [PubMed: 9643555]

Ward G, Tan L. The effect of the length of to-be-remembered lists and intervening lists on free recall: a reexamination using overt rehearsal. Journal of Experimental Psychology : Learning, Memory, and Cognition 2004;30(6):1196–210.

Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. Science 1993;261:1055–8. [PubMed: 8351520]

Winocur G, Moscovitch M, Fogel S, Rosenbaum RS, Sekeres M. Preserved spatial memory after hippocampal lesions: effects of extensive experience in a complex environment. Nature Neuroscience 2005;8(3):273–5.

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 2000;27(3):623–33. [PubMed: 11055443]

Wu X, Levy WB. Simulating symbolic distance effects in the transitive inference problem. Neurocomputing, 38– 2001;40:1603–1610.

Wu, X. B., & Levy, W. B. (1998). A hippocampal-like neural network model solves the transitive inference problem. In J. M. Bower (Ed.), *Computational neuroscience: Trends in research* (p. 567– 572). New York: Plenum Press.

Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. Journal of Neurophysiology 1998;79(2):1017–44. [PubMed: 9463459]
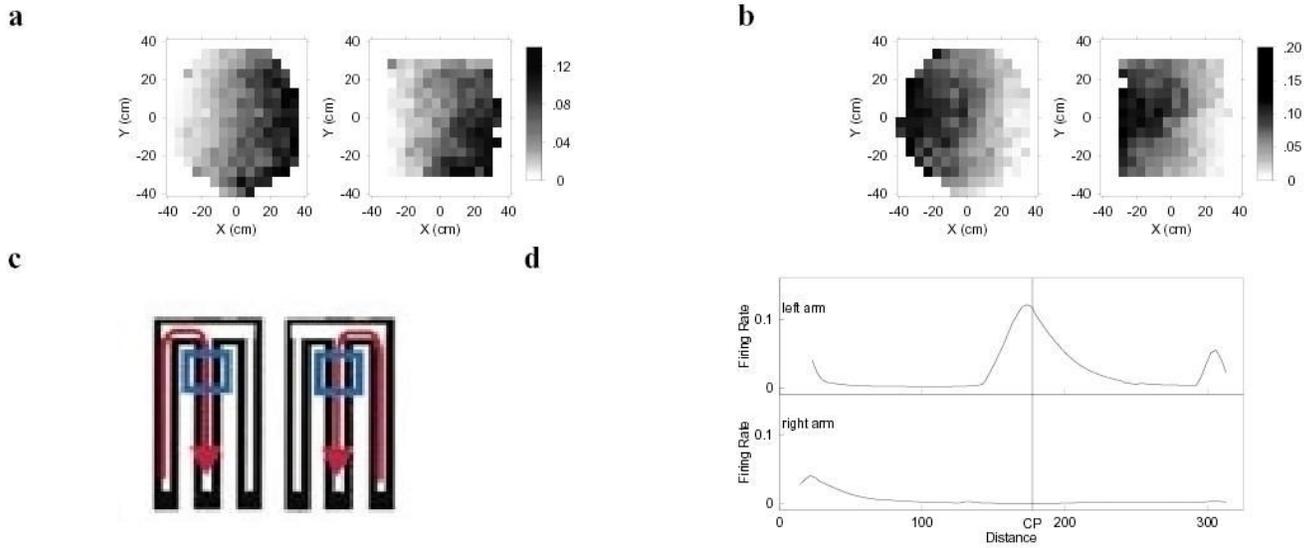
Figure 1. Simulated entorhinal place cells from Eq. 1 (Adapted from Howard *et al*, 2005)
a,b. Simulated entorhinal place fields from the open field. Like entorhinal cells (Quirk *et al*, 1992), the simulated cells showed noisy place fields that were correlated across similar environments. The cell in **a** took input from a head direction cell with a preferred direction that pointed towards the east. The cell in **b** received input from a head direction cell with a preferred direction that pointed west-by-northwest. **c.** Frank *et al*, (2000) trained rats to run in a W-maze in which the animal successively visited a food well in each arm in sequence. This paradigm provides the opportunity to examine firing of cells in the middle arm of the maze as part of different trajectories. **d.** Simulated entorhinal cells showed retrospective coding that distinguished the path the animal had taken to reach a point on the middle arm. A representative simulated retrospective cell shows different firing on the center arm on inbound paths depending on which arm the animal was coming from. The choice points, CP, is a point along the path that is several centimeters from the top of the middle arm. The choice point is indicated by the top of the box in **c**. Because position and heading are the same in both cases (see **c**), this indicates that the model, like the entorhinal cortex, is responding to more than position *per se*. See Howard *et al* 2005 for details.
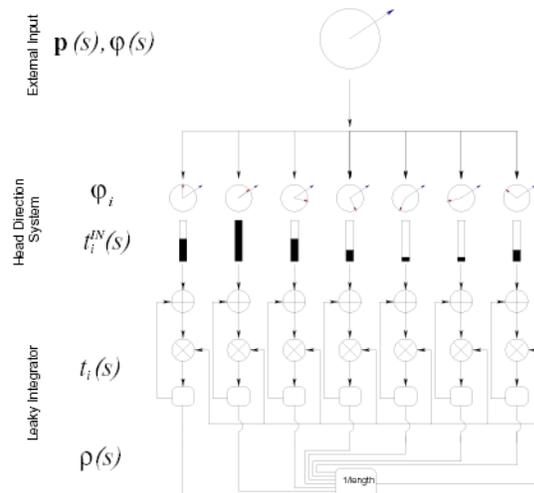
Figure 2. Schematic diagram of the cellular simulation
The input vector $\mathbf{t}^{IN}$ is constructed from a set of units that receive input from head direction cells. Each unit has a preferred direction that determines its activation. This input projects to a set of integrator cells. These maintain their firing. The activity of the integrator cells is constrained by divisive inhibition $\rho(s)$.
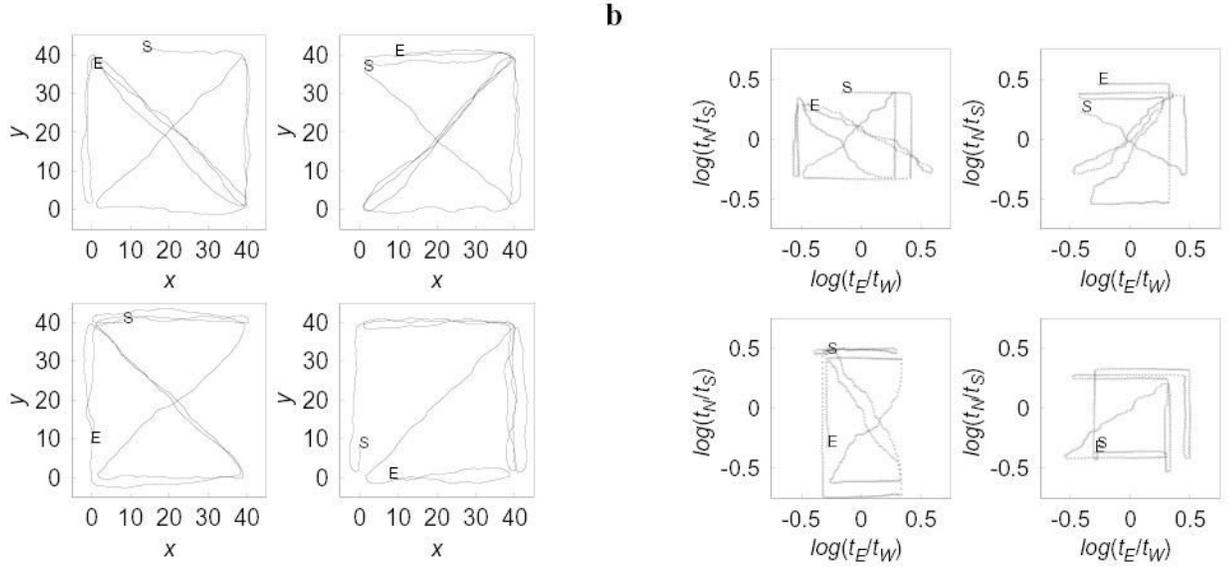
a

b

Figure 3. Physical space projects onto log firing rate space
A four-cell simulation with normalized inputs was run during simulated movements among the corners of a box. **a.** Successive paths of the simulated movement in physical space. **b.** The activity of the simulated neurons projected into log space for the sequences corresponding to **a**. The x axis in these plots is $\log(t_E/t_W)$, where t_E is the firing rate of the first unit, with preferred direction pointing toward the east. When the animal moves north/south, $\log t_E/t_W$ is unchanged, despite the fact that both t_E and t_W are subject to exponential decay.
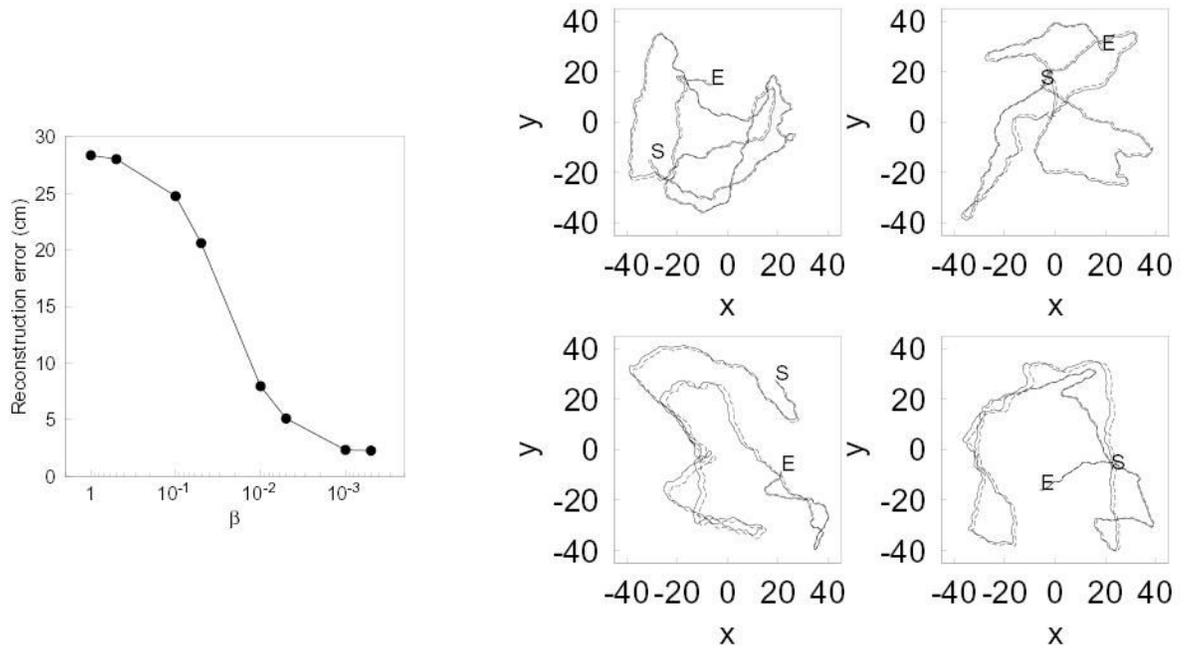
Figure 4. Error reconstruction
Left: Reconstruction error in cm as a function of β for reconstruction with eight cells. The reconstruction error decreases to about the size of a rat with β as large as .01. **Right:** Sample simulated paths (solid) with corresponding reconstructed paths (dashed) for β = .001.