



HHS Public Access

Author manuscript

Brain Res. Author manuscript; available in PMC 2016 September 24.

Published in final edited form as:

Brain Res. 2015 September 24; 1621: 345–354. doi:10.1016/j.brainres.2014.10.069.

Time and space in the hippocampus

Marc W. Howard and Howard Eichenbaum

Center for Memory and Brain and Department of Psychological and Brain Sciences, Boston University

Abstract

It has been hypothesized that one of the functions of the hippocampus is to enable the learning of relationships between different stimuli experienced in the environment. These relationships might be spatial (“the bathroom is about 5 m down the hall from the bedroom”) or temporal (“the coffee is ready about 3 minutes after the button was pressed”). Critically, these spatial and temporal relationships may exist on a variety of scales from, say a few hundred milliseconds up to minutes. In order to learn consistent relationships between stimuli separated by a variety of spatial and temporal scales using synaptic plasticity that has a fixed temporal window extending at most a few hundred milliseconds, information about the spatial and temporal relationships of distant stimuli must be available to the hippocampus in the present. Hippocampal place cells and time cells seem well-suited to represent the spatial and temporal locations of distant stimuli in order to support learning of these relationships. We review a recent computational hypothesis that can be used to construct both spatial and temporal relationships. We suggest that there is a deep computational connection between spatial and temporal coding in the hippocampus and that both serve the overarching function of learning relationships between stimuli—constructing a “memory space.”

1. Introduction

The modern era of research on hippocampal function began with Scoville and Milner’s (1957) identification of the hippocampus as critical to memory in humans. Subsequent pioneering studies that sought the basis of hippocampal memory function in the firing properties of hippocampal neurons revealed the existence of “place cells”, neurons that fire when a rat occupies a particular location in its environment (O’Keefe & Dostrovsky, 1971). This observation, plus studies showing selective deficits in spatial learning following hippocampal damage, inspired the notion that the hippocampus supports memory by creating a cognitive map of the physical environment into which memories are encoded (O’Keefe & Nadel, 1978). Decades of research since then have sought to characterize the organization of the cognitive map. Recent research in this direction has focused on how the hippocampus and neighboring interconnected structures supports navigation through space (McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Moser, Kropff, & Moser, 2008),

Invited submission for special issue of *Brain Research*, November 7, 2014

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

indicating this line of research is losing sight of the well established more general role of the hippocampal system in memory (Eichenbaum & Cohen, in press). Here we will review some of the key features of hippocampal place cells and suggest that they reflect a more general mechanism of temporal as well as spatial organization of events, and that this broader function is the basis of hippocampal function in memory.

1.1. Coding across space and time in the hippocampus

The hippocampal place code provides an exquisitely detailed representation of the animal's location in allocentric coordinate system at multiple spatial scales (Wilson & McNaughton, 1993; Solstad, Boccara, Kropff, Moser, & Moser, 2008). Eichenbaum (2001) hypothesized that the spatial map in the hippocampus is a special case of a more general function in representing relationships among all of the task-relevant stimuli, including both spatial and non-spatial variables (see also Eichenbaum, Dudchenko, Wood, Shapiro, & Tanila, 1999). If the hippocampus captures relationships among stimuli separated by macroscopic quantities of space and time, and if plasticity depends on simultaneous activation within a fixed temporal window, there must be some way to bridge these gaps. Put another way, if the hippocampus is to facilitate learning of a relationship such as, "these nuts are stored 20 cm from the big oak tree" or "the shock will come 5 s after the tone" then it must have some facility to represent these relationships. At the location where the nuts are being stored, the hippocampus must have some ability to describe the distant location of the oak tree; at the time the shock arrives the hippocampus must have some way to represent that a tone was experienced a certain time in the past. Once the relationship of the distant stimulus is represented, it is relatively straightforward to associate that representation to the other stimulus—the nut in the spatial example or the shock in the temporal example. In this way, the spatial and temporal representations can be used to cue the response. For instance, when the animal is 15 cm from the oak tree, the similarity between the current spatial representation and the stored representation can cause activation of the representation for the nuts. Similarly, if one is looking for nuts, nuts may be used as a cue to recover the spatial representation coding for "20 cm from the oak tree."

1.1.1. Spatial coding reflects spatial relationships to distant stimuli—These considerations suggest that the neural activity in the hippocampus at a particular time is not solely, nor even predominantly, caused by the stimuli available in the environment at that moment. For instance consider the experiment performed by Gothard and colleagues (Gothard, Skaggs, & McNaughton, 1996). They trained a rat to run back and forth on a linear track. Each of the journeys along the track started from a goal box. Across journeys the location of the box was moved relative to the track and the room coordinates. On journeys away from the box, place cells fired at a constant distance from the start box, moving across journeys relative to the room frame. Consider a place cell that fires 50 cm from the start box. Because the location of its firing is controlled by the box, this means that a stimulus that was experienced perhaps several seconds in the past—the box—controls firing of the cell in the present. Moreover, because the phenomenon also holds in the dark (Gothard, Hoffman, Battaglia, & McNaughton, 2001), we can be certain that the firing in this experiment is not dependent on the animal being able to see the goal box from its current location. Burgess and O'Keefe (1996) observed similar control of firing in a two-

dimensional environment; as the shape of the enclosure changed, place fields in the center of the environment changed location to maintain their relationship to one or more of the walls of the environment (see also Hartley, Burgess, Lever, Cacucci, & O'Keefe, 2000; Barry et al., 2006; Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009).

These place cells are apparently not coding for a place *per se* but a spatial relationship relative to a landmark, or set of landmarks. That is, a place cell in the Gothard et al. (1996) experiment is not coding for a location, but rather something more like “50 cm from the goal box”—a spatial relationship. Note that without establishing stimulus control, simply observing that a place cell fires in a particular location within an environment provides no information about the stimulus and relationship (or relationships) that it is being used to represent (Figure 1).

1.2. Temporal coding may reflect temporal relationships of temporally remote stimuli over long time scales

The principle that currently-available stimuli do not control hippocampal firing should also inform non-spatial firing as well. A large body of evidence suggests that the hippocampus constructs temporal and sequential relationships in addition to spatial relationships (e.g., Bunsey & Eichenbaum, 1996; Clark & Squire, 1998; Westmacott & Moscovitch, 2001). It is clear that neural activity supporting memory changes gradually over long periods of time. Although there are examples of some cells in the MTL that respond to non-spatial stimuli currently being perceived, such as the celebrated “Jennifer Aniston” cells (Quiroga, Reddy, Kreiman, Koch, & Fried, 2005), these cells coexist with other cells that show very slowly-changing activity. Figure 2 illustrates this phenomenon in recordings from human epileptic patients during performance of an item recognition experiment using pictures of faces and scenes (Howard, Viskontas, Shankar, & Fried, 2012; Viskontas, Knowlton, Steinmetz, & Fried, 2006). Note that the cells fire at an elevated level over periods of time much longer than the time any individual stimulus is presented. These cells are from human MTL; autocorrelated signals have been obtained from animal studies in a variety of brain regions (Manns, Howard, & Eichenbaum, 2007; Mankin et al., 2012; Hyman, Ma, Balaguer-Ballester, Durstewitz, & Seamans, 2012; Kim, Ghim, Lee, & Jung, 2013). While autocorrelation itself may result from any number of causes, including recording artifacts and stochastic noise, additional evidence suggests that gradually-changing firing could reflect, at least in part, memory for temporally remote events.

1.2.1. Time cells code for when an event took place—If the hippocampal ensemble changes gradually over long periods of time, this creates the possibility that the stimulus that caused the cell to begin firing was experienced far in the past, in much the same way that the place code can reflect stimuli experienced remotely in time. Time cells provide an excellent illustration of this phenomenon. Pastalkova, Itskov, Amarasingham, and Buzsaki (2008) trained rodents to navigate around a maze, stopping on each lap to run for a period of time in a running wheel. As expected, some hippocampal neurons fired in a circumscribed set of locations along the maze. They also observed that while running on the wheel, some hippocampal neurons fired during circumscribed periods of time. Similar results have been observed when the animal is not running at all (MacDonald, Lepage, Eden, & Eichenbaum,

2011; MacDonald, Carrow, Place, & Eichenbaum, 2013) and when the animal is running on a treadmill with controlled velocity (Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013).

These “time cells” carry information about the timing of past events. Figure 3a shows the firing of a set of time cells during the delay interval of an object memory task. The set of cells that fire early in the interval are different than the set of cells that fire during the middle of the interval which are, in turn, different than the set of cells that fire near the end of the interval. The ensemble of neurons carries information about the time since the beginning of the delay interval. The delay interval in this experiment began with a discrete event—the closing of a door to the enclosure where the animal waited during the delay. In these recordings, the accuracy of the timing information decreases as the delay interval proceeds (Howard et al., 2014). This can be seen by the spread in the “time fields” later in the delay interval. Kraus et al. (2013) observed a similar spreading. In addition, they were able to show that the distribution of peaks was not uniform across the delay interval. In addition, time cells can show history-dependence in which they fire differentially based on preceding events (Gill, Mizumori, & Smith, 2011; Pastalkova et al., 2008).

1.2.2. Hippocampal coding for the identity of temporally distant stimuli—Time cells may also be able to code for the time since specific non-spatial stimuli were experienced. A few of the time cells in the MacDonald et al. (2011) experiment coded for the identity of the object presented during the delay. MacDonald et al. (2013) observed different sequences of cells firing during the delay after an odor was presented in a recognition task, although the overall firing rates were quite low. Although it is not yet completely clear to what extent time cells carry information about preceding stimuli, there is abundant evidence that firing in the MTL carries information about non-spatial stimuli that occurred many seconds in the past. Figure 3b shows a non-spatial “splitter cell” observed by (Ginther, Walsh, & Ramus, 2011). In this experiment, the rat was presented one of two series’ of odors with an overlapping subsequence. Let us denote sequence 1 as M N A B C O P and sequence 2 as W X A B C Y Z, with odors A B C being common to the two sequences. During the time when the overlapping odors are presented, the locally available information is identical; only the past history of odor presentation distinguishes the two sequences from one another. Some hippocampal neurons fired differentially during the overlapping sequence (e.g., Fig. 3b), reflecting a sensitivity to previous stimuli. Hsieh, Gruber, Jenkins, and Ranganath (2014) observed analogous findings from humans using fMRI. In addition, Yakovlev, Fusi, Berman, and Zohary (1998) observed gradually-decaying stimulus-specific firing that persisted across the intertrial interval of a DNMS task in monkey perirhinal cortex. In humans, (Howard et al., 2012) showed that the gradually-changing representation in the MTL (Figure 2) was indeed caused by the stimulus presented to the patient. This rules out the possibility that the autocorrelation is purely a recording artifact or stochastic noise; the gradually-changing representation is at least in part a signal that carries information about stimuli that were experienced in the past.

While the significance of slowly-changing firing, as shown in Figure 2 is not yet clear, these findings raise the intriguing possibility that the representation carries information about the time of past events. Representations in the MTL change gradually over at least many

minutes (Manns et al., 2007; Hyman et al., 2012; Mankin et al., 2012; Howard et al., 2012). Thus far, time cells have not been observed firing over more than a few seconds. However, human and animal memory shows similar properties at a variety of time scales (Balsam & Gallistel, 2009; Glenberg et al., 1980; Howard, Youker, & Venkatadass, 2008; Lejeune & Wearden, 2006; Wearden & Lejeune, 2008; Lewis & Miall, 2009), suggesting that the slow changes in the neural representation could reflect a gradually-changing representation of history that supports behavior over a range of time scales (Howard & Eichenbaum, 2013; Howard, Shankar, Aue, & Criss, in press).

2. A computational approach to spatiotemporal context in the hippocampus

In order for the hippocampus to learn consistent spatial and temporal relationships, it must have the ability to represent temporal and spatial position of distant stimuli. Howard et al. (2014) described a computational mechanism to do just this. We start by describing, in a qualitative manner, how the model is able to represent temporal information, then generalize this ability to spatial relationships. After this description, we present a simulation illustrating the computational model's ability to describe the continuity between temporal and spatial variables by describing experimental results from Kraus et al. (2013).

2.1. A method for computing scale-invariant spatial and temporal relationships

The goal of the computational model introduced by Shankar and Howard (2012, 2013) is to, at each moment, estimate the history of stimuli leading up to the present. The history is encoded as a conjunctive representation of what stimulus happened at each time in the past. This timeline is encoded by a set of neurons each of which are activated when a particular stimulus was experienced a certain time in the past. Presenting a stimulus causes a sequence of neurons to be activated as the stimulus recedes further into the history (Fig. 3a). The "time fields" of the neurons are spread out over a range of times, with time fields representing times further in the past showing more spread, consistent with experimental findings (Howard et al., 2014; Kraus et al., 2013; MacDonald et al., 2013). As a consequence of this spread, the model represents distant events with less temporal accuracy. It turns out that, in the model, the decrease in accuracy is precisely scale-invariant (Shankar & Howard, 2012), a property that is very useful in explaining behavioral properties of learning and memory (Balsam & Gallistel, 2009; Chater & Brown, 2008; Gallistel & Gibbon, 2000; Howard et al., in press).

Constructing a representation of temporal history with these properties is computationally non-trivial (Shankar, 2014). The output of the model is generated from an intermediate state. The intermediate state in the temporal case is a set of leaky integrators. Each cell in the intermediate state accepts the input present in the environment at that moment and then decays gradually as time passes; each of the cells accepting input from the same stimulus decay at different rates. The model can be generalized from coding time to space by changing the rate of decay to depend on the animal's velocity in a particular direction (Howard et al., 2014). In this case, the intermediate state is a path integrator. A variable α controls the behavior of each of the cells. If α is a constant over time, then that cell is a leaky integrator and the output constructed from that intermediate state will be a temporal history. Cells in this representation code for the time since a stimulus was encoded. If α is

set at each moment to be the velocity of the animal, then that cell is a path integrator and the output constructed from that intermediate state will be a spatial representation. Cells in the spatial representation code for the net distance traveled, along some direction, since a stimulus—here a landmark—was encountered in the environment. This framework can be extended further to also generate representations of sequential position as well as temporal and spatial relationships (Howard et al., 2014).

Because an animal moves around an environment in smooth trajectories, space and time are highly correlated with one another under normal circumstances. Consider a linear track. If an animal starts from one end and travels to the other with constant velocity, position along the track and time since the departure are perfectly confounded with one another. Observing a cell that fires during a circumscribed region of the track cannot determine whether that cell is a “place cell” or a “time cell.” In order to systematically dissociate position from time *per se*, it is necessary to establish experimental control of velocity.

2.2. Demonstration: The continuity between spatial and temporal coding

Kraus et al. (2013) established control over velocity by having animals run on a treadmill. They varied velocity across trials and examined firing as a function of both time and of distance traveled on the treadmill. They characterized the sensitivity of each cell to distance traveled compared to time since the treadmill started. Some cells responded predominantly to position; others predominantly to time. Critically, they found a continuity of sensitivity to these variables recording from the same population of cells; many cells showed some sensitivity to both time elapsed and distance traveled.

This continuity between spatial coding and temporal coding can be readily explained in the computational framework we have described. Figure 4 shows simulated results of the constant distance condition of the Kraus et al. (2013) experiment. We started the interval with an input, corresponding to the start of the treadmill. Then, the simulated animal ran on the treadmill at one of three speeds until a certain distance was traversed on the treadmill. The treadmill advanced at different speeds. The cells in the intermediate state were allowed to evolve with their value of a set according to $a = a + b s$, where s is just the speed on that trial. The activity of a representative cell in the output representation is shown at each moment of trials run at the three speeds as a function of time (left) or distance traveled along the treadmill (right).

Figure 4a shows the results for a cell whose intermediate representation was dominated by temporal information, with a set close to 1 and b set close to zero. The firing of this cell is controlled by the time since the treadmill running started and does not depend directly on the distance traveled. It is fair to call the cell in Figure 4a a time cell. In contrast, the cell in Figure 4c is dominated by spatial information, with a set close to zero and b set close to 1. This cell fires at different times depending on the speed of the trial (left) but its firing is well controlled by the distance traveled on that trial (right). Interestingly, the model **a b** also admits hybrid cells that are partially modulated by temporal information and partially modulated by spatial information. The cell in Figure 4b shows a cell with $a = b = 0.5$. This cell is neither a pure time cell nor a pure distance cell, but is intermediate between the two. By choosing different values of a and b we can generate a cell with any degree of relative

sensitivity to temporal and spatial information. The continuity between temporal and spatial information observed by Kraus et al. (2013) amounts to being able to set a and b to a wide variety of values.

2.3 Exploiting spatial and temporal representations to behaviorally express relationships among stimuli

The computational framework described here is able to construct representations of distant stimuli. Here we discuss how one can use this spatiotemporal representation to build relationships between different stimuli that can be expressed behaviorally. This approach uses two forms of association which we refer to context-to-item binding and item-to-context binding. Each of these associations are assumed to depend on synaptic plasticity although the mechanisms for the two forms of association may be quite different.

Association is a longstanding theme in computational models of memory performance (e.g., Murdock, 1982; Raaijmakers & Shiffrin, 1980; Humphreys, Bain, & Pike, 1989). The idea is that a relationship between two stimuli, say A and B, is stored as a pattern of synaptic weights joining the representations of the two stimuli. The synapses joining the active units of A and the active units of B are strengthened during learning. When A is repeated as a probe, the active units in the representation of A cause the pattern in B to be reinstated *via* the synapses that were strengthened during learning. In the brain, associations could be instantiated by much more complex relationships than simple Hebbian association and the pathways supporting association may be much more detailed than a single set of weights that direction connect two sets of neurons. Nonetheless, this associationist framework provides a simple model to understand how a representation of spatiotemporal context could be used to express behavioral relationships between different stimuli.

Context-to-item bindings provide one way that behavioral associations can be supported by synaptic plasticity. Suppose that associations can be formed between a representation of spatiotemporal context and the currently-available stimulus—an item. To be more concrete, let the task be to learn that A predicts B not immediately, but at a delay of 10 s, as might be required in trace conditioning. Note that a simple association between the representations of A and B is insufficient to describe this temporal relationship. However, note that a temporal representation that codes for information about the time of occurrence of A could be used to support trace conditioning. During study, a series of stimulus-specific time cells is activated by presentation of A. When B is presented 10 s later, only the currently-active cells can be associated to the representation of B. When A is repeated, the same sequence of cells is activated. The cells firing 1 or 2 s after A have not been associated with the representation of B. As time passes, approaching 10 s, the overlap between the state of the temporal representation and the set of cells associated to B grows, resulting in increased activation of B at the appropriate time (see also Shankar & Howard, 2012; Howard & Eichenbaum, 2013).

In addition to the behavioral effects of a representation of spatiotemporal context on the activation of item representations, we can also consider the possibility that item representations can reinstate previous states of spatiotemporal context. The ability of items to recover previously encoded states of context is a common feature of several quantitative

models of memory performance (Dennis & Humphreys, 2001; Howard & Kahana, 2002). This provides a concrete definition of the “jump back in time” used to describe the phenomenological experience of episodic memory (e.g., Tulving, 1983). Although not yet definitive, a growing body of evidence from human single unit recordings, ECoG data and fMRI suggest that memory retrieval is accompanied by the reinstatement of a gradually-changing neural representation (DuBrow & Davachi, 2014; Howard et al., 2012; Manning, Polyn, Litt, Baltuch, & Kahana, 2011; Zeithamova, Dominick, & Preston, 2012), a qualitative prediction of these models. The mechanism by which contextual reinstatement is accomplished is not typically specified by mathematical models that describe behavior and may be quite complex (Howard, Fotedar, Datey, & Hasselmo, 2005).

If it exists, this hypothesized contextual retrieval mechanism provides a mechanism for learning behavioral associations between stimuli that are not presented close together in time. Consider a learning task in which subjects learn pairs of stimuli that include overlapping items, for instance the pairs A–B and B–C (e.g., Bunsey & Eichenbaum, 1996; Howard, Jing, Rao, Provn, & Datey, 2009; Primoff, 1938; Slamecka, 1976; Zeithamova et al., 2012). In a representative experiment, many pairs are learned in a random order so that there is no consistent temporal relationship between, say, A and C. Nonetheless they share a higher-order relationship by virtue of having both been experienced in proximity to B. Contextual retrieval makes this possible. Suppose that A–B is learned and then, much later B–C is studied. If during study of B–C, the stimulus B can reinstate its previous context, which includes information about the time of presentation of A, then the context in which C is studied includes information about A. If A is later repeated as a cue, this will tend to activate C resulting in an association between A and C. Because the representation of spatiotemporal context discussed here can capture spatial and temporal relationships, the mechanism of contextual reinstatement can account for the results of temporal mapping experiments in which subjects not only learn that A and C go together, but incorporate specific temporal relationships (Arcediano, Escobar, & Miller, 2003; Cole, Barnet, & Miller, 1995; Savastano & Miller, 1998). For instance, if the subject initially learns that A predicts a shock at a delay of 5 s and that A predicts B at a delay of four seconds, the subject will behave as if B predicts the shock at a delay of 1 s.

To summarize, reinstatement of a spatiotemporal context representation coupled with simple associations between the contextual representation and item representations could endow the hippocampal memory system with the ability to flexibly express a rich set of relationships between stimuli. Even though the elementary operations are simple, if contextual reinstatement is under strategic control, then a broad variety of task-appropriate relationships could be retrieved.

3. Discussion

This paper has reviewed an integrated computational hypothesis for spatial and temporal coding in the hippocampus. This approach is a quantitative implementation of longstanding ideas about hippocampal coding. Depending on the inputs to the model and the form of modulation, the method can construct both temporal and spatial representations. These representations can be used to store spatial and temporal relationships among stimuli

manifest in behavior. A model of flexible behavioral expression of relationships, requires that this representation is augmented with mechanisms for storing context-to-item associations and for reinstating prior states of spatiotemporal context. In this way, relationships between stimuli separated in space and time could be extracted to flexibly inform behavior.

3.1 Spatial and temporal input streams to the hippocampus

Information about experiences arrives in the hippocampus via two main cortical pathways (Van Essen, Felleman, DeYoe, Olavarria, & Knierim, 1990; Eichenbaum, Yonelinas, & Ranganath, 2007). The origins of these pathways are the well described “what” and “where” streams of sensory processing that represent, respectively, the identity and spatial location of perceptual objects (Mishkin & Ungerleider, 1982). These streams are best described in the visual system but also incorporate information from all of the other sensory modalities (Felleman & Van Essen, 1991; Burwell & Amaral, 1998; Suzuki & Amaral, 1994). These anatomical observations strongly suggest that a major function of the hippocampus is to integrate information about the nature of events and the spatial context in which they occur (Davachi, 2006; Eichenbaum et al., 2007). Several studies have now identified major end points of the what stream in the perirhinal cortex and lateral entorhinal cortex, and of the where stream in the parahippocampal cortex and medial entorhinal cortex, as specialized for processing event and spatial information (Eichenbaum et al., 2007). In addition, many studies have now shown that hippocampal networks integrate event and spatial contextual information (e.g., Komorowski, Manns, & Eichenbaum, 2009; McKenzie et al., 2014, reviewed in Davachi, 2006; Eichenbaum et al., 2007). Since there is now a wealth of evidence that the hippocampus also plays a key role in the temporal organization of events in memory (Eichenbaum, 2013), a fundamental question is how temporal coding is generated from information provided by the what and where streams (Eichenbaum, in press).

3.2 Extracting temporal and spatial regularities to construct semantic “maps”

Given that the hippocampus has access to a code that spatial and temporal relationships between stimuli that were experienced at distant times and places, semantic relationships can be learned by extracting regularities between stimuli experienced with consistent spatial and temporal relationships. Cognitive science has generated specific computational models of semantic memory (e.g., Griffiths, Steyvers, & Tenenbaum, 2007; Howard, Shankar, & Jagadisan, 2011; Landauer & Dumais, 1997) that describe mechanisms by which the temporal contexts of isolated experience could be integrated to form efficient networks of semantic memories. Neuroimaging work (Schapiro, Kustner, & Turk-Browne, 2012, see also Miyashita, 1988) makes it clear that MTL representations of stimuli come to reflect their temporal contexts, even when there are no explicit task demands to learn the relationships (see also Bunsey & Eichenbaum, 1996).

Human subjects can learn semantic relationships between stimuli by simply observing the contexts in which they appear. Consider being exposed to the novel word FLOOB in the sentence “The baker reached into the oven and pulled out the FLOOB.” Knowledge about other words that fit into the context lets us know that FLOOB is more similar to “bread” or

“cake” than to “democracy” or “spaceship.” The idea of learning from contexts has been formalized in computational models of semantic learning. These models observe large amounts of naturally occurring text and then use statistical techniques of varying sophistication to construct representations of the words that capture their contextual relationships (e.g. Griffiths et al., 2007; Landauer & Dumais, 1997; Shankar, Jagadisan, & Howard, 2009).

Neural evidence shows that long-term memory representations of stimuli in the MTL also reflects the temporal context in which the stimuli were experienced. For instance, Miyashita (1988) found neurons in inferotemporal (IT) cortex that responded to sample stimuli in a working memory task in which the animal was required to remember a sample stimulus during a delay interval. The sample stimuli were presented in the same order on each of several recording sessions. Remarkably, a cell that responded to one sample stimulus was more likely to respond to other sample stimuli presented close in the list despite the fact that the stimuli were separated by dozens of seconds during learning. This tendency to code stimuli experienced close together in time was subsequently observed in perirhinal cortex (Naya, Yoshida, & Miyashita, 2001; Erickson & Desimone, 1999; Messinger, Squire, Zola, & Albright, 2001) and appears to depend on connections from MTL regions back to IT cortex (Higuchi & Miyashita, 1996; Miyashita, Kameyama, Hasegawa, & Fukushima, 1998; Naya et al., 2001).

Schapiro et al. (2012) showed that the neural representations of visual stimuli in the MTL, as measured by representational similarity analysis, came to reflect their temporal context. At the beginning of a session, subjects were presented with a set of fractal images. Then the subjects viewed the images during a training session while performing a cover task that required them to detecting a slight alteration in some of the images. During training, some pairs of non-target images were repeatedly presented in sequence. Critically, pairs were assigned randomly from the pool. After training, the images were presented in a random order. After training, the pattern of BOLD response across hippocampal voxels was more similar for members of the same pair than for images that were not paired. This effect held throughout the subregions of the hippocampus.

Schapiro, Rogers, Cordova, Turk-Browne, and Botvinick (2013) showed that the ability to learn stimulus representations from temporal context extends to more complex networks of relationships. They found that stimuli were presented sequentially during presentation and shared common neighbors became more similar to one another than stimuli that were simply paired during presentation but whose temporal neighbors did not overlap to the same extent. The fact that hippocampal representations reflect these relatively subtle aspects of the statistical structure of experience suggests that the hippocampus learns relational networks from temporally structured experience. Time cells could play an essential role in representing the temporal structure of experience so that it can be extracted into lasting stimulus representations that reflect temporal structure.

Conclusions

We suggest that a general function of the hippocampus is to construct “memory spaces” that capture reliable spatial and temporal relationships between stimuli. In order to construct these networks within a fixed window of synaptic plasticity, the hippocampus must be able to, at each moment, have access to information about the location and temporal distance from stimuli that were experienced in the past. Place cells and time cells are ideally suited to provide this capacity. We briefly reviewed computational model that codes for spatial and temporal representations and showed that it describes a continuity between temporal and spatial representation, consistent with that observed experimentally (Kraus et al., 2013). This representation could be used to construct a learning model that encodes temporal and spatial relationships into a persistent semantic memory. This hypothesis is at least roughly consistent with neurophysiological studies of long-term memory representations and recent fMRI studies of learning of temporal contexts.

Acknowledgments

The authors gratefully acknowledge support from AFOSR FA9550-12-1-0369 and NSF BCS-1058937 to MWH and NIH grants MH094263, MH095297, MH052090 to HBE.

References

- Arceidiano F, Escobar M, Miller RR. Temporal integration and temporal backward associations in human and nonhuman subjects. *Learning & Behavior*. 2003; 31(3):242–56. [PubMed: 14577548]
- Balsam PD, Gallistel CR. Temporal maps and informativeness in associative learning. *Trends in Neuroscience*. 2009; 32(2):73–78.
- Barry C, Lever C, Hayman R, Hartley T, Burton S, O’Keefe J, Burgess N. The boundary vector cell model of place cell firing and spatial memory. *Reviews in Neuroscience*. 2006; 17(1–2):71–97.
- Bunsey M, Eichenbaum HB. Conservation of hippocampal memory function in rats and humans. *Nature*. 1996; 379(6562):255–257. [PubMed: 8538790]
- Burgess N, O’Keefe J. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*. 1996; 6(6):749–62. [PubMed: 9034860]
- Burwell RD, Amaral DG. Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *Journal of Comparative Neurology*. 1998; 398(2):179–205. [PubMed: 9700566]
- Chater N, Brown GDA. From universal laws of cognition to specific cognitive models. *Cognitive Science*. 2008; 32(1):36–67. 10.1080/03640210701801941 [PubMed: 21635331]
- Clark RE, Squire LR. Classical conditioning and brain systems: the role of awareness. *Science*. 1998; 280(5360):77–81. [PubMed: 9525860]
- Cole RP, Barnett RC, Miller RR. Temporal encoding in trace conditioning. *Animal Learning & Behavior*. 1995; 23(2):144–153.
- Davachi L. Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology*. 2006; 16(6):693–700. 10.1016/j.conb.2006.10.012 [PubMed: 17097284]
- Dennis S, Humphreys MS. A context noise model of episodic word recognition. *Psychological Review*. 2001; 108(2):452–78. [PubMed: 11381837]
- DuBrow S, Davachi L. Temporal memory is shaped by encoding stability and intervening item reactivation. *Journal of Neuroscience*. 2014 Oct; 34(42):13998–4005. 10.1523/JNEUROSCI.2535-14.2014 [PubMed: 25319696]
- Eichenbaum H. The hippocampus and declarative memory: cognitive mechanisms and neural codes. *Behavioural Brain Research*. 2001; 127(1–2):199–207. [PubMed: 11718892]
- Eichenbaum H. Memory on time. *Trends in Cognitive Sciences*. 2013; 17(2):81–8. 10.1016/j.tics.2012.12.007 [PubMed: 23318095]

- Eichenbaum H. Time cells in the hippocampus: A new dimension for mapping memories. *Nature Reviews Neuroscience*. (in press).
- Eichenbaum H, Cohen NJ. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron*. (in press).
- 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]
- Eichenbaum H, Yonelinas A, Ranganath C. The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*. 2007; 30 :123–152.
- Erickson CA, Desimone R. Responses of macaque perirhinal neurons during and after visual stimulus association learning. *Journal of Neuroscience*. 1999; 19(23):10404–10416. [PubMed: 10575038]
- Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*. 1991; 1(1):1–47. [PubMed: 1822724]
- Gallistel CR, Gibbon J. Time, rate, and conditioning. *Psychological Review*. 2000; 107(2):289–344. [PubMed: 10789198]
- Gill PR, Mizumori SJY, Smith DM. Hippocampal episode fields develop with learning. *Hippocampus*. 2011; 21(11):1240–9.10.1002/hipo.20832 [PubMed: 20665593]
- Ginther MR, Walsh DF, Ramus SJ. Hippocampal neurons encode different episodes in an overlapping sequence of odors task. *Journal of Neuroscience*. 2011; 31(7):2706–11.10.1523/JNEUROSCI.3413-10.2011 [PubMed: 21325539]
- 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.
- Gothard KM, Hoffman KL, Battaglia FP, McNaughton BL. Dentate gyrus and cal ensemble activity during spatial reference frame shifts in the presence and absence of visual input. *Journal of Neuroscience*. 2001; 21(18):7284–92. [PubMed: 11549738]
- Gothard KM, Skaggs WE, McNaughton BL. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *Journal of Neuroscience*. 1996; 16(24):8027–40. [PubMed: 8987829]
- Griffiths TL, Steyvers M, Tenenbaum JB. Topics in semantic representation. *Psychological Review*. 2007; 114(2):211–44. [PubMed: 17500626]
- 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]
- 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, Eichenbaum H. The hippocampus, time, and memory across scales. *Journal of Experimental Psychology: General*. 2013; 142(4):1211–30.10.1037/a0033621 [PubMed: 23915126]
- 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, Jing B, Rao VA, Probyn JP, Datey AV. Bridging the gap: Transitive associations between items presented in similar temporal contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2009; 35 :391–407.
- Howard MW, Kahana MJ. A distributed representation of temporal context. *Journal of Mathematical Psychology*. 2002; 46(3):269–299.
- Howard MW, MacDonald CJ, Tiganj Z, Shankar KH, Du Q, Hasselmo ME, Eichenbaum H. A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*. 2014; 34(13):4692–707.10.1523/JNEUROSCI.5808-12.2014 [PubMed: 24672015]
- Howard MW, Shankar KH, Aue W, Criss AH. A distributed representation of internal time. *Psychological Review*. (in press).

- Howard MW, Shankar KH, Jagadisan UKK. Constructing semantic representations from a gradually-changing representation of temporal context. *Topics in Cognitive Science*. 2011; 3 :48–73. [PubMed: 21552364]
- Howard MW, Viskontas IV, Shankar KH, Fried I. Ensembles of human mtl neurons “jump back in time” in response to a repeated stimulus. *Hippocampus*. 2012; 22(PMC3407826):1833–1847. [PubMed: 22488671]
- Howard MW, Youker TE, Venkatadass V. The persistence of memory: Contiguity effects across several minutes. *Psychonomic Bulletin & Review*. 2008; 15(PMC2493616):58–63. [PubMed: 18605480]
- Hsieh L-T, Gruber MJ, Jenkins LJ, Ranganath C. Hippocampal activity patterns carry information about objects in temporal context. *Neuron*. 2014; 81(5):1165–1178. [PubMed: 24607234]
- Humphreys MS, Bain JD, Pike R. Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*. 1989; 96 :208–233.
- Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences USA*. 2012; 109 :5086–91.10.1073/pnas.1114415109
- Kim J, Ghim J-W, Lee JH, Jung MW. Neural correlates of interval timing in rodent prefrontal cortex. *Journal of Neuroscience*. 2013; 33(34):13834–47.10.1523/JNEUROSCI.1443-13.201 [PubMed: 23966703]
- Komorowski RW, Manns JR, Eichenbaum H. Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *Journal of Neuroscience*. 2009; 29(31):9918–29.10.1523/JNEUROSCI.1378-09.2009 [PubMed: 19657042]
- Kraus BJ, Robinson RJ 2nd, White JA, Eichenbaum H, Hasselmo ME. Hippocampal “time cells”: time versus path integration. *Neuron*. 2013; 78(6):1090–101.10.1016/j.neuron.2013.04.015 [PubMed: 23707613]
- Landauer TK, Dumais ST. Solution to Plato’s problem : The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*. 1997; 104 :211–240.
- Lejeune H, Wearden JH. Scalar properties in animal timing: conformity and violations. *Quarterly Journal of Experimental Psychology*. 2006; 59(11):1875–908.
- Lever C, Burton S, Jeewajee A, O’Keefe J, Burgess N. Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience*. 2009; 29(31):9771–7. [PubMed: 19657030]
- Lewis PA, Miall RC. The precision of temporal judgement: milliseconds, many minutes, and beyond. *Philosophical Transactions of the Royal Society London B: Biological Sciences*. 2009; 364(1525):1897–905.10.1098/rstb.2009.0020
- MacDonald CJ, Carrow S, Place R, Eichenbaum H. Distinct hippocampal time cell sequences represent odor memories immobilized rats. *Journal of Neuroscience*. 2013; 33(36):14607–14616. [PubMed: 24005311]
- MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron*. 2011; 71 :737–749. [PubMed: 21867888]
- Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK. Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences*. 2012; 109 :19462–7.10.1073/pnas.1214107109
- Manning JR, Polyn SM, Litt B, Baltuch G, Kahana MJ. Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Science, USA*. 2011; 108(31):12893–7.
- Manns JR, Howard MW, Eichenbaum HB. Gradual changes in hippocampal activity support remembering the order of events. *Neuron*. 2007; 56(PMC2104541):530–540. [PubMed: 17988635]
- McKenzie S, Frank AJ, Kinsky NR, Porter B, Rivière PD, Eichenbaum H. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron*. 2014; 83(1):202–15.10.1016/j.neuron.2014.05.019 [PubMed: 24910078]

- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser M-B. Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*. 2006; 7(8):663–78.10.1038/nrn1932 [PubMed: 16858394]
- Messinger A, Squire LR, Zola SM, Albright TD. Neuronal representations of stimulus associations develop in the temporal lobe during learning. *Proceedings of the National Academy of Science, USA*. 2001; 98(21):12239–12244.
- Mishkin M, Ungerleider LG. Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behavioural Brain Research*. 1982; 6(1):57–77. [PubMed: 7126325]
- Miyashita Y. Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*. 1988; 335(6193):817–820. [PubMed: 3185711]
- Miyashita Y, Kameyama M, Hasegawa II, Fukushima T. Consolidation of visual associative long-term memory in the temporal cortex of primates. *Neurobiology of Learning and Memory*. 1998; 70(1/2):197–211. [PubMed: 9753597]
- Moser EI, Kropff E, Moser M-B. Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*. 2008; 31 :69–89.10.1146/an-nurev.neuro.31.061307.090723
- Murdock BB. A theory for the storage and retrieval of item and associative information. *Psychological Review*. 1982; 89 :609–626.
- Naya Y, Yoshida M, Miyashita Y. Backward spreading of memory-retrieval signal in the primate temporal cortex. *Science*. 2001; 291(5504):661–664. [PubMed: 11158679]
- 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. *The hippocampus as a cognitive map*. New York: Oxford University Press; 1978.
- Pastalkova E, Itskov V, Amarasingham A, Buzsaki G. Internally generated cell assembly sequences in the rat hippocampus. *Science*. 2008; 321(5894):1322–7. [PubMed: 18772431]
- Primoff E. Backward and forward associations as an organizing act in serial and in paired-associate learning. *Journal of Psychology*. 1938; 5 :375–395.
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I. Invariant visual representation by single neurons in the human brain. *Nature*. 2005; 435(7045):1102–7. [PubMed: 15973409]
- Raaijmakers, JGW.; Shiffrin, RM. SAM: A theory of probabilistic search of associative memory. In: Bower, GH., editor. *The psychology of learning and motivation: Advances in research and theory*. Vol. 14. New York: Academic Press; 1980. p. 207-262.
- Savastano HI, Miller RR. Time as content in Pavlovian conditioning. *Behavioural Processes*. 1998; 44 :147–162. [PubMed: 24896972]
- Schapiro AC, Kustner LV, Turk-Browne NB. Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*. 2012; 22 :1622–1627. [PubMed: 22885059]
- Schapiro AC, Rogers TT, Cordova NI, Turk-Browne NB, Botvinick MM. Neural representations of events arise from temporal community structure. *Nature Neuroscience*. 2013; 16(4):486–92.10.1038/nn.3331 [PubMed: 23416451]
- Scoville WB, Milner B. Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*. 1957; 20 :11–21.
- Shankar, KH. Generic construction of scale-invariantly coarse grained memory. 2014. arXiv: 1406.3185 [q-bio.NC]
- Shankar KH, Howard MW. A scale-invariant representation of time. *Neural Computation*. 2012; 24 : 134–193. [PubMed: 21919782]
- Shankar KH, Howard MW. Optimally fuzzy scale-free memory. *Journal of Machine Learning Research*. 2013; 14 :3753–3780.
- Shankar KH, Jagadisan UKK, Howard MW. Sequential learning using temporal context. *Journal of Mathematical Psychology*. 2009; 53 :474–485.
- Slamecka NJ. An analysis of double-function lists. *Memory & Cognition*. 1976; 4 :581–585. [PubMed: 21286984]

- Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI. Representation of geometric borders in the entorhinal cortex. *Science*. 2008; 322(5909):1865–8. [PubMed: 19095945]
- Suzuki WA, Amaral DG. Perirhinal and parahippocampal cortices of the macaque monkey: cortical afferents. *Journal of Comparative Neurology*. 1994; 350(4):497–533.10.1002/cne.903500402 [PubMed: 7890828]
- Tulving, E. *Elements of episodic memory*. New York: Oxford; 1983.
- Van Essen DC, Felleman DJ, DeYoe EA, Olavarria J, Knierim J. Modular and hierarchical organization of extrastriate visual cortex in the macaque monkey. *Cold Spring Harbor Symposium on Quantitative Biology*. 1990; 55 :679–96.
- Viskontas IV, Knowlton BJ, Steinmetz PN, Fried I. Differences in mnemonic processing by neurons in the human hippocampus and parahippocampal regions. *Journal of Cognitive Neuroscience*. 2006; 18(10):1654–62. [PubMed: 17014370]
- Wearden JH, Lejeune H. Scalar properties in human timing: conformity and violations. *Quarterly Journal of Experimental Psychology*. 2008; 61(4):569–87.
- Westmacott R, Moscovitch M. Names and words without meaning: incidental post-morbid semantic learning in a person with extensive bilateral medial temporal damage. *Neuropsychology*. 2001; 15(4):586–96. [PubMed: 11761048]
- Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science*. 1993; 261 :1055–8. [PubMed: 8351520]
- Yakovlev V, Fusi S, Berman E, Zohary E. Inter-trial neuronal activity in inferior temporal cortex: a putative vehicle to generate long-term visual association. *Nature Neuroscience*. 1998; 1(4):310–317. [PubMed: 10195165]
- Zeithamova D, Dominick AL, Preston AR. Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*. 2012; 75(1):168–179. [PubMed: 22794270]




Figure 1.

Firing of hippocampal cells must be controlled by events distant in space and time. Left: A place cell fires in a location within an environment with four landmarks, A B C D. It is unclear whether the cell fires because it is SE of A, NE of B, SW of C, NW of D, or any conjunction of those variables. In order to determine this, it would be necessary to systematically move the landmarks and establish control over the cell's firing. Right: A cell fires during a temporal sequence of events A B C D E. It is not clear whether the cell fires because D has just been presented, because it is one time step after C, or because it is two time steps after B. If the sequence is consistent and well-learned it is also possible that it is firing because E is predicted. In order to determine which of these relationships is responsible for the cell firing, it would be necessary to present the stimuli many times in many different orders to establish stimulus control of the firing.

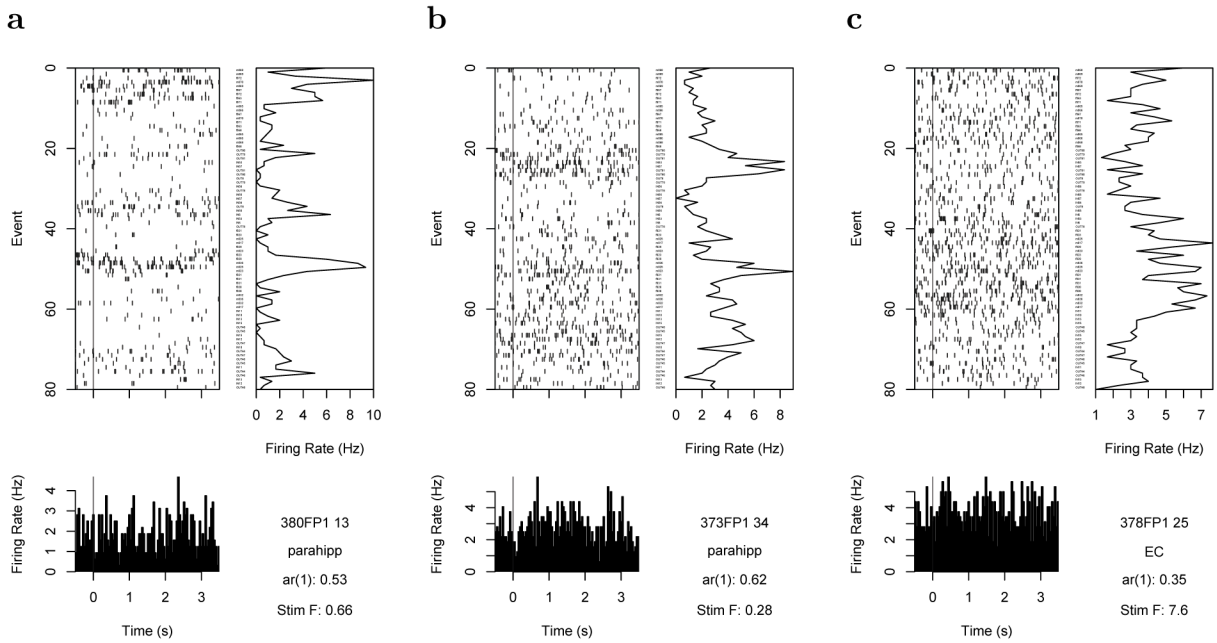


Figure 2.

Examples of neural firing correlated over long periods of time. Each panel gives the spikes fired over a recording session. Each row gives a different trial of a continuous recognition task, aligned to the presentation of the stimulus (vertical line). At the bottom a PSTH is shown. On the right of each panel, the cell's firing rate for each trial is shown. Note the robust autocorrelation—the firing of each cell changes gradually over many many trials. The entire experiment takes about 300 s. **a.** This cell shows autocorrelated firing over a few trials. Each trial is approximately 4 s, so this constitutes a change over about a dozen seconds. **b.** This cell shows autocorrelation over about 10 trials, corresponding to several dozen seconds. **c.** This cell shows autocorrelation over a few dozen trials, corresponding to a few minutes. After Howard, et al. (2012).




Figure 3.

A relational memory code in the hippocampus could rely on temporal memory representations. **a.** Neurons in the rodent hippocampus during the delay of an object memory task. The cells are sorted according to their peak of firing. Across cells, the time into the delay interval can be reconstructed. After MacDonald, et al., (2011). **b.** Non-spatial “splitter cells” retain information about past events. Rats were presented with sequences of odors with overlapping subsequences such as M N A B C O P and W X A B C Y Z. Some cells responded differentially to the odors in the subsequence depending on which sequence they were presented in. This cell fires more to B and C when they were preceded by the other odors of sequence 1 than when they were preceded by the other odors of sequence 2. After Ginther et al., (2011).

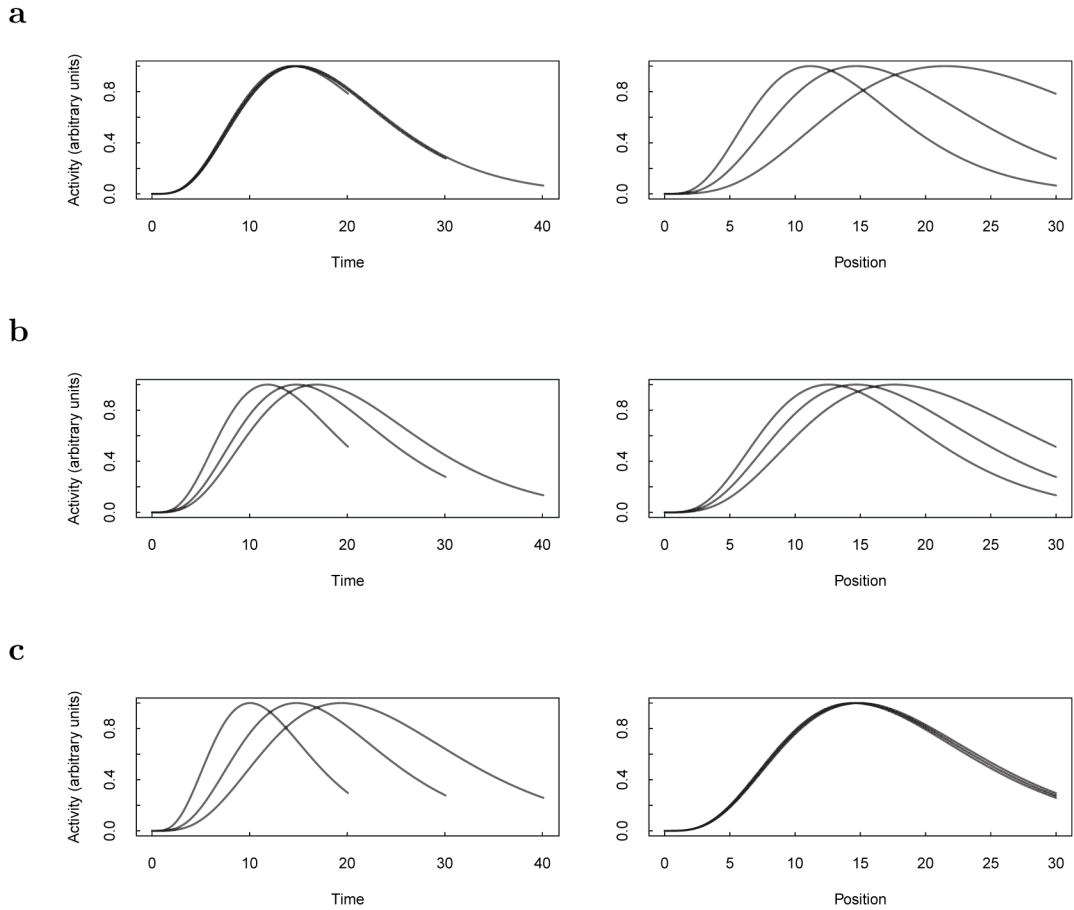


Figure 4.

A continuity between spatial and temporal coding. We simulated results from the Kraus et al., (2013) experiment using the computational framework introduced in Howard et al., (2014). In the simulated experiment, the rat ran on a treadmill at a constant speed (0.75, 1, or 1.5) each trial. The trials continued for different durations until 30 units of distance had been traversed. In each panel we plot the activity of the cell as a function of time (left) and as a function of distance traveled (right). **a.** A “time cell” that codes for time since the treadmill started. We set a close to 1 and b close to zero (with a slight offset so the different lines would be visible). Note that the response of the cell is almost completely controlled by time (left) whereas the activity as a function of position varies greatly depending on the speed (right). **b.** A cell showing coding intermediate between a temporal and spatial coding. Here $a = 0.5$ and $b = 0.5$. **c.** A “distance cell” that codes for elapsed distance since the treadmill started. We set a close to 0 and b close to 1 (with a slight offset so the different lines would be visible). Note that the response of the cell is almost completely controlled by position (right) whereas the activity as a function of time varies greatly depending on the speed (left).