THEORETICAL REVIEW

The Hippocampus, Time, and Memory Across Scales

Marc W. Howard and Howard Eichenbaum
Boston University

A wealth of experimental studies with animals have offered insights about how neural networks within the hippocampus support the temporal organization of memories. These studies have revealed the existence of "time cells" that encode moments in time, much as the well-known "place cells" map locations in space. Another line of work inspired by human behavioral studies suggests that episodic memories are mediated by a state of temporal context that changes gradually over long time scales, up to at least a few thousand seconds. In this view, the "mental time travel" hypothesized to support the experience of episodic memory corresponds to a "jump back in time" in which a previous state of temporal context is recovered. We suggest that these 2 sets of findings could be different facets of a representation of temporal history that maintains a record at the last few thousand seconds of experience. The ability to represent long time scales comes at the cost of discarding precise information about when a stimulus was experienced—this uncertainty becomes greater for events further in the past. We review recent computational work that describes a mechanism that could construct such a scale-invariant representation. Taken as a whole, this suggests the hippocampus plays its role in multiple aspects of cognition by representing events embedded in a general spatiotemporal context. The representation of internal time can be useful across nonhippocampal memory systems.

Keywords: time, hippocampus, scale-invariant memory

Objects fade out of consciousness slowly. If the present thought is of A B C D E F G, the next one will be of B C D E F G H, and the one after that of C D E F G H I—the lingerings of the past dropping successively away, and the incomings of the future making up the loss. These lingerings of old objects, these incomings of new, are the germs of memory and expectation, the retrospective and the prospective sense of time. They give that continuity to consciousness without which it could not be called a stream.

—William James, *The Principles of Psychology*

It has long been appreciated that time and memory are intimately related. Aristotle (trans. 1930), considering the temporal property of memory wrote, "Only those animals which perceive time remember, and the organ whereby they perceive time is also that whereby they remember." In more modern times, psychologists have long noted the temporal nature of episodic memory (Tulving, 1983, 1985). This relationship goes both ways, however. Not only does episodic memory depend on our perception of time, as the quote from James above illustrates, our perception of the passage of time is critically dependent on memory. If all we had access to was the present moment of perception, there is no possibility of constructing an ordered set of perceptions. "Now" must necessarily be later than "past"; perception of this relationship requires some sort of memory for past events. Quantitative behavioral models of time and memory have been developed in the fields of animal conditioning (e.g., Gallistel & Gibbon, 2000; Gibbon, 1977) and human list learning (e.g., G. D. Brown, Preece, & Hulme, 2000; G. D. Brown, Neath, & Chater, 2007). Our contention is that a satisfactory quantitative model of time and memory can only be constructed and evaluated if it is constrained by both neurophysiological and behavioral results.

Recent years have seen two lines of work that speak to the question of how the brain represents time and supports episodic memory. The first line of work, from neurobiology and cognitive neuroscience, suggests that the medial temporal lobe (MTL) operates on a representation that includes information about not only which stimuli have been experienced but also the time and order in which they were experienced. The second line of work, from cognitive psychology and neurobiology, argues for a neural representation that changes gradually over time scales ranging up to at least tens of minutes. Moreover, this gradually changing neural representation can be recovered, resulting in a jump back in time

that is accompanied by an internal experience of vivid recollection of a specific instance. These two lines of work suggest the possibility that they reflect a common computational mechanism—if a representation of the sequence leading up to the present moment could somehow change over very long periods of time and be recovered by remembered stimuli, then this would be sufficient to account for both classes of findings. However, the long time scales implied by episodic memory results raise serious computational challenges for obvious models of sequences. If representing a sequence of stimuli A B C depends on a chain of associations from A to B to C, how can one maintain a representation of a sequence that extends for tens of minutes? We address a solution to this challenge based on recent work in computational neuroscience after detailing the evidence resulting from each of these lines of work.

Temporal Organization in the MTL

A large and growing literature indicates a key role for the hippocampus in encoding and retrieving the order of events in sequence memories, in humans and animals and across a broad range of behavioral paradigms. This role is consistent with neural representations of the recent past that capture information about the order and timing of events. We can define two types of representation that emphasize different types of information. Consider the state of memory after presentation of stimulus A, followed 3 s later by stimulus B, which was followed 6 s later by stimulus C. An ordinal representation would retain information about the sequence of stimuli—C was more recent than B, which was more recent than A—but would discard the information about the durations of the delays between them. A temporal representation would retain the entire timeline, retaining both the stimuli and the gaps between them. Ordinal information can be extracted from a temporal representation, but the converse is not true. Many aspects of hippocampal processing suggest a role for an ordinal or sequential representation, whereas other evidence speaks directly to a role in intrinsically temporal processing.

Sequential Representation in the MTL

Several studies suggest that the hippocampus has a special role to play in maintaining and retrieving information about the order in which stimuli were experienced. Such evidence comes from individuals with selective hippocampal damage (Mayes et al., 2001; Spiers, Maguire, & Burgess, 2001) and functional magnetic resonance imaging (fMRI) studies (e.g., Ekstrom & Bookheimer, 2007; Kumaran & Maguire, 2006; Lehn et al., 2009; Tubridy & Davachi, 2011; see also Staresina & Davachi, 2006). Studies in animals have also explored their capacity for remembering the order of stimulus sequences and have investigated the role of the hippocampus in temporal order memory.

A particularly striking example (Fortin, Agster, & Eichenbaum, 2002; see also Kesner, Gilbert, & Barua, 2002) showed that rats could remember the order of once-presented sequences of odors and that this ability depended critically on the integrity of the hippocampus. On each trial animals were initially presented with a unique list of five odors. After study of the list, they were given pairs of nonadjacent odors from the list and rewarded for selecting the probe that was presented earlier in the list. Rats performed well

above chance on this test of temporal order memory, and selective hippocampal damage impaired their memory for order. For comparison, in separate blocks of trials, rats studied the lists of five odors but, in subsequent recognition tests, simply had to distinguish an odor that had appeared in the list from one that did not. On these recognition tests, rats with hippocampal damage performed as well as normal rats. Furthermore, rats with hippocampal damage, like control rats, showed better performance on memory for more recently experienced items in the recognition task (i.e., a recency effect). This suggests that the performance of normal rats in temporal order memory was not based solely on differences in the relative recency of the odors. Confirming this conclusion, Templer and Hampton (2013) recently showed definitively that memory for order in a similar test in monkeys is not based on differences in the strengths of earlier and later experienced items, or by list position, and they confirmed that accuracy in order judgments was improved by greater separation of items within a list.

Our everyday experiences often overlap in the contents of events, for example, when we walk through common streets on the way to different locations. A representation of a sequence is one way to address the associative interference that might otherwise result from simple associations between elements. Consistent with the hypothesis that it operates on a representation of the sequence of stimuli leading up to the present, the hippocampus plays an essential role in disambiguation of overlapping sequence memories. Human fMRI studies show that the hippocampus is preferentially activated by learning and retrieving overlapping sequences composed of a variety of different stimuli (T. I. Brown, Ross, Keller, Hasselmo, & Stern, 2010; Kumaran & Maguire, 2006; Ross, Brown, & Stern, 2009). Similarly, animal work has shown that rodents could disambiguate overlapping odor sequences posed as a series of stimulus choices, some of which involved identical items in two lists (Agster, Fortin, & Eichenbaum, 2002). Rats with hippocampal damage were impaired in making a critical choice just following the overlapping elements only when a delay was imposed between those choices. Thus, other brain systems can support reproduction of overlapping continuous sequences, but the hippocampus is required when the sequence is halted and then must be remembered to continue. Taken together, these studies show that the hippocampus plays a critical role in memory for sequential events and for disambiguation of overlapping sequences in humans and animals.

Hippocampal neural ensembles also disambiguate overlapping sequence memories, using different neuronal populations that represent each of the overlapping memories. This capacity has been demonstrated in several studies using variants of spatial alternation tasks in which rats traverse the "stem" of a T-maze then alternately choose between left and right turn arms (Wood, Dudchenko, Robitsek, & Eichenbaum, 2000). On each trial, as the rat traverses the stem of the maze that is common to both left- and right-turn paths, it must remember its last choice and maintain this information until it reaches the choice point. Several studies have observed that hippocampal place cells fire differentially as the rat passes through successive locations on the maze stem, depending on either the previous choice or the upcoming choice (Ainge, van der Meer, Langston, & Wood, 2007; Ferbinteanu & Shapiro, 2003; Frank, Brown, & Wilson, 2000; Wood et al., 2000; see Shapiro, Kennedy, & Ferbinteanu, 2006, for a review). Similarly, in the

nonspatial domain, hippocampal principal neurons fire differentially during the overlapping portions of odor sequences (Ginther, Walsh, & Ramus, 2011).

Temporal Representation in the MTL

The results reviewed up to now do not require an inherently temporal representation. For instance, an ordinal neural representation that discards information about the delay between stimuli would be sufficient to judge the relative order of two probes. Additional evidence is needed to establish that the hippocampus retains and utilizes inherently temporal information.

It has been known for some time that the hippocampus is required for trace conditioning (Solomon, Vander Schaaf, Thompson, & Weisz, 1986). In trace conditioning, a conditioned stimulus (CS) predicts an unconditioned stimulus (US) after a delay has passed. For instance, a tone CS might sound and then an air puff US would be presented several hundred milliseconds after the tone. Clark and Squire (1998) noted that acquisition of trace conditioning is associated with awareness of the contingency between the CS and the US in human amnesia patients. In trace fear conditioning, and delay fear conditioning with long delay intervals, hippocampal-lesioned animals are impaired at expressing any form of conditioned response (CR) in response to the CS (e.g., Bangasser, Waxler, Santollo, & Shors, 2006; Beylin et al., 2001). Critically, recent studies of the effect of dorsal hippocampal lesions on appetitive delay and trace conditioning show that even though lesioned animals can acquire a CR, its timing is disrupted in subtle ways. After training using a 40 s delay conditioning protocol, Tam and Bonardi (2012b) examined timing of the CR using a peak procedure. In the peak procedure, subjects are trained that after a CS the first response after a certain delay is rewarded. On peak trials, the response is omitted and the temporal pattern of when the subject responds is observed. The typical finding is that the responses peak around the time reward would be expected. Tam and Bonardi (2012b) found that animals with lesions to the dorsal hippocampus were less accurate in the times at which they responded, showing a peak response earlier than control animals with a wider temporal spread. Tam and Bonardi (2012a) observed similar results in appetitive trace conditioning; on trials where the reward was omitted, rodents with dorsal hippocampal lesions again expressed less accurate information about the time of the reward.
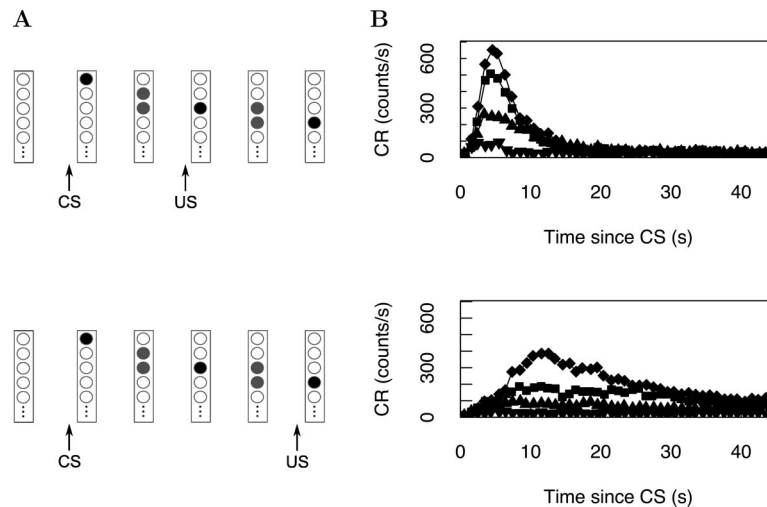
The demands of trace conditioning and the role of the hippocampus in supporting it suggest that the hippocampus should have access to a record of the recent temporal history. To respond at the correct time requires information about how long in the past the CS was presented. It is extremely challenging to account for the timing of the CR with a simple strength model in which cells are activated by the CS and then gradually decay in their firing rate. It is even more challenging to imagine an account of well-timed CRs using cells that are activated by the CS and then simply maintain their activity at the same high rate throughout the trace interval as expected in models of short-term memory. In contrast, well-timed CRs can be readily accounted for if the system has a set of cells that become sequentially activated following presentation of the CS. Figure 1 illustrates the utility of such a distributed representation of how far in the past the CS was presented in accounting for appropriately timed CRs.

Recent findings strongly suggest that the hippocampus has access to a distributed representation that contains information about when a delay interval started. Pastalkova, Itskov, Amarasingham, and Buzsaki (2008) had animals run laps along a linear maze, spending 15 s of each lap running in a wheel in a fixed location. Pastalkova et al. (2008) observed that different individual hippocampal neurons fired sequentially on the wheel, analogous to hippocampal place cells but during an interval where the animal's physical location is unchanged. Additional recent studies on hippocampal cells that encode time have been pursued in rats performing a variety of sequence memory tasks. In one of these tasks, rats learned associations between specific objects and odors separated by a 10 s interval; learning to bridge these sequential associations across a delay depends on the hippocampus (Kesner, Hunsaker, & Gilbert, 2005). Using this paradigm, MacDonald, Lepage, Eden, and Eichenbaum (2011) reported that individual hippocampal neurons fire for brief moments in sequence, composing a neural ensemble that filled the empty interval. They called these cells "time cells" (see Figure 2). Examining the simultaneously recorded time cells shown in Figure 2A, it is clear that this ensemble carries information about when the delay interval started. This would, in principle at least, be sufficient to account for results from a trace conditioning experiment. Because different neurons fire at different delays after the stimulus initiating the delay interval, these can be conditioned to different responses.

Several forms of evidence indicate that hippocampal neurons can represent time passing independent of their role in encoding space and other aspects of behavioral events. MacDonald et al. (2011) used general linear model analyses to show that these ensemble firing patterns reflected time passing even after accounting for variations in firing pattern due to the precise location of the rat and its head direction and speed. Time cells also bridge empty delays in a matching to sample task in head-fixed rats, showing time-related activity in the absence of even small variations in location, coincident or planned movement, or path integration (MacDonald, Place, & Eichenbaum, 2012). In addition, subsequent studies have provided compelling evidence that time cells integrate both time as well as distance traveled as rats run in place in a treadmill at different speeds (Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013). Finally, Naya and Suzuki (2011) described cells in the hippocampus of head-fixed monkeys as they studied a two stimulus list of visual stimuli. They reported that the hippocampal cells fired in response to the serial position of the stimuli without coding the identity of the stimulus itself and contained continuously varying information about the temporal interval between the presentation of the two list stimuli. Taken together, there is substantial and growing evidence that hippocampal neurons encode the flow of time in the service of the temporal organization of sequence memories.

Conjunctive Representation of What and When

In the real world, many different stimuli occur, signaling a variety of different temporal contingencies about future events. In order to predict the future correctly, we require not only information about when a stimulus occurred but also information about what stimulus was experienced. Not surprisingly, animals can show evidence for behavior that distinguishes information about what happened when. For instance, Roberts, Cheng, and Cohen

Figure 1.   Utility of a distributed timing signal. A: Schematic of a distributed timing signal as a function of time for two conditions in which a conditioned stimulus (CS) precedes an unconditioned stimulus (US) at two different delays. The state of a set of nodes is shown at various times during a trial of a trace conditioning experiment. Shading is meant to represent the activation of each node, with dark shading indicating that the node is more active. Top: Presentation of the CS starts the representation through a series of states in which different nodes are sequentially activated. The US is presented after a particular delay. The activated node can be associated to the node corresponding to an appropriate conditioned response (CR). When the CS is repeated, the CR is not predicted right away but only becomes activated after an appropriate delay. Bottom: When the delay before the US is longer, other nodes are activated. As a consequence, different nodes are conditioned to the CR than when the delay is shorter. B: Experimental data adapted from "Temporal Control of Conditioned Responding in Goldfish," by M. R. Drew, P. A. Couvillon, B. Zupan, A. Cooke, and P. Balsam, 2005, *The Journal of Neuroscience, 31,* p. 33. Copyright 2005 by the Society for Neuroscience. A certain time after a light CS is presented a shock US was administered to goldfish. The *y* axis gives the frequency of a CR during probe trials in which there was no shock. Different lines show responses after different numbers of learning trials (each block had 50 learning trials). The top plot shows performance when the delay was 5 s. The bottom plot shows performance when the delay was 10 s. Note that the time of peak responding varies according to the delay interval, even at early stages of learning.
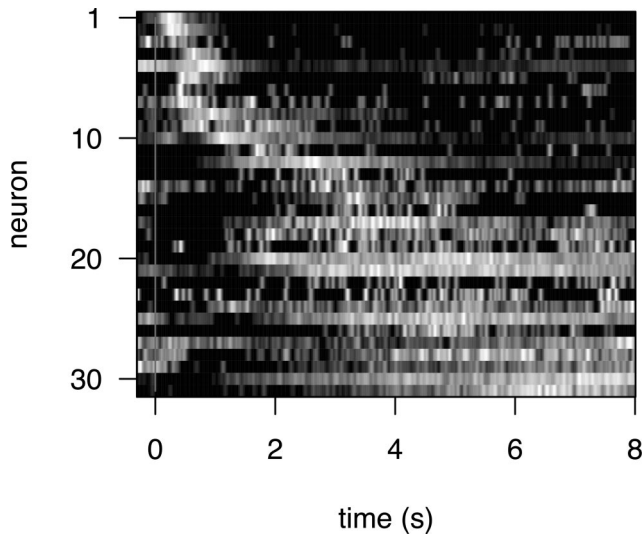
(1989) trained pigeons with two conditioned stimuli, a tone and a light. The first peck 15 s after presentation of CS1 resulted in the delivery of a food reward. In contrast, the first peck 30 s after presentation of CS2 resulted in food reward. After training, peak trials presented each CS for a long period of time without reward. Subjects showed a peak of responding at the time appropriate to the CS that was presented. That is, on peak trials when CS1 was presented, the CR was maximal around 15 s, whereas on peak trials when CS2 was presented, the CR was maximal around 30 s. Similar results are observed using rats as subjects (Yi, 2007). Referring to Figure 1, the results are just as one would expect if there were different nodes corresponding to each CS. That is, in order to account for behavior in this task the nodes in Figure 1 would have to be expanded into two columns of nodes. One column would correspond to a set of nodes that are sequentially activated by presentation of CS1, whereas the other column would be sequentially activated by presentation of CS2. That is, any particular node fires only when one of the stimuli has been experienced a certain time in the past. Taken together, a population of such nodes is a conjunctive representation of what and when information.

Consistent with what one would expect from a conjunctive representation of what happened when, hippocampal time cells are

sequentially activated and reflect previously experienced information. In the Pastalkova et al. (2008) study, the period on the treadmill was preceded (and followed) by two distinct movements. Different sequences of hippocampal cells were activated on the treadmill by these two sequences of movements. Gill, Mizumori, and Smith (2011) observed sequentially activated hippocampal neurons while animals waited on an elevated platform between trials on a plus maze. The trials on the plus maze were blocked; on a series of 15 trials the east arm was baited, whereas on the next 15 trials the west arm was baited. The sequence of neurons that fired on the platform between trials was different between the blocks of trials. While the results of these two studies are consistent with a conjunctive representation of what happened when, there are other possibilities.

MacDonald et al. (2012) examined hippocampal time cells in a delayed matching to sample task using odor stimuli. Because the animals were head-fixed, there is no possibility that any findings could be attributable to spatial or overt behavioral correlates. The animal's task was to sample an odor, wait for a 3 s delay and then determine whether a second odor was the same or different as the sample odor. Critically, distinct sequences of hippocampal neurons spanned the delay period when each of the four odors served as the sample odor. This result implies that the hippocampal ensemble

Figure 2. The firing rate of simultaneously recorded cells in dorsal CA1 as a function of time into the delay period of a memory task. Adapted from "Hippocampal 'Time Cells' Bridge the Gap in Memory for Discontiguous Events," by C. J. MacDonald, K. Q. Lepage, U. T. Eden, and H. Eichenbaum, 2011, *Neuron, 71,* p. 739. Copyright 2011 by Cell Press. Each row is the temporal profile of one cell averaged across trials; the cells have been ordered according to the time at which their firing peaks. CA = cornus ammoni.

carried information both about what odor was presented how far in the past.

Conventional Methods for Representing Sequences of What and When Information

While a variety of ideas and models have been suggested for representing sequential information, two classes of hypotheses for representing sequential and temporal information have received the bulk of the attention from hippocampal neurophysiologists: short-term memory buffers and synfire chains. The short-term memory buffer is a temporary store in which a fixed number of slots are available to hold recently presented information. From the pioneering work of Atkinson and Shiffrin (1968), short-term memory buffers have been used to account for behavioral results in tasks as diverse as verbal recall (Atkinson & Shiffrin, 1968; Raaijmakers & Shiffrin, 1980; Sirotin, Kimball, & Kahana, 2005) and statistical learning (Gureckis & Love, 2010). Short-term memory buffers have long been assumed to be interference-based (Waugh & Norman, 1965). That is, the contents of short-term memory are assumed to be unchanged by the passage of time per se, but displaced by incoming stimuli. When a new stimulus enters the buffer, an old item in the buffer is displaced. If the oldest item in short-term memory is displaced and the other items shift down one slot, then the buffer forms an ordinal representation that retains the order of recently presented stimuli. Noting the property of neurons to fire repeatedly due to an after-depolarization (ADP) Lisman, Idiart, and Jensen (Lisman & Idiart, 1995; Jensen & Lisman, 1996; Lisman & Jensen, 2013) proposed a short-term memory buffer based on brain oscillations. In this model, the slots of the buffer are neurons firing synchronously within a gamma cycle; the buffer is renewed each theta cycle by intrinsic firing properties. The capacity of the buffer is set by the number of gamma cycles that can fit into a theta cycle (see also Koene & Hasselmo, 2007).

Because they are designed to be sensitive to interference rather than time per se, short-term memory buffers are typically ill-suited to describe temporally dependent behavior. Synfire chains are an alternative computational hypothesis that can construct a distributed representation of what happened when (Hayon, Abeles, & Lehmann, 2005; Goldman, 2009). In synfire chains, an external stimulus causes some neurons to fire. These neurons activate other neurons, which activate other neurons and so on. Because it takes some time to traverse each link in the chain, the set of neurons active, say, 100 ms after the stimulus will tend to be different from the set of neurons activated 500 ms after the stimulus. Chains of sequentially activated neurons could in principle account for the existence of time cells (Itskov, Curto, Pastalkova, & Buzséki, 2011).

Unfortunately, both of these methods face significant challenges in scaling up. The buffer is constrained by the number of slots. Synfire chains can in principle be arbitrarily long, but in practice, this is a significant challenge. The number of links in the chain necessary to represent a particular delay goes up linearly with the duration of the longest interval to be maintained (Goldman, 2009).[1] Episodic memory shows similar temporal properties on the scale of seconds and hundreds of seconds. Similarly, later in the text we will see that neural representations in the MTL change gradually over at least thousands of seconds. Neither short-term memory buffers nor synfire chains seem well-suited for generating temporal or sequential representations over a wide variety of time scales.

Episodic Memory Across Time Scales

The world we inhabit contains structure at a wide variety of time scales (e.g., Voss & Clarke, 1975). Here we consider mnemonic and neural effects over cognitive time scales that can be assessed under controlled conditions in the laboratory, from 100 or 200 ms up to a few thousand seconds. From an evolutionary perspective, time-scale invariance is extremely adaptive. A memory buffer with a fixed capacity entails a huge risk. If a stimulus predicts a biologically important outcome at a lag that exceeds the fixed capacity of short-term memory, this could be disastrous. For decades, behavioral studies of a variety of memory tasks have suggested that memory persists and has similar properties over a wide range of time scales (Gibbon, 1977; Gallistel & Gibbon, 2000; G. D. Brown et al., 2000; G. D. Brown et al., 2007). This work has suggested that time scale invariance is a general principle that can be observed in a wide variety of tasks.

We review work from cognitive psychology suggesting that temporal effects in episodic memory tasks persist over periods of time up at least a few hundred seconds. In parallel, recent neurophysiological work has shown that neural representations in the medial temporal lobe change gradually over at least tens of minutes.
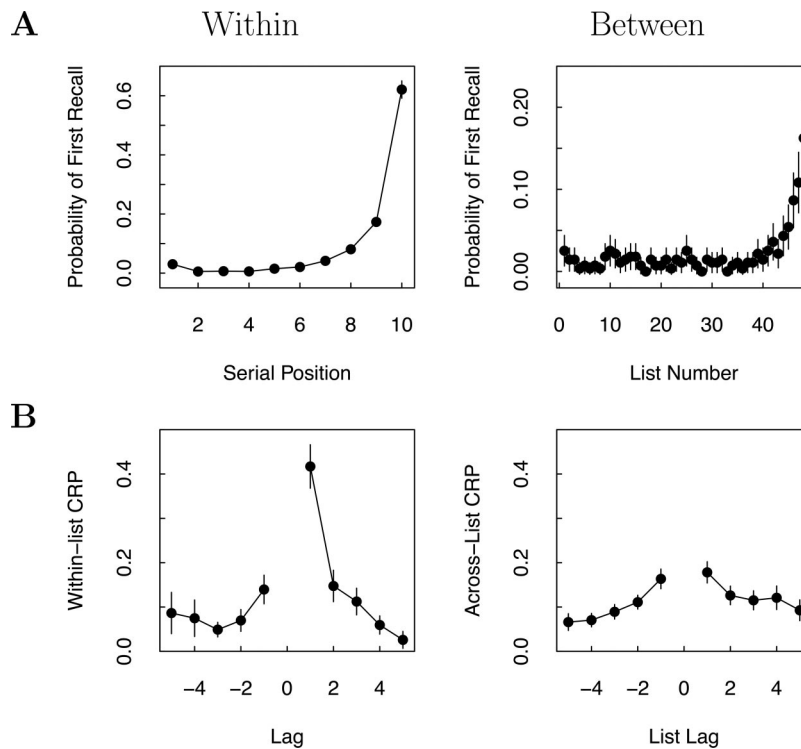
---

[1] Also, the law of large numbers makes it difficult to have Weber law variability in chains of sequentially activated neurons.

Recency and Contiguity Effects in Episodic Memory

Episodic memory has been suggested to reflect "mental time travel" in which the subject reexperiences a particular event (Tulving, 1983). Free recall is an experimental task used to study episodic memory in the laboratory. In free recall, the subject is presented with a random list of words, then asked to recall as many words as possible from the list in the order they come to mind. Free recall exhibits two phenomena that are believed to be general properties of episodic memory. The recency effect refers to the finding that, all other things being equal, recent words from the end of the list are better recalled than other words (Murdock, 1962; Bjork & Whitten, 1974; Glenberg et al., 1980). The contiguity effect refers to the finding that, all other things being equal, when a word from the list is recalled, the next word recalled will tend to

come from a nearby position within the list (Kahana, 1996; Howard & Kahana, 1999). Both recency and contiguity are quite general, being observed not only in free recall but in a wide variety of other laboratory memory experiments as well (see Kahana, Howard, & Polyn, 2008, for a review).

While recency and contiguity effects over short time scales could be attributed to short-term memory, a growing body of evidence suggests that recency and contiguity effects persist over much longer time scales, presenting a challenge for models based on a fixed capacity short-term memory. For instance, Howard, Youker, and Venkatadass (2008) presented subjects with lists of 10 words for immediate free recall. At the end of the session, subjects free-recalled all the words they could from all the lists. A recency effect was observed in both the immediate recall of one list as well

Figure 3. Recency and contiguity effects across time scales in final free recall. A: The recency effect within- and across-lists. Adapted from "The Persistence of Memory: Contiguity Effects Across Several Minutes," by M. W. Howard, T. E. Youker, and V. Venkatadass, 2008, *Psychonomic Bulletin & Review, 15,* p. 60. Copyright 2008 by Springer. In Howard et al. (2008), subjects performed immediate free recall on 48 lists of 10 words. At the end of the session, they were asked to recall all of the words in the entire experiment in the order they came to mind. The within-list curve gives the probability that the first word recalled during immediate recall of a particular list came from each serial position. There is a strong tendency recall words from the end of the list. The across-list curve gives the probability of first recall from the final free recall session as a function of *list.* Here there is again a recency effect favoring the recall of words from the last several *lists.* B: Contiguity effect within and across lists. Adapted from "Exploring the Retrieval Dynamics of Delayed and Final Free Recall: Further Evidence for Temporal-Contextual Search," by N. Unsworth, 2008, *Journal of Memory and Language, 59,* p. 230. Copyright 2006 by Elsevier. Subjects studied 10 lists of 10 words each (Unsworth, 2008). Lists were initially tested with delayed recall. During the final free recall session, subjects recalled many words from many lists. The within-list effect gives the probability of a final free recall transition between words from the same list as a function of the difference in serial position within that list. Given that a pair of successive recalls came from different lists, the across-list effect gives an estimate of the relative probability of making a transition from one list to another as a function of the distance *in lists* between those two words. CRP = conditional response probability.

as across lists (see Figure 3A, see also Bjork & Whitten, 1974; Glenberg et al., 1980; Howard & Kahana, 1999). Unsworth (2008) observed a similar persistence of the contiguity effect both within- and across-lists in final free recall (see Figure 3B, see also Howard & Kahana, 1999; Howard et al., 2008; Kılıç, Criss, & Howard, 2013). While it is of course possible to retain the idea of short-term memory by assuming that distinct mechanisms account for different time scales (Atkinson & Shiffrin, 1968; Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Sirotin et al., 2005; Lehman & Malmberg, 2013), the similarity of the recency and contiguity effects across scales are also consistent with a memory store that is not as tightly constrained as a short-term memory buffer.

Retrieved context models (Howard & Kahana, 2002; Dennis & Humphreys, 2001; Sederberg, Howard, & Kahana, 2008; Polyn, Norman, & Kahana, 2009) explain recency and contiguity effects in episodic memory as the result of a gradually changing state of temporal context. Like short-term memory, temporal context contains information about recently presented items. Rather than dropping out abruptly, however, information caused by each presented word decays gradually over time. Temporal context mediates recall of the items. Each word is encoded in the state of temporal context available when it was presented. The current state of context is used as a cue to initiate free recall. Because the state of temporal context changes gradually, a recency effect naturally results. In retrieved context accounts, the contiguity effect is a result of the ability of remembered items to recover the state of temporal context that obtained when they were originally encoded. Because this state resembles the encoding state for items nearby in the list, a contiguity effect naturally results. Unfortunately, although the equations of the temporal context model (TCM) do not decay as abruptly as short-term memory, they are not scale-invariant (Howard, 2004). Moreover, the equations describing the temporal context vector in TCM describe a decaying activation of recently presented stimuli rather than a conjunctive representation of what happened when, greatly limiting the model's utility in describing a variety of memory effects.

Recency and Contiguity in Neural Representations in the MTL

If neural representations in the MTL are responsible for episodic recall, then given that the recency effect extends over hundreds of seconds, one would expect neural representations in the MTL to change gradually over similarly long periods of time. There is a growing body of evidence that neural representations in the MTL change over thousands of seconds.
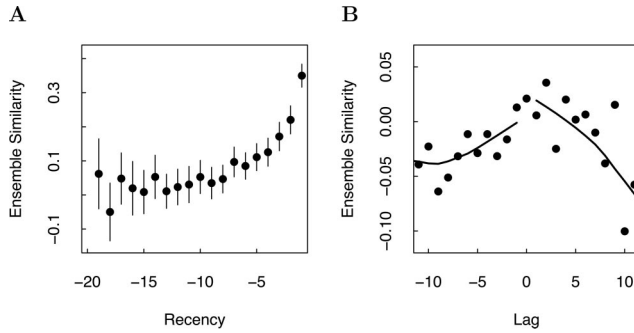
Manns, Howard, and Eichenbaum (2007) recorded from ensembles of CA1 neurons as rats performed a judgment of recency task wherein rats encoded and remembered unique sequences of odors (Fortin et al., 2002). Manns et al. (2007) measured population vectors for each study event and found that the neural representation in the hippocampus changed gradually within a list, correlating with performance on the judgment of recency task. That is, the similarity of the population vectors from the first and third odors within a list were more similar than the population vectors from the first and fifth odors with a list. Moreover, despite the fact that there was no requirement to remember the odors from one list to the next, Manns et al. (2007) also found that the ensemble changed

gradually *across lists.* That is, pairs of population vectors from events from nearby lists (e.g., List 5 vs. List 7) were closer to one another than pairs of events from distant lists (e.g., List 5 vs. List 15). Lists were separated by about a minute; there were reliable changes continuing even after several hundred seconds. Even in the absence of an explicit temporal memory task, other recent studies have observed that hippocampal ensembles change their activity gradually over thousands of seconds during spatial exploration (Hyman, Ma, Balaguer-Ballester, Durstewitz, & Seamans, 2012; Mankin et al., 2012). The MTL apparently maintains a representation that changes gradually over periods of time sufficiently long to account for the recency effect in episodic memory. Preliminary evidence also suggests that a gradually changing signal can be recovered by repetition of a stimulus, suggesting a jump back in time.

A recent study attempted to directly measure a jump back in time during a study of recognition memory. Subjects in the Howard, Viskontas, Shankar, and Fried (2012) study were epileptic patients who studied lists of pictures. In each block, the pictures in the list would be presented twice. The first time the picture was presented, the subject should respond "new"; the second time the picture was presented, the subject should respond "old." While the patients performed the task, extracellular recordings were taken from a variety of locations in the MTL. Howard et al. (2012) measured a population vector across the 3 s following presentation of each stimulus. In a first analysis, they computed a similarity measure between each pair of event vectors and aggregated the results as a function of the recency of the comparison. For instance, the similarity of the 8th event in a block to the 5th event in a block would go into the bin corresponding to recency −3. The results, in Figure 4A, showed a gradual change in the ensemble across time within a block, here over about 20 s.[2]

Repeating a list item allows us to discriminate among the many processes that could have caused the autocorrelation in Figure 4A. For simplicity, imagine that a list of items A B C D E was presented one at a time and then later C is repeated. Consider the similarity of the population vector following the repetition of C to the neighbors of C from the list. If the autocorrelation was a recording artifact, or otherwise unrelated to the stimuli being presented, then we would expect no correlation between the population vector when C is repeated and the population vectors during study of the list items. If the autocorrelation was a consequence of stimulus-specific persistent firing, we would expect the retrieved information to resemble the pattern after the initial presentation of C. Because the stimulus-specific activity persists, the pattern after repetition of C should resemble the pattern after presentation of D more so than the pattern after E, but not resemble the pattern after presentation of B. But what if repetition of C caused a "jump back in time" that would recover the state of temporal history available when C was initially studied? Under these circumstances we would expect the recovered information to resemble the pattern before study of C as well as after study of C. That is, the recovered population vector would resemble B and D more so than A or E, respectively.

---

[2] Howard et al. (2012) also observed gradual changes across blocks, extending several minutes.

Figure 4. Neural recency and contiguity effects. Adapted from "A Neural Signature of Mental Time Travel in the Human MTL," by M. W. Howard, I. V. Viskontas, K. H. Shankar, and I. Fried, 2012, *Hippocampus, 22,* p. 1839. Copyright 2012 by Wiley. Multiple neurons were recorded from human medial temporal lobe during performance of a continuous recognition task. A: The population vector during presentation of each stimulus was compared to the population vector during presentation of previous stimuli and aggregated as a function of recency between the two stimulus presentations. The ensemble state changed gradually over at least a few dozen seconds. B: When an item was repeated, that state was compared to the neighbors of the original presentation, subtracting out the contribution due to recency. The results suggest a jump back in time caused by repetition of the stimulus.

Figure 4B shows the results of a neural contiguity analysis on the same neurons. The results were consistent with our expectations for a "jump back in time," suggesting that the repeated item caused recovery of gradually changing information available before presentation of the repeated item. Several limitations of this study preclude a definitive conclusion. First, unit isolation from epilepsy patients is not nearly as reliable as can be achieved using tetrode recordings from animal preparations. Second, the design of the continuous recognition study meant that the repetition of C was not always at a long delay relative to the initial presentation of C. This necessitated removal of the recency effect via statistical methods.

While this story still has some gaps that need to be filled in, the overarching theme of this line of research is recognizable. A growing body of evidence suggests that episodic memory shows temporal effects—recency and contiguity—that persist over long periods of time, consistent with scale-invariance. Computational models predict that this is due to a gradually changing state of temporal context that serves as a cue for episodic recall; the contiguity effect is due to partial recovery of previous states of temporal context. MTL representations change over time scales long enough to account for behavioral recency and contiguity effects. Moreover, a temporally varying neural representation may be recovered by a repeated stimulus.

A Scale-Invariant Representation of Temporal History

Thus far we have reviewed two streams of research. The first suggests that the hippocampus is central in constructing and utilizing a representation that captures sequential and temporal information. The second argues for a representation of temporal context that changes over relatively long periods of time and that can be recovered in service of recency and contiguity effects in episodic

memory. These perspectives on the function of the MTL are not at all contradictory. A sequential or temporal representation could change gradually over time. If this representation could be recovered by a remembered stimulus, then this would result in a jump back in time. The challenge in unifying these two lines of research is to construct a sequential representation that exhibits the same properties over short time scales as over long time scales, from a few hundred milliseconds up to perhaps a few thousand seconds. However, we cannot take seriously a hypothesis that is impossible to physically realize. What is needed is a scale-invariant sequential representation that can be computed.

Both of the computational methods for representing sequences we have discussed thus far—a short-term memory buffer and a synfire chain—are ill-suited to scale up to long time scales. Consider a short-term memory buffer with a certain number of slots. How can the number of slots be simultaneously a few items, as required to account for immediate recall of a single list and also several dozen items, as required to account for recall of a set of lists (see Figure 3)? This problem is especially acute for biophysical models that ascribe the capacity of short-term memory to neurophysiological time constants (Lisman & Idiart, 1995; Jensen & Lisman, 1996). Similarly, consider the problems involved with scaling up a synfire chain to the time scale of thousands of seconds. The number of cells necessary to code for these long delays goes up linearly with the length of the longest delay, resulting in a significant demand on resources. Moreover, in a chain of sequentially activating cells, disruption to the activity of one node will propagate down the chain to a great many other nodes. While it is possible that there are satisfactory solutions to these problems within these extant models, a different computational paradigm that avoids these problems may be better suited for constructing a scale-invariant sequential representation.

To illustrate the problem in more detail and provide a sense of what a solution would look like, let us imagine how a scale-invariant memory buffer would appear. For simplicity, let's assume that stimuli come once every unit of time, A B C D E F G. Let's first consider how a short-term memory buffer would represent this sequence. After presentation of G, the first (most recent) slot in a short-term memory buffer would hold G, the second slot would hold F, the third would hold E and so on until all of the slots are filled. If a scale-invariant buffer could be constructed, the first (most recent) slot would hold G, but the second slot would hold E and F, the third slot would hold A B C and D and so on until all of the slots are filled. Because each slot holds information smeared out in time, events much further in the past can be represented with the same number of slots compared to the short-term memory buffer. The ability to represent longer time scales comes at the cost of discarding precise temporal information about stimuli further in the past. That is, because they are in the same slot, the difference in time between E and F is obscured relative to the difference between F and G, which are in different slots. This approach can be seen as change in the "width" of a slot. Rather than coding for the stimulus, say, four time steps in the past, the contents of a slot are more like a temporal average centered four steps in the past but smeared out in either direction.

But in the context of a buffer representation, this spread leads to a computational paradox. After presentation of G, the scale-invariant buffer has E and F in the second slot. Recall that precise temporal information about the stimuli within a slot has been

discarded. At the next time step, after presentation of H, the scale-invariant buffer should have F and G in the second slot. How can the buffer "know" that E should move to the third slot but F should not? Similarly, how can the buffer know that A should move out of the third slot, but B, C, and D should not? One could solve this problem by keeping a buffer with one slot for A, one slot for B, and so on. But this solution simply returns us to where we started, faced with the original challenge of constructing a short-term memory buffer that extends over very long time scales. Constructing a scale-invariant representation of temporal history with the properties we sketched above is a nontrivial computational challenge.

Recent computational cognitive neuroscience work describes a mechanism for computing a scale-invariant representation of temporal history leading up to the present moment (Shankar & Howard, 2012; Howard, Shankar, Aue, & Criss, 2013; Howard, MacDonald, et al., 2013; Shankar & Howard, 2013). The method relies on two populations of cells. The goal is to construct a population of cells that maintains a scale-invariant history analogous to the scale-invariant buffer sketched above. Different cells in this population code for different points in the history, analogous to the different slots in the buffer. But the population holding the record of the temporal history isn't constructed directly from the incoming stimuli. Rather, the stimuli drive a second population of cells that forms an intermediate step in constructing the history. In a buffer, incoming stimuli activate the node representing the slot corresponding to the most recent experience. In the intermediate representation, incoming stimuli send information to all the cells coding for that stimulus. It turns out that we can show that this intermediate representation contains a great deal of information about the history. We can extract the information from the intermediate representation using a set of fixed connection weights and record the estimate of history in the other population of cells. The resolution of the history that is extracted naturally gets worse as the cells in the intermediate representation become less and less active. We explain the details of this method in the next subsection. We explore the behavioral implications of a scale-invariant representation of temporal history in the following subsections.

Encoding and Extracting a Temporal History

A representation of temporal history maintains information about what happened when. As with other methods, such as a short-term memory buffer or a synfire chain, information about what stimulus was presented at a particular time is carried by which neurons fire. In a synfire chain, one stimulus sets off one chain, whereas another stimulus sets off a different chain. Similarly, the representation of temporal history has distinct populations of cells corresponding to each stimulus. In explaining the model, we first focus our attention on the calculation of when information and restrict our attention to one stimulus. In general, the representation consists of many sets of cells with similar temporal properties but receiving different inputs.[3] This property is essential for enabling behavioral models of a number of important phenomena. Neurophysiologically, this property corresponds to the stimulus-specificity of hippocampal time cells (MacDonald et al., 2012).

With our attention focused on one stimulus, our job is to reconstruct when in the past that stimulus was experienced. Let us

assume that at each moment in time we have access only to information about the current value of the stimulus. We refer to the function that describes whether the stimulus is present or not at each time τ as $f(\tau)$. For simplicity one can imagine that $f(\tau)$ is one for the moments when the stimulus is present and zero for all other times. The basic problem that needs to be solved in constructing a temporal history is to make past values of the stimulus function available at the present moment.
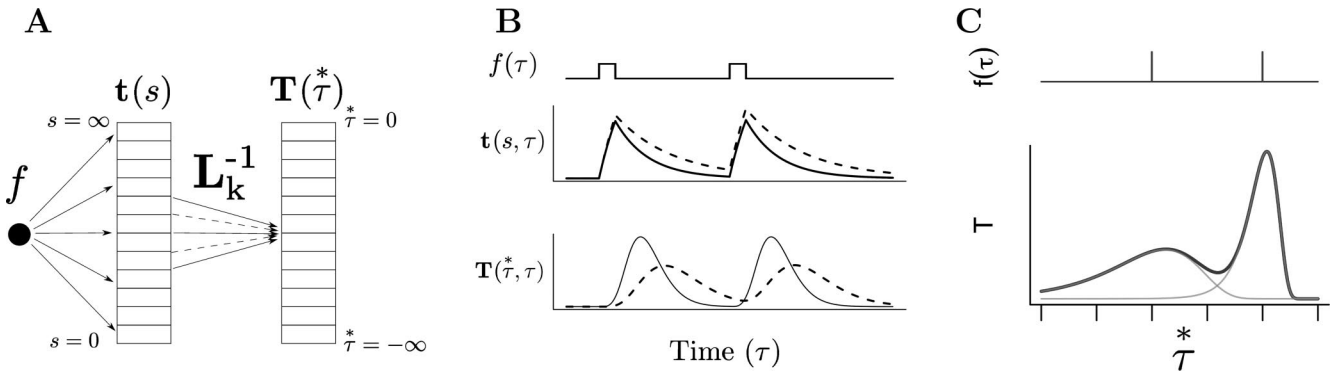
The method introduced by Shankar and Howard (2012) constructs a temporal history in two steps utilizing two sets of cells. The first set of cells, which serves as an intermediate step in constructing the temporal history, is referred to as $\mathbf{t}(s)$. Each of the cells in $\mathbf{t}(s)$ has a different value of s. The second set of cells, which will end up containing a scale-invariant representation of temporal history, is referred to as $\mathbf{T}(\overset{*}{\tau})$. Each of the cells in $\mathbf{T}(\overset{*}{\tau})$ has a different value of $\overset{*}{\tau}$ (see Figure 5). It will turn out that the different values of $\overset{*}{\tau}$ correspond roughly to the "slots" in the scale-invariant buffer described above—different values of $\overset{*}{\tau}$ correspond to different points in the past. At each moment in time, each cell in $\mathbf{t}(s)$ takes as input the current value of f. At each moment in time the scale-invariant history $\mathbf{T}(\overset{*}{\tau})$ is constructed from the current pattern across cells in $\mathbf{t}(s)$. Let us trace through the steps in more detail.

The firing rate of each cell in $\mathbf{t}(s)$ obeys the following differential equation that describes how the firing rate of each cell in $\mathbf{t}$ changes from one moment to the next.

$$\frac{d\mathbf{t}}{d\tau} = -s\mathbf{t} + f(\tau) \tag{1}$$

Let us take a moment and unpack what this equation tells us. First, the change in a cell's firing rate depends on the input at that moment, that cell's current firing rate, and that cell's particular value of s. Suppose that initially $\mathbf{t}$ is zero and $f(\tau)$ is also zero. Under these circumstances, the right hand side of Equation 1 is zero and $\mathbf{t}$ does not change. This will persist until the stimulus is presented leading to a positive value of $f(\tau)$. At this point the right hand side of Equation 1 will be positive and $\mathbf{t}$ will increase from zero. Note that this is true for all values of (s)—when the stimulus is presented, all of the cells in $\mathbf{t}(s)$ receive this input (see Figure 5A) and increase their firing at a rate controlled by their value of s, like a charging capacitor (see Figure 5B). Cells with small values of s respond more slowly, whereas cells with large values of s approach their asymptotic value more quickly. Now let us assume that a cell in $\mathbf{t}$ has reached its peak value and we turn off the stimulus such that $f(\tau)$ is zero. Now, because $\mathbf{t}$ is greater than zero, the right hand side of Equation 1 is negative, meaning that $\mathbf{t}$ decreases from its peak value (see Figure 5B). Note that as $\mathbf{t}$ gets smaller, so does the rate at which $\mathbf{t}$ changes. As $\mathbf{t}$ gets closer to zero, it slows more and more following an exponential function with rate parameter s. Cells with small values of s decay back toward zero very slowly; cells with larger values of s decay more quickly. The values of s that are available across cells will end up controlling the time scale that can be reconstructed in the history.

---

[3] More precisely, one can imagine an input vector $\mathbf{f}(\tau)$ that specifies the stimulus as a point in a space of appropriate dimensionality. We require a different population of cells to code the history for each dimension of $\mathbf{f}(\tau)$.

Figure 5. Schematics illustrating a computational hypothesis for a scale-invariant representation of temporal history. A: The cells in an intermediate representation **t** are driven by an input *f* describing the presence or absence of a particular stimulus at that moment. Adapted from "A Scale-Invariant Representation of Time," by K. H. Shankar and M. W. Howard, 2012, *Neural Computation, 24,* p. 145. Copyright 2012 by MIT Press. At each moment, only the current value of *f* is available. Each cell in **t**(*s*) has a different value of *s* that controls the rate at which that cell responds to an input. A reconstruction of temporal history **T** is constructed at each moment from **t**. Each cell in **T** is aligned with a paired cell in **t**(*s*). Each cell in **T** receives inputs from several cells in **t** adjacent to its paired cell. B: Adapted from "A Scale-Invariant Representation of Time," by K. H. Shankar and M. W. Howard, 2012, *Neural Computation, 24,* p. 143. Copyright 2012 by MIT Press. Top: The stimulus function *f* providing input to a set of cells is nonzero for two periods of time. Middle: Cells in the intermediate representation respond immediately to nonzero *f* like charging capacitors and then decay exponentially after the stimulus is removed. Different cells in the intermediate representation respond at different rates but starting at the same time. Bottom: Cells in the reconstruction **T** do not respond to the stimulus immediately, but after some characteristic delay. Different cells respond at different delays and with different temporal spreads. C: This plot shows a snapshot of the activation across nodes at one point in time after two presentations of the stimulus. Each presentation contributes to the representation. The representation of history is imperfect, with accuracy that decreases for events further in the past. Adapted from *A Quantitative Model of Time in Episodic Memory*, by M. W. Howard, K. H. Shankar, W. Aue, and A. H. Criss, 2013, manuscript submitted for publication.

It is clear that **t**(*s*) is not a temporal history. At any particular moment we cannot look at any one cell in **t**(*s*) and learn anything certain about any particular moment in the history.[4] However, the pattern of firing rates distributed across the cells in **t**(*s*) contains a tremendous amount of information about the stimulus history. As a matter of fact, it can be shown that if we had cells with all possible values of *s,* the pattern **t**(*s*) across cells would uniquely specify a single history function.[5] This means that if we had all values of *s* in our set of cells, the pattern of activity distributed across **t**(*s*) would contain *all* of the information about the stimulus history. The representation **T**($\overset{*}{\tau}$) is an attempt to recover the stimulus history hidden in **t**(*s*). Let us think for a moment about what it would mean if we could recover the stimulus history precisely.

As mentioned earlier, it will turn out that the representation **T**($\overset{*}{\tau}$) stores a blurry representation of the temporal history of *f*(*τ*) leading up to the present moment. Cells in **T**($\overset{*}{\tau}$) are each indexed by a value *τ**. Each value of $\overset{*}{\tau}$ is in register with a particular value of *s* (see Figure 5A). It will turn out that the value of $\overset{*}{\tau}$ corresponds to the part of the history that cell represents. If, at time *τ*, **T**($\overset{*}{\tau}$) was a perfect reconstruction of the stimulus function, then the firing rate of the cell with **T**($\overset{*}{\tau}$) = − 2 would be exactly equal to the value of the input 2 units of time in the past, *f*(*τ* − 2); the firing rate of the cell with $\overset{*}{\tau}$ = − 4 would be equal to the value of the input 4 units of time in the past, *f*(*τ* − 4). In this way, the history of all points in time *f*(*τ*′ < *τ*) would be copied into a spatial pattern of cells across the column of **T**($\overset{*}{\tau}$). Although the cells in **T**($\overset{*}{\tau}$) are sequen-

tially activated as time passes, they do not directly connect to one another.

Each cell in **T**($\overset{*}{\tau}$) has a particular value of $\overset{*}{\tau}$ and is aligned with a cell in **t**(*s*) with a corresponding value of *s*. At each moment, the firing rate of a target cell in **T** is constructed from the pattern of activity across the neighbors of its pair in **t**(*s*). The connection weights alternate between positive and negative and are balanced to sum to zero (for a more precise description, see Shankar & Howard, 2012). This form of connections estimates a higher order derivative across values of *s*. It can be shown that connection weights of this form result in a "smeared" approximation of the stimulus history.[6] Rather than a cell in **T**($\overset{*}{\tau}$) being equal to one value of *f*(*τ*) shifted in time, it is like an average of the values of *f* around a point in time. In this way, **T**($\overset{*}{\tau}$) of history exhibits

---

[4] To see why this is the case consider that there are many histories that can lead to the same firing rate. For instance, a particular nonzero firing rate for a cell could be caused by a small activation a short time in the past or a large activation a longer time in the past.

[5] More precisely, **t**(*s*) is the Laplace transform of *f*(*τ*) with real coefficients *s*. In principle, the Laplace transform of well-behaved function can be perfectly inverted to recover the original function *f*(*τ*).

[6] More precisely, the weights for the 2*k* − 1 near neighbors of the paired value of *s* approximate the *k*th derivative with respect to *s*. In the limit as *k* goes to infinity, it can be shown (Post, 1930) that the set of cells in **T**($\overset{*}{\tau}$) is proportional to the inverse Laplace transform of *f*(*τ*) leading up that moment (see Shankar & Howard, 2012, 2013, for details).

uncertainty in the time at which the stimulus was experienced. Moreover, the amount of this smear is larger for points further in the past (see Figure 5C). It can be shown that the width of the smear goes up linearly with the delay represented meaning that $\mathbf{T}(\overset{*}{\tau})$ is precisely scale-invariant (Shankar & Howard, 2012). Although the representation of the stimulus history is not precise, the error is distributed proportionally across all time scales for which values of s are available. Because each cell in $\mathbf{T}(\overset{*}{\tau})$ is constructed from a particular value of s and its neighbors, there is no requirement that the values of $\overset{*}{\tau}$ be evenly spaced. As a consequence, one can represent much longer delays with many fewer cells than would be possible with a shift register (Shankar & Howard, 2013).[7]

Ordinal and Spatial Representation

The current approach constructs a temporal history from an intermediate representation. The intermediate representation changes continuously in time, enabling the reconstruction to reflect the intervals between presentation of various stimuli. This framework is sufficiently general to accommodate an ordinal representation, in which only the sequence of stimuli are recorded discarding information about the intervals, and even a spatial representation that keeps track of the net distance traveled since a stimulus was encountered (Howard et al., 2013). These changes can be accomplished by altering the way in which the intermediate representation changes over time.

In constructing a temporal history, the intermediate representation changes continuously in time. If the intermediate representation stopped changing for some interval of time, so would the reconstruction of history. A scale-invariant ordinal representation can be constructed by only allowing the intermediate representation to change when a stimulus is presented. During periods when there is no stimulus, there is no change in the intermediate representation and thus no change in the reconstruction. The most recent stimulus only advances to the next "slot" when the next stimulus is presented. More formally, to implement a scale-invariant ordinal representation, we replace Equation 1 with the following generalization

$$\frac{d\mathbf{t}}{d\tau} = \alpha(\tau)[-s\mathbf{t} + f(\tau)]. \quad (2)$$

It is clear that Equation 1 is just Equation 2 with $\alpha(\tau)$ fixed at one. To construct a scale-invariant ordinal representation, we just set $\alpha(\tau)$ to be equal to the total magnitude of the input across stimuli—during times when there is no input $\alpha(\tau) = 0$ and there is no change in the intermediate representation (see Howard et al., 2013, for a more precise treatment).

It is well-known that as a rat moves around a spatial environment, neurons in the hippocampus code for the animal's allocentric position (e.g., O'Keefe & Dostrovsky, 1971; Wilson & McNaughton, 1993; Muller, 1996; O'Keefe & Burgess, 1996). It has long been suggested that the place code reflects a specific implementation of a more general computational function for representing spatiotemporal context (e.g., O'Keefe & Nadel, 1978; Eichenbaum, 1999; Hasselmo, 2012). The computational framework developed above can also accommodate a representation of location relative to environmental landmarks (Howard et al., 2013).

Consider the temporal history that develops as a subject moves with constant velocity, say driving along a road. At a particular moment, the subject encounters a billboard and continues onward. During the time after the billboard is encountered, the temporal history includes a representation of the billboard at successively more remote positions in time. Now, because the velocity is constant, during this interval time and position are perfectly correlated. The temporal representation of the billboard functions exactly the same as a scale-invariant spatial representation of location relative to the billboard would have. The temporal and spatial representations diverge if velocity changes. Suppose that the driver abruptly stops. While the driver remains stopped, the temporal representation keeps changing, whereas the spatial representation does not. Here the problem is analogous to the one faced in constructing an ordinal representation. If we can arrange for the intermediate representation to stop changing when the position stops changing, then the spatial representation of location relative to the billboard will stop changing as well. More generally, allowing the intermediate representation to change proportionally to the (signed) velocity in a particular direction enables the same algorithm used to construct a temporal history to also represent spatial location relative to fixed landmarks that are experienced in the environment. More formally, we can use Equation 2 to construct a scale-invariant representation of position relative to a landmark if $f(\tau)$ is only nonzero when the animal is in contact with a particular environmental landmark, such as the wall of an enclosure, and $\alpha(\tau)$ is equal to the velocity along a particular direction at each time step (see Howard et al., 2013, for details).

The mathematical framework described here is sufficiently general to accommodate construction of a temporal history, an ordinal representation, and even a spatial representation. What do these three modes of representing aspects of spatiotemporal context have in common? A temporal history provides information about relationships among stimuli that are not experienced simultaneously. Similarly, an ordinal representation carries information about the order in which stimuli are experienced while discarding information about the temporal intervals between them. A spatial representation reflects information about the spatial relationships between various stimuli experienced as the animal moves around an environment discarding both temporal information and sequential information. These three forms of representation can all be seen as complementary computational strategies to gather information about relationships between stimuli that are not experienced in the same place and at the same time. The intimate involvement of the hippocampus in temporal, ordinal, and spatial processing is consistent with a general function supporting relational memory (Cohen & Eichenbaum, 1993).

Evaluating Neurophysiological and Behavioral Predictions of the Model

As a computational hypothesis about the representation of stimulus history, this approach can be evaluated both by comparing it to neurophysiological findings, on the one hand, and behavioral

---

[7] It turns out that Weber-Fechner law spacing of the nodes is optimal in an information theoretic sense and also enables the description of several behavioral findings (see Howard et al., 2013; Shankar & Howard, 2013, for details).

results, on the other. We describe preliminary efforts to evaluate these comparisons as well as additional predictions for future empirical work that the model generates.

Comparison to hippocampal time cells. There are a number of striking similarities between the properties of hippocampal time cells (Kraus et al., 2013; MacDonald et al., 2012; MacDonald et al., 2011) and $\mathbf{T}(\overset{*}{\tau})$, which contains an estimate of the history leading up to the present moment. This linking hypothesis makes the following predictions about the properties of hippocampal time cells that have been empirically observed:
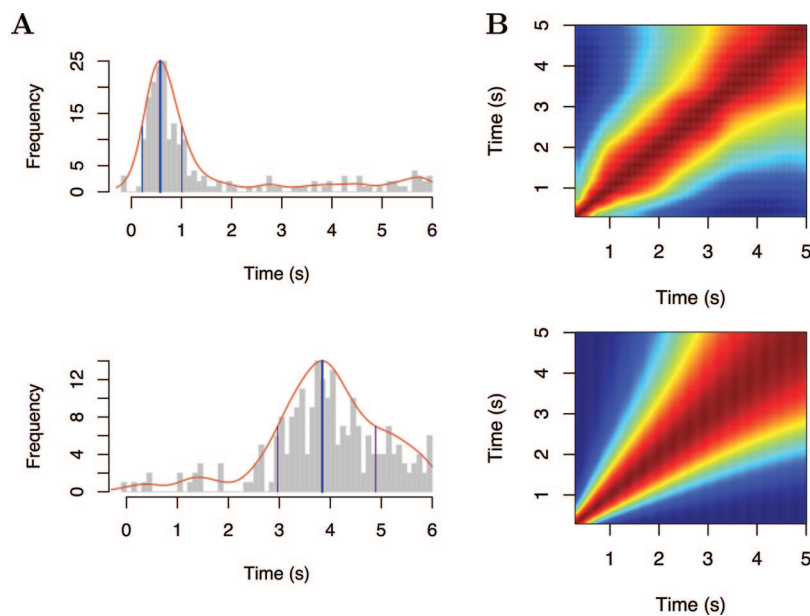
1. The width of time fields should expand with their time of peak firing.

2. Different stimuli should activate distinct sequences of time cells.

3. Time cells can respond to both time per se and self-motion information.

Prediction 1 (see Figure 6; Howard et al., 2013; Kraus et al., 2013) is necessary but not sufficient for mathematical scale-invariance. The computational model predicts the spread in the temporal width of time fields as a consequence of the decreasing accuracy for events that happened further in the past (see Figures 5B and 5C).

It is hard to overstate the importance of the spread in time fields in enabling the description of a wide variety of behavioral memory effects, including the recency effect and the decrease in order judgments with delay. Notably, the ensemble similarity of a set of time cells that did not have this property would not change gradually over a long time period. If time fields all had a fixed width, then the ensemble similarity would only fall off over the time scale corresponding to that width and then stop changing (Pastalkova et al., 2008). The observation of Prediction 1 keeps open the possibility that gradual changes in MTL firing patterns observed over long periods of time (Howard et al., 2012; Hyman et al., 2012; Manns et al., 2007) are a consequence of gradual changes in a population of time cells.
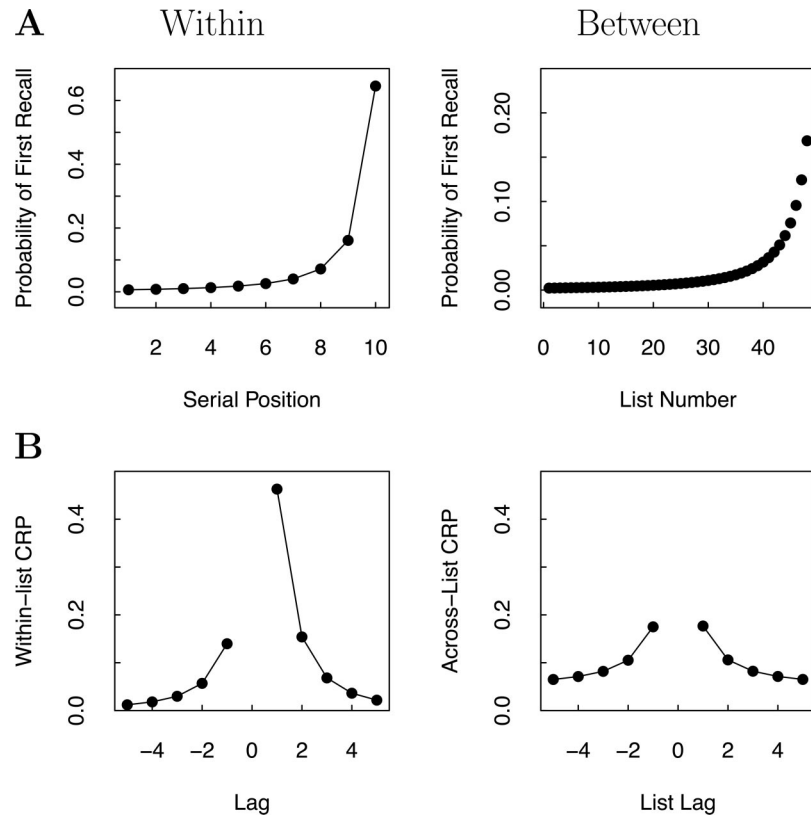
Prediction 2 (MacDonald et al., 2012) means that hippocampal time cells maintain a conjunctive record of what and when information. This capacity is extremely important behaviorally. This property is essential in order to model behavioral associations between stimuli that are separated in time. This property is manifest in behavioral models of trace conditioning with appropriately timed responses and in the asymmetric part of the contiguity effect in free recall (see Figure 7B).

Prediction 3 (Kraus et al., 2013) is essential if one is to use the algorithm for constructing a temporal history to provide a common computational account of hippocampal time cells and the place

Figure 6. Time cells show decreasing accuracy for times further in the past. From *A Unified Mathematical Framework for Coding Time, Space, and Sequences in the Medial Temporal Lobe* by M. W Howard, C. J. MacDonald, K. H. Shanker, Q. Du, M. E. Hasselmo, and H. Eichenbaum, 2013, p. 10. Copyright 2013 by M. W Howard, C. J. MacDonald, K. H. Shanker, Q. Du, M. E. Hasselmo, and H. Eichenbaum. A: The firing of two representative time cells as a function of time. In each plot, the total number of spikes fired in each bin is shown as a function of time synchronized to the onset of the delay period. The smooth red line gives an estimate of the density of spikes as function of time. The thick vertical blue line gives the estimate of the cell's mode; the two thin vertical blue lines give the estimate of the half-height region. The spread is statistically reliable across cells, as is the asymmetry in the shape of the time field. Compare to Figure 5B. B: Ensemble similarity during the delay period. The color scale gives the ensemble similarity (cosine of the angle between the vectors) of the smoothed population vectors across sessions for each pair of times during the delay. Top: Empirical data. Bottom: Ensemble similarity of **T** at different times.

A Within   Between

B

Figure 7. Recency and contiguity effects across scales using a scale-invariant representation of temporal history. Panels correspond to Figure 3. A: The recency effect was generated with a simple associative model between the current state of history and the list stimuli. The probability of choosing a stimulus to be recalled was generated using a power law softmax rule. From *A Quantitative Model of Time in Episodic Memory,* by M. W. Howard, K. H. Shankar, W. Aue, and A. H. Criss, 2013, p. 40. Copyright 2013 by M. W. Howard, K. H. Shankar, W. Aue, and A. H. Criss. B: The contiguity effect was generated by assuming that a recalled item causes partial recovery of the state of temporal history that obtained when it was studied. This recovered state overlapped with the state of history when nearby items were encoded, resulting in a contiguity effect. Because the representation of temporal history is scale-invariant, the model can simultaneously account for contiguity effects both within and between lists. See Howard et al. (2013) for details. From *A Quantitative Model of Time in Episodic Memory,* by M. W. Howard, K. H. Shankar, W. Aue, and A. H. Criss, 2013, p. 42. Copyright 2013 by M. W. Howard, K. H. Shankar, W. Aue, and A. H. Criss. CRP = conditional response probability.

code. Kraus et al. (2013) examined the firing of hippocampal cells while animals ran along a treadmill. The velocity of the treadmill was manipulated across trials such that time on the treadmill could be decoupled from the distance traveled. They found some cells that responded best to time per se, some cells that responded best to distance traveled, and some cells that showed sensitivity to both variables. This finding suggests that the computation that gives rise to time cells is part of a more general function for computing a spatiotemporal context, consistent with the view that episodic memory requires binding items with this general contextual information (e.g., Eichenbaum, Yonelinas, & Ranganath, 2007).

There are a number of additional predictions that follow from the mathematical formalism that have not yet been tested. Some of these predictions are quantitatively specific and relatively challenging to empirically evaluate. For instance, the spread in time fields should not merely increase with the delay that is being represented, the spread should be mathematically scale-invariant.

Similarly, the distribution of times of peak firing should exhibit Weber-Fechner law scaling. Other predictions are more qualitative. For instance, thus far time cells have been observed spanning all of the delays that have thus far been examined, up to about 15 s. Our hypothesis predicts that time cells should be observed over even longer time scales, extending to hundreds, and perhaps even a few thousand, seconds.

There are other ways to construct a conjunctive representation of what happened when (e.g., Grossberg & Merrill, 1992, 1996; Levy, 1996; Wu & Levy, 2001). The present computational hypothesis for how to construct a scale-invariant representation of history makes a very specific set of predictions about the intermediate representation. The model predicts that there should be populations of stimulus-specific cells that show exponentially decaying persistent firing (see Figure 5B). According to the hypothesis, these cells should have a variety of rate constants and provide a necessary precursor to the construction of hippocampal time

cells. While examination of some previous studies are suggestive of some of these properties (Fuster & Jervey, 1982; Naya & Suzuki, 2011; Yakovlev, Fusi, Berman, & Zohary, 1998), this topic has not been systematically investigated. Consideration of MTL anatomy and functional specificity suggests that the lateral entorhinal cortex (Eichenbaum et al., 2007; Kerr, Agster, Furtak, & Burwell, 2007; van Strien, Cappaert, & Witter, 2009) is a strong candidate for a population of exponentially decaying stimulus-specific cells.

Behavioral findings across time scales. This hypothesis for constructing a scale-invariant temporal history corresponds well to neurophysiological results from hippocampal time cells. The properties of the representation also make it extremely well-suited for describing behavioral effects as well. Here we discuss two classes of behavioral findings from tasks believed to depend on the hippocampus, trace conditioning and recency and contiguity effects from episodic memory.

One of the primary functions of memory is to learn and exploit temporal relationships between stimuli available in the environment. Conditioning experiments can be seen as a highly refined test of this ability. Previously unrelated stimuli can be paired; behavioral expression of the expectation of a US can be observed even in nonverbal animals. In episodic memory experiments, human subjects study a list of randomly chosen stimuli and learn that they go together on the basis of a single learning episode. In order to account for behaviorally observed associations between stimuli, we can construct a simple behavioral model in which states of history mediate the relationships between nominal stimuli. We form an association between a study item and the state of temporal history present when it is experienced. To the extent that state of history is recreated, by whatever means, that stimulus will be predicted to occur. Associations between a cue and a response are a consequence of the effect the cue has on the state of history and then the efficacy of that state of history as a probe for recovery of the response.

Because a stimulus causes a series of distinct internal states, this model of temporal history can support behavior that discriminates different temporal intervals. Shankar and Howard (2012) showed that an associative model utilizing temporal history can readily account for the appropriate timing of CRs learned after varying delays (Drew, Couvillon, Zupan, Cooke, & Balsam, 2005, see Figure 1). The modeled CRs peaked at different times for animals trained with different delays. Moreover, they grew stronger with increased learning trials as the association between the state of temporal history and the CR grew progressively stronger. Moreover, because the sequence of states is caused by a particular stimulus, generating a conjunctive representation of what happened when, there is no conceptual challenge in associating different stimuli to different appropriately timed responses (Roberts et al., 1989). Because the representation of stimulus history is mathematically scale-invariant, it is straightforward to account for scale-invariant behavior in conditioning tasks (Balsam & Gallistel, 2009; Gallistel & Gibbon, 2000).

All other things being equal, a stimulus that has been recently experienced is more likely to recur than one that has not. A scale-invariant recency effect is an adaptive response to a world with structure at a wide variety of time scales (Anderson & Schooler, 1991). Using the scale-invariant representation of temporal history in the role of temporal context allows a scale-

invariant account of recency and contiguity effects (see Figure 7A, Howard et al., 2013). The recency effect follows from the gradually changing nature of the temporal history; the contiguity effect is a consequence of a "jump back in time" in which a previous state of history is recovered.

According to this behavioral model, each stimulus in the list is bound to the stimulus history available when it was presented. This history combines what and when information about all the other stimuli leading up to presentation of that particular word in the list. At recall, the current state of history is used as a cue to probe each of the list items. One of the list items is probabilistically selected based on its activation and the activation of the other candidates for recall. Because the state of history changes gradually, the current state of history is more similar to the encoding state of recently presented words, and they are more strongly cued, leading to a behavioral recency effect. Because the change over time is scale-invariant, this representation of internal time naturally predicts that the recency effect should also be scale invariant (see Howard et al., 2013, for details).

If remembering a list item also causes a "jump back in time" such that the state of history when that item was initially studied is partially recovered, then it is also straightforward to account for a behavioral contiguity effect (see Figure 7B). Consider a list A B C D E. When C is remembered, it recovers the state of history when it was initially presented. Because this state was part of a gradual flow of time during study of the list, it is similar to both the preceding and following states of history, serving as an effective cue for both backward and forward associations. Because the representation changes gradually across all time scales, there is no difficulty in describing contiguity effects both within- and across-lists. According to the model, the asymmetry arises because C also serves as an input to the state of history, just as it did when it was initially presented. The history when D was studied includes C one time step in the past; in contrast, when B was studied C was not part of the history. The result is an asymmetric association. This account also naturally explains the disappearance of the asymmetry in the across-list condition. This is because the time scale over which recalls take place is similar to the time scale at which words are studied but much faster than the time scale at which lists are studied. That is, if the words are studied for 1 s each, then the encoding history for D includes C at a delay of 1 s in the past. In contrast, the history for a word from List 10 includes words from List 9 about a minute in the past. When a word is recovered during retrieval, the delay to the next recall is much closer to a second than a minute. As a consequence, this asymmetric cue is much more effective within-list than across-lists (see also Howard et al., 2008).

We've shown that a computational mechanism for recovering a previous state of history—jumping back in time—can provide a mechanistic account of the contiguity effect in episodic recall tasks. This same computational mechanism may not only be manifest in episodic recall tasks but may play an important role in recognition memory as well. In recognition memory, subjects are tested for their memory for probe stimuli, some of which have been previously experienced and some of which are new. It has long been proposed that there are two distinct processes that cause s subjects to endorse previously experienced probes as old (Atkinson & Juola, 1974; Mandler, 1980; Tulving, 1985). One of these processes, recollection, is a recall-like, vivid reexperience of some

particular details of the study event. The other process is referred to as familiarity. There is not a consensus about whether it makes sense to think of them as distinct processes rather than points on a continuum (e.g., Dunn, 2008; Fortin, Wright, & Eichenbaum, 2004; Staresina, Fell, Dunn, Axmacher, & Henson, 2013; Wixted & Squire, 2011), and there are serious issues involved with measuring these two processes (Onyper, Zhang, & Howard, 2010; Province & Rouder, 2012; Wixted, 2007; Yonelinas & Parks, 2007). Nonetheless there can be little question that recognition is supported by sources of information that vary in their time course of retrieval (e.g., Hintzman & Curran, 1994; McElree, Dolan, & Jacoby, 1999), their electrophysiological correlates (e.g., Davachi, Mitchell, & Wagner, 2003; Rugg & Yonelinas, 2003; Staresina et al., 2013; Uncapher, Otten, & Rugg, 2006), and the ability to recover details about the context in which the probe was experienced (e.g., Meiser, Sattler, & Weisser, 2008; Onyper et al., 2010; Starns & Hicks, 2005).

Eichenbaum et al. (2007) argued that the hippocampus causes recollection by enabling reconstruction of the encoding context in which a probe item was originally presented. This contextual representation was argued to be present in extrahippocampal MTL regions. This function, which may depend on pattern completion in hippocampal region CA3 (e.g., Hasselmo & Wyble, 1997, Marr, 1971; McClelland, McNaughton, & O'Reilly, 1995; McNaughton & Morris, 1987; Norman & O'Reilly, 2003), enables the reconstruction of a complete state of activation across entorhinal cortex. This seems similar to the computational demands of the reconstruction of a previous state of temporal history—given a repeated item, the remainder of the neural state available when that probe was presented is reconstructed. What, then, is the relationship between recollection believed to support recognition and the jump back in time believed to support the contiguity effect in recall?

We argue that recollection and the jump back in time are reflections of the same neural phenomenon. Indeed recollection in item recognition is associated with a behavioral contiguity effect like that expected from a jump back in time. Schwartz, Howard, Jing, and Kahana (2005) presented subjects a series of travel pictures were presented during study. At test, subjects rated a series of probes. When successive old probes were from nearby positions within the study list, there was enhanced memory for the second probe. However, this contiguity effect was only observed when the first probe received a highest-confidence old item rating, believed to be preferentially supported by recollection. We propose that recollection is precisely the recovery of a previous state of spatiotemporal context given an item as a probe and that recollection and recall are different ways of describing the same kind of reconstruction event.

Timing Beyond the MTL

Thus far we have focused on the role of the hippocampus in temporal representation and temporal memory. Of course, there is a tremendous amount of evidence for contributions to timing and memory that are not strictly dependent on the hippocampus but rely on a variety of regions including prefrontal cortex (PFC), basal ganglia, and the cerebellum (e.g., Mangels, Ivry, & Shimizu, 1998; Matell & Meck, 2004). Moreover, a growing body of evidence points to interactions between brain regions as essential for learning and memory tasks that more closely resemble real-world

levels of complexity. For instance, in real-world situations subjects incorporate feedback about their choices over various time scales to correctly assign value to the various options they will face in the future. Accordingly, the MTL and reward systems in the striatum seem to interact in a relatively subtle way to enable learning about feedback and generalization about the value of stimuli (Bornstein & Daw, 2012; Foerde, Race, Verfaellie, & Shohamy, 2013; Foerde & Shohamy, 2011; Wimmer & Shohamy, 2012). Similarly, in the relative judgment of recency (JOR) task, subjects must retrieve temporal information about a pair of probe stimuli and then compare and evaluate the retrieved information to make a task-appropriate judgment. In the JOR task, connections between PFC and the MTL appear to be extremely important in successful performance of order judgments on a pair of probe stimuli (e.g., Barker, Bird, Alexander, & Warburton, 2007; DeVito & Eichenbaum, 2011; Jenkins & Ranganath, 2010).

A scale-invariant representation of conjunctive what and when information could be tapped in a variety of timing and memory tasks, not just those dependent on the MTL, to support performance across a wide variety of time scales in task-appropriate ways. As an example, we discuss how directly accessing a scale-invariant representation of conjunctive what and when information can be used to construct a behavioral model of judgments of recency across scales. Then we discuss evidence that neural representations in regions outside the MTL show gradual changes over time scales up to a few thousand seconds and that brain regions outside the MTL have access to a conjunctive representation of what and when. Broadly speaking, our contention is that the algorithm we have described constructs a scale-invariant representation of conjunctive what and when information that is available to many brain regions that exploit temporal organization in support of a variety of behaviors.

Judgments of recency across time scales. In the applications to free recall and conditioning described above, we treated the representation of stimulus history as a form of temporal context mediating associations between stimuli. But this is not the only way this representation could support behavior. In principle, one could directly examine the current state of the temporal history to answer questions about the history of a probe stimulus. In the context of a buffer model, one could account for the relative JOR task by examining the contents of one slot at a time and comparing the contents of each slot to the probe stimuli, stopping when a match is found. We can construct a similar behavioral model using the scale-invariant representation of temporal history in place of a short-term memory buffer. Here, the "slots" correspond to rows of $\mathbf{T}(\overset{*}{\tau})$ with values of $\overset{*}{\tau}$ successively further in the past. This form of sequential access to various points in the past would not be possible if all that memory contained was a decaying trace of past events (Hinrichs, 1970; Howard & Kahana, 2002).

Behavioral data from JORs over short time scales strongly support the idea of sequential organization of memory. Hacker (1980) presented subjects with a series of consonants one at a time sufficiently rapidly that the entire list took only a few seconds. After a mask was presented, subjects were asked to choose which one of two probe stimuli was presented more recently. Not surprisingly, accuracy was better for more recent stimuli and for pairs of probes that were widely separated in the list. Remarkably, however, the reaction time (RT) on trials when the subject correctly chose the more recent probe depended on the recency of the

more recent stimulus but not on the recency of the less recent probe. In contrast, when the subject incorrectly chose the less recent probe the RT depended on the recency to the less recent probe but not on the recency of the more recent probe. These results, which have been replicated and extended many times (Hockley, 1984; McElree & Dosher, 1993; Muter, 1979), are exactly what one would expect if the subject were sequentially examining a representation of the recent past. This interpretation requires, of course, that a representation of the recent past ordered along an internal timeline exists and can be purposefully accessed. Consistent with this interpretation, Chan, Ross, Earle, and Caplan (2009) found that when subjects were asked to report the probe that occurred first in the list, they found RT results consistent with serial scanning from the beginning of the list.

A simple behavioral scanning model using $\mathbf{T}(\overset{*}{\tau})$ can readily account for the behavioral results from the Hacker (1980) experiment (Howard et al., 2013). This is not surprising as sequential scanning models have been shown by several authors to account well for those results (e.g., G. D. Brown et al., 2000; Hacker, 1980). More notably, the same serial scanning model also accounts for a number of JOR phenomena across much longer scales. For instance, Yntema and Trask (1963) introduced the JOR task in an experiment in which subjects had to choose which of two words was presented most recently in a long continuous stream of stimuli that lasted for many minutes. The recencies judged ranged from a few seconds (similar to the duration of the entire list in the Hacker, 1980, experiment) up to about a hundred seconds. Howard et al. (2013) showed a satisfactory description of the (Yntema & Trask, 1963) data using the same behavioral model that was used to describe the much shorter time scales in the Hacker (1980) experiment. Thus far, reaction times have not been reported for JORs on the time scale of the Yntema and Trask (1963) experiment. If JORs are based on a scale-invariant representation of history, then one might expect to see scanning-like RT results over that time scale as well.

Temporal organization of neural representations outside of the MTL. In animals, the JOR task depends on connectivity between the PFC and the MTL (Barker et al., 2007; DeVito & Eichenbaum, 2011). Neuropsychological studies also confirm that the PFC is important in judgments of temporal order (McAndrews & Milner, 1991; Milner, Corsi, & Leonard, 1991; Shimamura, Janowsky, & Squire, 1990). Using fMRI, Jenkins and Ranganath (2010) observed that the pattern of activation in several subregions of the PFC changed gradually across presentation of multiple study items. In particular, changes in the left rostrolateral PFC correlated with performance on a subsequent JOR task. In addition to these results showing changes in PFC activation over the scale of a few seconds, Hyman et al. (2012) observed gradual changes in population vectors of rodent PFC neurons across thousands of seconds during spatial exploration. These results suggest that like the MTL, the PFC also has access to neural representations that change gradually over time scales both on the order of a few seconds and also on the order of thousands of seconds.

Neural correlations have also been observed over relatively long periods of time in regions outside of the MTL and the PFC. Smith and Sommer (2013) recently observed autocorrelations in the responsiveness of V4 neurons to moving gratings over at least a few minutes. Learning studies suggest that a variety of brain regions respond to the history of reward and temporal anticipations of reward. Bernacchia, Seo, Lee, and Wang (2011) recorded from monkeys and found neurons in a variety of regions (anterior cingulate, dorsolateral PFC, lateral intraparietal cortex) that had sensitivity to the past sequence of rewards with different time constants. Notably, the distribution across neurons of time constants was power-law-distributed, consistent with a scale-invariant record of past reward. In addition to sensitivity to past rewards, it is also possible to measure sensitivity to expected future rewards at different time points. For instance, Tanaka et al. (2004) correlated blood oxygenated level dependent (BOLD) response with the amount of reward expected at various time points in the future as predicted by reinforcement learning models with different parameters controlling the rate of temporal discounting. They saw a continuum of regions within the insula and striatum responding to different points in future time.

A remarkable recent series of studies suggests that the striatum contains a conjunctive representation of what and when information strikingly similar to hippocampal time cells. Adler et al. (2012) recorded from striatal neurons while monkeys performed a simple Pavlovian conditioning task. One of several visual stimuli was presented. After a 2 s delay, an outcome corresponding to the stimulus was presented. Adler et al. reported that medium spiny neurons in the putamen were sequentially activated for a circumscribed period of time during the delay. These time cells appear to have specificity to the outcome that will occur (and thus the stimulus that started the interval) and to span the entire delay interval, and cells that peaked later in the delay interval appeared to show an increased spread in firing. Subsequent work has shown that this remarkable correspondence with hippocampal time cells apparently holds for medium spiny neurons not only in the putamen but also in the caudate and the ventral striatum (Adler, Finkes, Katabi, Prut, & Bergman, 2013). Notably, this follow-up study reported that the time cells continue after presentation of the outcome into the intertrial interval (up to 6 s after initial presentation of the stimulus), arguing against the possibility that they are simply coding for the upcoming outcome. These studies raise the extremely intriguing possibility that both the hippocampus and the striatum have access to a conjunctive representation of what happened when that grows less accurate for events further in the past. If this type of representation can be observed in the hippocampus and striatum, two regions that are relatively widely separated, it is possible that similar findings could also be observed many other brain regions as well.

Conclusions

We have reviewed two streams of research suggesting on the one hand that the hippocampus maintains order information about recent stimuli and on the other that episodic memory depends on neural representations that change gradually over time scales up to thousands of seconds. We suggest that these two streams of research are reconciled if the MTL maintains a scale-invariant representation of stimulus history. We described a computational mechanism for constructing such a representation and compared its properties to hippocampal time cells and canonical behavioral results from studies of recency and contiguity effects. We suggest that this method for coding stimulus history may describe findings from other behavioral tasks, including judgments of recency, and that neural representations in other brain regions, notably the PFC

and striatum, may prove to maintain a similar scale-invariant representation of stimulus history.

References

Adler, A., Finkes, I., Katabi, S., Prut, Y., & Bergman, H. (2013). Encoding by synchronization in the primate striatum. *The Journal of Neuroscience, 33,* 4854–4866. doi:10.1523/JNEUROSCI.4791-12.2013

Adler, A., Katabi, S., Finkes, I., Israel, Z., Prut, Y., & Bergman, H. (2012). Temporal convergence of dynamic cell assemblies in the striato-pallidal network. *Journal of Neuroscience, 32,* 2473–2484. doi:10.1523/JNEUROSCI.4830-11.2012

Agster, K. L., Fortin, N. J., & Eichenbaum, H. (2002). The hippocampus and disambiguation of overlapping sequences. *The Journal of Neuroscience, 22,* 5760–5768.

Ainge, J. A., van der Meer, M. A. A., Langston, R. F., & Wood, E. R. (2007). Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. *Hippocampus, 17,* 988–1002. doi:10.1002/hipo.20301

Anderson, J., & Schooler, L. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396–408. doi:10.1111/j.1467-9280.1991.tb00174.x

Aristotle. (1930). *On memory and reminiscence* (J. I. Beare, Trans.), Retrieved from http://psychclassics.asu.edu/Aristotle/memory.htm

Atkinson, R. C., & Juola, J. F. (1974). Search and decision processes in recognition memory. In D. H. Krantz, R. C. Atkinson, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (pp. 243–290). San Francisco, CA: Freeman.

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, pp. 89–105). New York, NY: Academic Press. doi:10.1016/S0079-7421(08)60422-3

Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences, 32,* 73–78. doi:10.1016/j.tins.2008.10.004

Bangasser, D. A., Waxler, D. E., Santollo, J., & Shors, T. J. (2006). Trace conditioning and the hippocampus: The importance of contiguity. *The Journal of Neuroscience, 26,* 8702–8706. doi:10.1523/JNEUROSCI.1742-06.2006

Barker, G. R. I., Bird, F., Alexander, V., & Warburton, E. C. (2007). Recognition memory for objects, place, and temporal order: A disconnection analysis of the role of the medial prefrontal cortex and perirhinal cortex. *The Journal of Neuroscience, 27,* 2948–2957. doi:10.1523/JNEUROSCI.5289-06.2007

Bernacchia, A., Seo, H., Lee, D., & Wang, X. J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience, 14,* 366–372. doi:10.1038/nn.2752

Beylin, A., Gandhi, C., Wood, G., Talk, A., Matzel, L., & Shors, T. (2001). The role of the hippocampus in trace conditioning: Temporal discontinuity or task difficulty? *Neurobiology of Learning and Memory, 76,* 447–461. doi:10.1006/nlme.2001.4039

Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology, 6,* 173–189. doi:10.1016/0010-0285(74)90009-7

Bornstein, A. M., & Daw, N. D. (2012). Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience, 35,* 1011–1023. doi:10.1111/j.1460-9568.2011.07920.x

Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review, 114,* 539–576.

Brown, G. D., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review, 107,* 127–181. doi:10.1037/0033-295X.107.1.127

Brown, T. I., Ross, R. S., Keller, J. B., Hasselmo, M. E., & Stern, C. E. (2010). Which way was I going? Contextual retrieval supports the disambiguation of well learned overlapping navigational routes. *The Journal of Neuroscience, 30,* 7414–7422. doi:10.1523/JNEUROSCI.6021-09.2010

Chan, M., Ross, B., Earle, G., & Caplan, J. B. (2009). Precise instructions determine participants' memory search strategy in judgments of relative order in short lists. *Psychonomic Bulletin & Review, 16,* 945–951. doi:10.3758/PBR.16.5.945

Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: The role of awareness. *Science, 280,* 77–81. doi:10.1126/science.280.5360.77

Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system.* Cambridge, MA: MIT Press.

Davachi, L., Mitchell, J. P., & Wagner, A. D. (2003). Multiple routes to memory: Distinct medial temporal lobe processes build item and source memories. *Proceedings of the National Academy of Science of the United States of America, 100,* 2157–2162. doi:10.1073/pnas.0337195100

Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: Empirical and computational investigations of recency effects. *Psychological Review, 112,* 3–42. doi:10.1037/0033-295X.112.1.3

Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review, 108,* 452–478. doi:10.1037/0033-295X.108.2.452

DeVito, L. M., & Eichenbaum, H. (2011). Memory for the order of events in specific sequences: Contributions of the hippocampus and medial prefrontal cortex. *The Journal of Neuroscience, 31,* 3169–3175. doi:10.1523/JNEUROSCI.4202-10.2011

Drew, M. R., Couvillon, P. A., Zupan, B., Cooke, A., & Balsam, P. (2005). Temporal control of conditioned responding in goldfish. *Journal of Experimental Psychology: Animal Behavior Processes, 31,* 31–39. doi:10.1037/0097-7403.31.1.31

Dunn, J. C. (2008). The dimensionality of the remember-know task: A state-trace analysis. *Psychological Review, 115,* 426–446. doi:10.1037/0033-295X.115.2.426

Eichenbaum, H. (1999). The hippocampus and mechanisms of declarative memory. *Behavioural Brain Research, 103,* 123–133. doi:10.1016/S0166-4328(99)00044-3

Eichenbaum, H., Yonelinas, A., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience, 30,* 123–152. doi:10.1146/annurev.neuro.30.051606.094328

Ekstrom, A. D., & Bookheimer, S. Y. (2007). Spatial and temporal episodic memory retrieval recruit dissociable functional networks in the human brain. *Learning & Memory, 14,* 645–654. doi:10.1101/lm.575107

Ferbinteanu, J., & Shapiro, M. L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron, 40,* 1227–1239. doi:10.1016/S0896-6273(03)00752-9

Foerde, K., Race, E., Verfaellie, M., & Shohamy, D. (2013). A role for the medial temporal lobe in feedback-driven learning: Evidence from amnesia. *Journal of Neuroscience, 33,* 5698–5704. doi:10.1523/JNEUROSCI.5217-12.2013

Foerde, K., & Shohamy, D. (2011). Feedback timing modulates brain systems for learning in humans. *The Journal of Neuroscience, 31,* 13157–13167. doi:10.1523/JNEUROSCI.2701-11.2011

Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience, 5,* 458–462.

Fortin, N. J., Wright, S. P., & Eichenbaum, H. (2004). Recollection-like memory retrieval in rats is dependent on the hippocampus. *Nature, 431,* 188–191. doi:10.1038/nature02853

Frank, L. M., Brown, E. N., & Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron, 27,* 169–178. doi:10.1016/S0896-6273(00)00018-0

Fuster, J. M., & Jervey, J. P. (1982). Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. *Journal of Neuroscience, 2,* 361–375.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review, 107,* 289–344. doi:10.1037/0033-295X.107.2.289

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279–325. doi:10.1037/0033-295X.84.3.279

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

Ginther, M. R., Walsh, D. F., & Ramus, S. J. (2011). Hippocampal neurons encode different episodes in an overlapping sequence of odors task. *The Journal of Neuroscience, 31,* 2706–2711. doi:10.1523/JNEUROSCI.3413-10.2011

Glenberg, A. M., Bradley, M. M., Stevenson, J. A., Kraus, T. A., Tkachuk, M. J., & Gretz, A. L. (1980). A two-process account of long-term serial position effects. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 355–369. doi:10.1037/0278-7393.6.4.355

Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron, 61,* 621–634. doi:10.1016/j.neuron.2008.12.012

Grossberg, S., & Merrill, J. (1992). A neural network model of adaptively timed reinforcement learning and hippocampal dynamics. *Cognitive Brain Research, 1,* 3–38. doi:10.1016/0926-6410(92)90003-A

Grossberg, S., & Merrill, J. (1996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience, 8,* 257–277. doi:10.1162/jocn.1996.8.3.257

Gureckis, T. M., & Love, B. C. (2010). Direct associations or internal transformations? Exploring the mechanisms underlying sequential learning behavior. *Cognitive Science, 34,* 10–50. doi:10.1111/j.1551-6709.2009.01076.x

Hacker, M. J. (1980). Speed and accuracy of recency judgments for events in short-term memory. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 651–675.

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

Hasselmo, M. E., & Wyble, B. P. (1997). Free recall and recognition in a network model of the hippocampus: Simulating effects of scopolamine on human memory function. *Behavioural Brain Research, 89,* 1–34. doi:10.1016/S0166-4328(97)00048-X

Hayon, G., Abeles, M., & Lehmann, D. (2005). A model for representing the dynamics of a system of synfire chains. *Journal of Computational Neuroscience, 18,* 41–53. doi:10.1007/s10827-005-5479-1

Hinrichs, J. V. (1970). A two-process memory-strength theory for judgment of recency. *Psychological Review, 77,* 223–233. doi:10.1037/h0029101

Hintzman, D. L., & Curran, T. (1994). Retrieval dynamics of recognition and frequency judgments: Evidence for separate processes of familiarity and recall. *Journal of Memory and Language, 33,* 1–18. doi:10.1006/jmla.1994.1001

Hockley, W. E. (1984). Analysis of response time distributions in the study of cognitive processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 598–615. doi:10.1037/0278-7393.10.4.598

Howard, M. W. (2004). Scaling behavior in the temporal context model. *Journal of Mathematical Psychology, 48,* 230–238.

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 923–941. doi:10.1037/0278-7393.25.4.923

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology, 46,* 269–299. doi:10.1006/jmps.2001.1388

Howard, M. W., MacDonald, C. J., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2013). *A unified mathematical framework for coding time, space, and sequences in the medial temporal lobe.* Manuscript submitted for publication.

Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (2013). *A quantitative model of time in episodic memory.* Manuscript submitted for publication.

Howard, M. W., Viskontas, I. V., Shankar, K. H., & Fried, I. (2012). A neural signature of mental time travel in the human MTL. *Hippocampus, 22,* 1833–1847. doi:10.1002/hipo.22018

Howard, M. W., Youker, T. E., & Venkatadass, V. (2008). The persistence of memory: Contiguity effects across several minutes. *Psychonomic Bulletin & Review, 15,* 58–63. doi:10.3758/PBR.15.1.58

Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D., & Seamans, J. K. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences of the United States of America, 109,* 5086–5091.

Itskov, V., Curto, C., Pastalkova, E., & Buzséki, G. (2011). Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus. *The Journal of Neuroscience, 31,* 2828–2834. doi:10.1523/JNEUROSCI.3773-10.2011

Jenkins, L. J., & Ranganath, C. (2010). Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *The Journal of Neuroscience, 30,* 15558–15565. doi:10.1523/JNEUROSCI.1337-10.2010

Jensen, O., & Lisman, J. E. (1996). Hippocampal CA3 region predicts memory sequences: Accounting for the phase precession of place cells. *Learning & Memory, 3,* 279–287. doi:10.1101/lm.3.2-3.279

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition, 24,* 103–109.

Kahana, M. J., Howard, M., & Polyn, S. (2008). Associative processes in episodic memory. In H. L. Roediger III (Vol. Ed.) & J. Byrne (Series Ed.), *Cognitive psychology of memory: Vol. 2. Learning and memory—A comprehensive reference* (pp. 476–490). Oxford, England: Elsevier.

Kerr, K. M., Agster, K. L., Furtak, S. C., & Burwell, R. D. (2007). Functional neuroanatomy of the parahippocampal region: The lateral and medial entorhinal areas. *Hippocampus, 17,* 697–708. doi:10.1002/hipo.20315

Kesner, R. P., Gilbert, P. E., & Barua, L. A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behavioral Neuroscience, 116,* 286–290. doi:10.1037/0735-7044.116.2.286

Kesner, R. P., Hunsaker, M. R., & Gilbert, P. E. (2005). The role of CA1 in the acquisition of an object-trace-odor paired associate task. *Behavioral Neuroscience, 119,* 781–786. doi:10.1037/0735-7044.119.3.781

Kılıç, A., Criss, A. H., & Howard, M. W. (2013). A causal contiguity effect that persists across time scales. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 39,* 297–303.

Koene, R. A., & Hasselmo, M. E. (2007). First-in-first-out item replacement in a model of short-term memory based on persistent spiking. *Cerebral Cortex, 17,* 1766–1781. doi:10.1093/cercor/bhl088

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

Kumaran, D., & Maguire, E. A. (2006). An unexpected sequence of events: Mismatch detection in the human hippocampus. *PLoS Biology, 4,* e424. doi:10.1371/journal.pbio.0040424

Lehman, M., & Malmberg, K. J. (2013). A buffer model of memory encoding and temporal correlations in retrieval. *Psychological Review, 120,* 155–189.

Lehn, H., Steffenach, H. A., van Strien, N. M., Veltman, D. J., Witter, M. P., & Håberg, A. K. (2009). A specific role of the human hippocampus in recall of temporal sequences. *The Journal of Neuroscience, 29,* 3475–3484. doi:10.1523/JNEUROSCI.5370-08.2009

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus,*

6, 579–590. doi:10.1002/(SICI)1098-1063(1996)6:6<579::AID-HIPO3>3.0.CO;2-C

Lisman, J. E., & Idiart, M. A. (1995). Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science, 267,* 1512–1515. doi:10.1126/science.7878473

Lisman, J. E., & Jensen, O. (2013). The theta-gamma neural code. *Neuron, 77,* 1002–1016. doi:10.1016/j.neuron.2013.03.007

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron, 71,* 737–749. doi:10.1016/j.neuron.2011.07.012

MacDonald, C. J., Place, R., & Eichenbaum, H. (2012). *Hippocampal time cells map out a sequence of events independently of movement and space* (2012 Neuroscience Meeting Planner, Program No. 203.11). Retrieved from the Society for Neuroscience website: http://www.sfn.org

Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review, 87,* 252–271. doi:10.1037/0033-295X.87.3.252

Mangels, J. A., Ivry, R. B., & Shimizu, N. (1998). Dissociable contributions of the prefrontal and neocerebellar cortex to time perception. *Cognitive Brain Research, 7,* 15–39. doi:10.1016/S0926-6410(98)00005-6

Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences of the United States of America, 109,* 19462–19467.

Manns, J. R., Howard, M. W., & Eichenbaum, H. B. (2007). Gradual changes in hippocampal activity support remembering the order of events. *Neuron, 56,* 530–540. doi:10.1016/j.neuron.2007.08.017

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London: Series B. Biological Sciences, 262,* 23–81. doi:10.1098/rstb.1971.0078

Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: Coincidence detection of oscillatory processes. *Cognitive Brain Research, 21,* 139–170. doi:10.1016/j.cogbrainres.2004.06.012

Mayes, A., Isaac, C., Holdstock, J., Hunkin, N., Montaldi, D., Downes, J., . . . Roberts, J. N. (2001). Memory for single items, word pairs, and temporal order of different kinds in a patient with selective hippocampal lesions. *Cognitive Neuropsychology, 18,* 97–123.

McAndrews, M., & Milner, B. (1991). The frontal cortex and memory for temporal order. *Neuropsychologia, 29,* 849–859. doi:10.1016/0028-3932(91)90051-9

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102,* 419–457. doi:10.1037/0033-295X.102.3.419

McElree, B., Dolan, P. O., & Jacoby, L. L. (1999). Isolating the contributions of familiarity and source information to item recognition: A time course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 563–582. doi:10.1037/0278-7393.25.3.563

McElree, B., & Dosher, B. A. (1993). Serial recovery processes in the recovery of order information. *Journal of Experimental Psychology: General, 122,* 291–315. doi:10.1037/0096-3445.122.3.291

McNaughton, B. L., & Morris, R. G. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences, 10,* 408–415. doi:10.1016/0166-2236(87)90011-7

Meiser, T., Sattler, C., & Weisser, K. (2008). Binding of multidimensional context information as a distinctive characteristic of remember judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 32–49. doi:10.1037/0278-7393.34.1.32

Milner, B., Corsi, P., & Leonard, G. (1991). Frontal-lobe contribution to recency judgments. *Neuropsychologia, 29,* 601–618. doi:10.1016/0028-3932(91)90013-X

Muller, R. (1996). A quarter of a century of place cells. *Neuron, 17,* 813–822. doi:10.1016/S0896-6273(00)80214-7

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology, 64,* 482–488. doi:10.1037/h0045106

Muter, P. (1979). Response latencies in discriminations of recency. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 160–169. doi:10.1037/0278-7393.5.2.160

Naya, Y., & Suzuki, W. (2011). Integrating what and when across the primate medial temporal lobe. *Science, 333,* 773–776. doi:10.1126/science.1206773

Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review, 110,* 611–646. doi:10.1037/0033-295X.110.4.611

O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature, 381,* 425–428. doi:10.1038/381425a0

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research, 34,* 171–175.

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

Onyper, S. V., Zhang, Y. X., & Howard, M. W. (2010). Some-or-none recollection: Evidence from item and source memory. *Journal of Experimental Psychology: General, 139,* 341–364. doi:10.1037/a0018926

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science, 321,* 1322–1327. doi:10.1126/science.1159775

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review, 116,* 129–156. doi:10.1037/a0014420

Post, E. (1930). Generalized differentiation. *Transactions of the American Mathematical Society, 32,* 723–781. doi:10.1090/S0002-9947-1930-1501560-X

Province, J. M., & Rouder, J. N. (2012). Evidence for discrete-state processing in recognition memory. *Proceedings of the National Academy of Sciences of the United States of America, 109,* 14357–14362. doi:10.1073/pnas.1103880109

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, pp. 207–262). New York, NY: Academic Press. doi:10.1016/S0079-7421(08)60162-0

Roberts, W. A., Cheng, K., & Cohen, J. S. (1989). Timing light and tone signals in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 15,* 23–35. doi:10.1037/0097-7403.15.1.23

Ross, R. S., Brown, T. I., & Stern, C. E. (2009). The retrieval of learned sequences engages the hippocampus: Evidence from fMRI. *Hippocampus, 19,* 790–799. doi:10.1002/hipo.20558

Rugg, M. D., & Yonelinas, A. P. (2003). Human recognition memory: A cognitive neuroscience perspective. *Trends in Cognitive Sciences, 7,* 313–319. doi:10.1016/S1364-6613(03)00131-1

Schwartz, G., Howard, M. W., Jing, B., & Kahana, M. J. (2005). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychological Science, 16,* 898–904. doi:10.1111/j.1467-9280.2005.01634.x

Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review, 115,* 893–912. doi:10.1037/a0013396

Shankar, K. H., & Howard, M. W. (2012). A scale-invariant representation of time. *Neural Computation, 24,* 134–193. doi:10.1162/NECO_a_00212

Shankar, K. H., & Howard, M. W. (2013). *Optimally fuzzy scale-free memory.* Manuscript submitted for publication.

Shapiro, M. L., Kennedy, P. J., & Ferbinteanu, J. (2006). Representing episodes in the mammalian brain. *Current Opinion in Neurobiology, 16,* 701–709. doi:10.1016/j.conb.2006.08.017

Shimamura, A., Janowsky, J., & Squire, L. (1990). Memory for the temporal order of events in patients with frontal lobe lesions and amnesic patients. *Neuropsychologia, 28,* 803–813. doi:10.1016/0028-3932(90)90004-8

Sirotin, Y. B., Kimball, D. R., & Kahana, M. J. (2005). Going beyond a single list: Modeling the effects of prior experience on episodic free recall. *Psychonomic Bulletin & Review, 12,* 787–805. doi:10.3758/BF03196773

Smith, M. A., & Sommer, M. A. (2013). Spatial and temporal scales of neuronal correlation in visual area v4. *The Journal of Neuroscience, 33,* 5422–5432. doi:10.1523/JNEUROSCI.4782-12.2013

Solomon, P. R., Vander Schaaf, E. R., Thompson, R. F., & Weisz, D. J. (1986). Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. *Behavioral Neuroscience, 100,* 729–744. doi:10.1037/0735-7044.100.5.729

Spiers, H. J., Maguire, E. A., & Burgess, N. (2001). Hippocampal amnesia. *Neurocase, 7,* 357–382. doi:10.1076/neur.7.5.357.16245

Staresina, B. P., & Davachi, L. (2006). Differential encoding mechanisms for subsequent associative recognition and free recall. *The Journal of Neuroscience, 26,* 9162–9172. doi:10.1523/JNEUROSCI.2877-06.2006

Staresina, B. P., Fell, J., Dunn, J. C., Axmacher, N., & Henson, R. N. (2013). Using state-trace analysis to dissociate the functions of the human hippocampus and perirhinal cortex in recognition memory. *Proceedings of the National Academy of Sciences of the United States of America, 110,* 3119–3124. doi:10.1073/pnas.1215710110

Starns, J. J., & Hicks, J. L. (2005). Source dimensions are retrieved independently in multidimensional monitoring tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 1213–1220. doi:10.1037/0278-7393.31.6.1213

Tam, S. K. E., & Bonardi, C. (2012a). Dorsal hippocampal involvement in appetitive trace conditioning and interval timing. *Behavioral Neuroscience, 126,* 258–269. doi:10.1037/a0027164

Tam, S. K. E., & Bonardi, C. (2012b). Dorsal hippocampal lesions disrupt Pavlovian delay conditioning and conditioned-response timing. *Behavioural Brain Research, 230,* 259–267. doi:10.1016/j.bbr.2012.02.016

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., & Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience, 7,* 887–893. doi:10.1016/j.bbr.2012.02.016

Templer, V. L., & Hampton, R. R. (2013). Cognitive mechanisms of memory for order in rhesus monkeys (*Macaca mulatta*). *Hippocampus, 23,* 193–201.

Tubridy, S., & Davachi, L. (2011). Medial temporal lobe contributions to episodic sequence encoding. *Cerebral Cortex, 21,* 272–280. doi:10.1093/cercor/bhq092

Tulving, E. (1983). *Elements of episodic memory*. New York, NY: Oxford University Press.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology, 26,* 1–12. doi:10.1037/h0080017

Uncapher, M. R., Otten, L. J., & Rugg, M. D. (2006). Episodic encoding is more than the sum of its parts: An fMRI investigation of multifeatural contextual encoding. *Neuron, 52,* 547–556. doi:10.1016/j.neuron.2006.08.011

Unsworth, N. (2008). Exploring the retrieval dynamics of delayed and final free recall: Further evidence for temporal-contextual search. *Journal of Memory and Language, 59,* 223–236. doi:10.1016/j.jml.2008.04.002

van Strien, N. M., Cappaert, N. L. M., & Witter, M. P. (2009). The anatomy of memory: An interactive overview of the parahippocampal-hippocampal network. *Nature Reviews Neuroscience, 10,* 272–282. doi:10.1038/nrn2614

Voss, R. F., & Clarke, J. (1975). 1/f noise in music and speech. *Nature, 258,* 317–318. doi:10.1038/258317a0

Waugh, N. C., & Norman, D. (1965). Primary memory. *Psychological Review, 72,* 89–104. doi:10.1037/h0021797

Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science, 261,* 1055–1058.

Wimmer, G. E., & Shohamy, D. (2012). Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science, 338,* 270–273. doi:10.1126/science.1223252

Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review, 114,* 152–176. doi:10.1037/0033-295X.114.1.152

Wixted, J. T., & Squire, L. R. (2011). The familiarity/recollection distinction does not illuminate medial temporal lobe function: Response to Montaldi and Mayes. *Trends in Cognitive Sciences, 15,* 340–341. doi:10.1016/j.tics.2011.06.006

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron, 27,* 623–633. doi:10.1016/S0896-6273(00)00071-4

Wu, X., & Levy, W. B. (2001). Simulating symbolic distance effects in the transitive inference problem. *Neurocomputing, 38–40,* 1603–1610. doi:10.1016/S0925-2312(01)00512-4

Yakovlev, V., Fusi, S., Berman, E., & Zohary, E. (1998). Inter-trial neuronal activity in inferior temporal cortex: A putative vehicle to generate long-term visual association. *Nature Neuroscience, 1,* 310–317. doi:10.1038/1131

Yi, L. (2007). Applications of timing theories to a peak procedure. *Behavioural Processes, 75,* 188–198. doi:10.1016/j.beproc.2007.01.010

Yntema, D. B., & Trask, F. P. (1963). Recall as a search process. *Journal of Verbal Learning and Verbal Behavior, 2,* 65–74. doi:10.1016/S0022-5371(63)80069-9

Yonelinas, A. P., & Parks, C. M. (2007). Receiver operating characteristics (ROCs) in recognition memory: A review. *Psychological Bulletin, 133,* 800–832. doi:10.1037/0033-2909.133.5.800