Jeremy R. Manning, Michael J. Kahana, and Kenneth A. Norman

August 31, 2013

The role of context in episodic memory

Corresponding author:
Kenneth A. Norman
Department of Psychology, Princeton University
Green Hall, Washington Road
Princeton, NJ 08544
Telephone: 609 258 9694
Fax: 609 258 1113

knorman@princeton.edu

Word count: 6475
Figure count: 3
Index keywords: Computational models, context, ECoG, electrophysiology, episodic memory, fMRI, free recall, situation model

Abstract

In this chapter we discuss the role of context in organizing episodic memories. We define context as *slowly drifting* information (i.e. information that persists over a relatively long time scale in the person's brain; for example, a representation of the person's location). First, we use the Temporal Context Model of memory search (Howard & Kahana, 2002) to illustrate how binding slowly-drifting contextual information to more transient representations (e.g. of studied words) serves to organize our memories of the more transient information. We next present electrophysiological studies examining the role of slowly drifting representations in organizing episodic memories, and we provide an overview of the brain systems involved in representing slowly drifting information. Finally, we discuss sources of variability in the rate of contextual drift.

1 Introduction

Context is simultaneously one of the most fundamental and elusive concepts in memory research. Memory researchers often define context by exclusion: in a memory experiment, there is a set of *items* that the participant is being asked to memorize (e.g. a list of words), and then there is *context*, which reflects everything else that is represented in the person's brain during the experiment. Context might include, for example, information about the external environment, mood, thoughts about recently encountered items, plans concerning the future, and incidental features of the stimuli such as the color and the spatial location of a word on the screen (for a review see Smith & Vela, 2001). This list makes it clear that defining context in terms of the type of information being represented is futile. Under this definition, anything can be context – for example, if you test memory for which colors were seen, rather than word identity, then the background colors become the items and word identities become part of the context.

Instead of defining context in terms of specific types of information (and the roles they play in particular experiments), we focus here on the time scale of information representation. Figure 1 illustrates, in timeline form, what might being going through a person's head as they learn a list of words in a typical memory experiment. The figure shows how information is represented at different time scales. At the bottom, there are

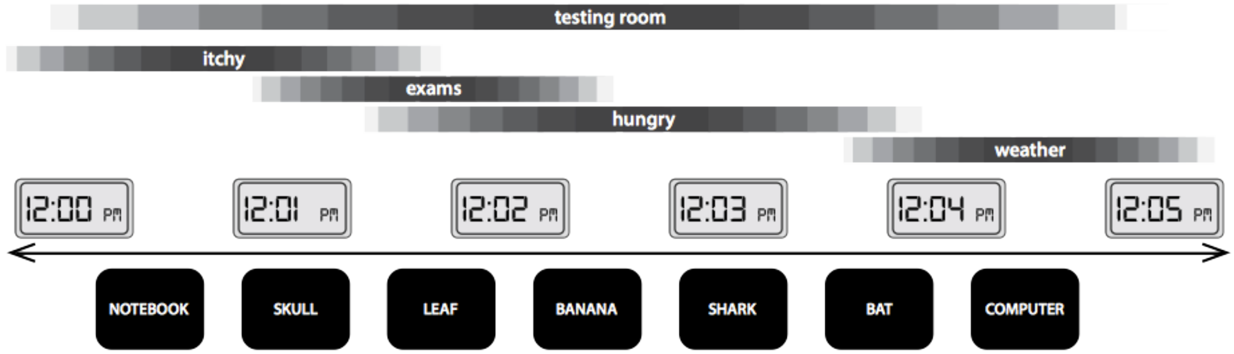

Figure 1: Illustration of contextual drift. A hypothetical participant studies the 7 words in the bottom row over the course of 5 minutes (times are indicated by the digital clock faces). In addition to thoughts related to the studied words, the participant experiences thoughts related to other external stimuli (e.g. the testing room) as well as internal states (feeling itchy, worrying about upcoming exams, etc.). These thoughts persist for different amounts of time and can overlap.

sensory representations of words that persist only while the word is being presented; above, there are representations of thoughts or mental states that persist over longer time scales.

The central idea of this chapter is that slowly drifting information (i.e., information that persists over relatively long time scales) can be used to *time stamp* and organize more quickly drifting information. This time stamping is accomplished by means of the hippocampus binding co-active representations together (O'Reilly & Rudy, 2001), including representations that are drifting at different rates. In the scenario shown in Figure 1, these bindings allow words to cue retrieval of co-active *contextual threads* and vice-versa. For example, if the participant recalls the word **shark**, they might also recall being hungry when that word was presented, which in turn might trigger retrieval of the fact that **leaf** was also presented while the person was hungry. When **leaf** is retrieved, this might trigger retrieval of having been worried about exams, and then the participant might recall also hearing the word **skull** when worrying about their exams, and so on. The set of contextual threads that were active during a previously experienced event (e.g. studying the word **skull**) constitutes a unique time stamp for that event, and we will refer to the process of reactivating these contextual threads as *contextual reinstatement*. Effectively, this contextual reinstatement process allows the participant to mentally "jump back in time" to when the contextual threads were initially active. If the participant succeeds in reinstating a large number of contextual threads that were linked to an item at study, they will all convergently cue the associated event, thereby boosting the probability of retrieving the event.

To understand the role of context in episodic memory, we need to understand the factors that give rise to slow drift in the brain, and the role that slow drift plays in time stamping memories. Speaking generally, there are two different (and non-mutually exclusive) ways to get slow drift. The first possibility is that slow drift can arise from the brain representing slowly drifting features of the world. For example, because we cannot teleport between locations, our location in one moment will be similar to our location in the next moment. Therefore brain areas that represent current location information will show the requisite slow drift property. The second possibility is *intrinsic maintenance*. Even if features of the world disappear quickly, the brain has the ability to sustain patterns of neural firing corresponding to both external features of the world and internal thoughts (see Section 3 below and also Ranganath, Hasselmo, and Stern, this volume).

The rest of this chapter is divided into three main parts. In Section 2, we describe a computational model of memory search to illustrate how the aforementioned psychological principles (slow drift and contextual reinstatement) can account for detailed patterns of memory data. We also discuss electrophysiological evidence for slow drift and contextual reinstatement. In Section 3, we provide an overview of the brain systems involved in representing slowly drifting information. In Section 4, we discuss sources of variability in contextual drift.

2 Explaining behavioral memory data using the Temporal Context Model

The worth of a psychological theory may be measured by how well it explains detailed patterns of behavioral and neural data. Towards this end, researchers have built computational models that instantiate the principles outlined above (slow drift and contextual reinstatement) and fit the models to detailed patterns of memory data. Here we focus on modeling data from the *free recall* paradigm, followed by a brief survey of other relevant paradigms. In the free recall paradigm, participants study lists of items (typically words) and then attempt to recall the items in any order. Free recall is a useful testbed for models of context and memory because it provides data regarding both accuracy (i.e., which memories were recalled) and the order in which items were recalled.

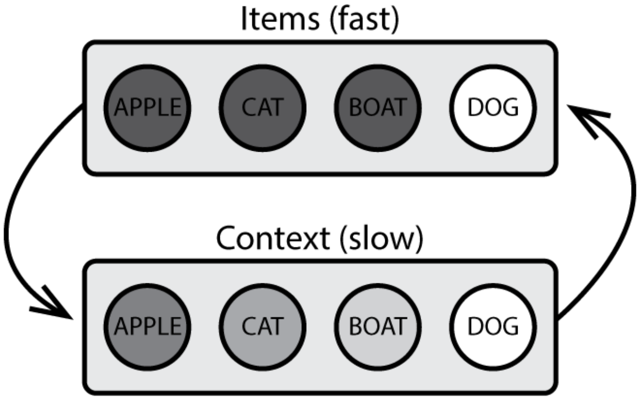

Figure 2. The Temporal Context Model. Nodes (represented by circles) in the item layer correspond to the sensory representations of the words on the study list. When an item (in this case, **dog**) is studied, its node in the item layer is activated. The shadings of the nodes reflect their activations, where brighter shading reflects greater activation. Nodes in the context layer represent a running average of item-layer activations. Here the activations in the context layer reflect the state of the model as the participant studies the fourth word on a list, **dog**, after studying **apple**, **cat**, and **boat**. The arrows denote interactions between the item and context layers.

To provide an intuition for how slow drift and contextual reinstatement can account for behavioral memory data, we will focus our discussion on one specific model, the Temporal Context Model (TCM; Howard & Kahana, 2002), schematized in Figure 2. TCM is a member of a large family of models that incorporate slow drift and contextual reinstatement (e.g. Davelaar et al., 2005; Dennis & Humphreys, 2001; Polyn et al., 2009; Sederberg et al., 2008; Shankar & Howard, 2012). The purpose of this section is not to differentiate between models within this family, but rather to illustrate these models' shared predictions.

In TCM, there are two interconnected layers of nodes (or computing elements). These two layers represent the same information but at different time scales. The *item layer* represents the item that is currently being studied on the list, and the *context layer* represents a running average of recently studied items. As the participant studies each item on a list (e.g. the words **apple**, **cat**, **boat**, **dog**), the corresponding node in the item layer is activated. In the figure, the participant is currently viewing the word **dog**, so the **dog** node is active (denoted by bright shading) and the **apple**, **cat**, and **boat** nodes are inactive (denoted by darker shading). As each node of the item layer is activated, the item becomes associated with the current state of the context layer. The context layer is continually updated by averaging together the current item representation with the previous state of context. Because the context layer computes a running average, the state of context evolves gradually as the participant studies the list.

Recall is simulated in TCM by using the current state of context as a retrieval cue. Each item's node in the item layer is activated according to how similar that item's associated state of context is to the current state of context. Items stochastically compete to be recalled, with more active items serving as stronger competitors. When an item wins the recall competition, two things happen. First, the recalled item is incorporated into the current state of context (e.g. if you recall **dog**, that is blended into the context vector). Second, the current state of context is also updated with *retrieved mental context* (i.e. the state of mental context that was linked to the just-retrieved item at study). The updated state of context is then used to probe memory for other items. As discussed in the Introduction, reinstating an item's study-phase context can be construed as mentally jumping back to the moment when the just-recalled item was studied. Cuing with reinstated contextual information increases the chances that the next item recalled will be one associated with a similar context.

2.1 Modeling key regularities in free recall behavior

TCM and similar models have enjoyed extensive success at explaining many stereotyped behaviors observed during free recall and other episodic memory-related tasks. One fundamental regularity is termed the *recency effect*, which refers to participants' ability to more easily retrieve information pertaining to recent experiences (e.g. items from the end of a just-studied list) than information pertaining to long-ago experiences (Murdock, 1962). The *contiguity effect* is another fundamental regularity in free recall, and refers to participants' tendency to successively recall items that occupied nearby positions in the study lists (Kahana, 1996). For example, if a list contained the sub-sequence **apple cat boat** and the participant recalled the word **cat**, it is far more likely that the next response would be either **boat** or **apple** than some other list item.

Both the recency effect and the contiguity effect exhibit *time scale invariance*, meaning that these same patterns are observed at short and long time scales. For example, recency effects are observed in immediate free recall experiments (where participants are given a list of words and are tested immediately afterwards) and also in situations where the items on the list are spaced out across multiple days (Glenberg et al., 1983) or even longer

(Moreton & Ward, 2010). Likewise, contiguity effects are observed in immediate free recall and also in *continual distractor free recall*, where participants perform a distracting task after studying each item (Bjork & Whitten, 1974; Howard & Kahana, 1999). Contiguity effects have even been observed across word lists (i.e., given that a participant has just recalled an item from the fourth list, they are more likely to recall the next item from the third or fifth list than from the second or sixth lists; Howard et al., 2008).

The recency and contiguity effects (along with some degree of time scale invariance) emerge naturally out of retrieved context models of episodic memory like TCM. According to these models, recency effects arise because the context present at test is *relatively* more similar to the context associated with the end-of-list items than the context associated with earlier items. This fact is true regardless of whether the items were presented within a few seconds of each other or days apart, which explains why the model predicts recency effects at both short and long time scales.

Retrieved context models like TCM explain the contiguity effect in terms of slow drift at study and contextual reinstatement at test. When participants retrieve an item, they retrieve contextual features associated with that item, and then they use the retrieved contextual features to probe their memories for more items. Due to slow drift at study, items studied nearer in time will have been associated with (relatively) more similar states of context than far-apart items. Just as with the recency effect, this *relative match* property holds regardless of the time delay between successive item presentations, explaining why retrieved context models predict contiguity effects at a variety of time scales.

The fact that context models can explain recency and contiguity does not mean that contextual drift is always responsible for these effects. For example, it may be possible to explain recency effects in immediate free recall in terms of participants actively maintaining end-of-list items in working memory, and then directly recalling these items from working memory (e.g. Atkinson & Shiffrin, 1968). Likewise, it is possible to explain contiguity effects in immediate free recall in terms of participants directly forming links between adjacently studied items. However, these alternative accounts can not explain why recency and contiguity effects persist across long time scales.

Importantly, many retrieved context models (including TCM) are not fully time scale invariant. For example, TCM has a characteristic contextual drift rate which determines how rapidly the state of context evolves to incorporate new information. Events that occur at time scales slower than this drift rate will not show substantial recency or contiguity effects. A recently developed theory from Shankar and Howard (2012) extends TCM to account for precise timing and full time scale invariance, by positing that drift occurs across a spectrum of time scales (also see Howard & Eichenbaum, 2013).

2.2 Other findings that can be addressed using this framework

While modelers interested in how context shapes episodic memory have focused primarily on free recall, it is important to note that contextual tagging and contextual reinstatement contribute to some degree to virtually every episodic memory paradigm. For example, in tests of recognition memory, presenting information about a past event (e.g. displaying a previously studied word) can reinstate the context in which the word was studied, such that words studied in similar contexts are subsequently remembered more easily (Schwartz et al., 2005). Contextual reinstatement can also lead to memory errors. For example, suppose that a participant studies two lists of items, *A* and *B*. If the participant is reminded of studying list *A* prior to studying list *B*, the contextual threads associated with list *A* may be reinstated and bound to the list *B* items. This binding can lead participants to *misattribute* memories of list *B* items to list *A* (Gershman et al., 2013; Hupbach et al., 2007; Sederberg et al., 2011). Just as contextual reinstatement can facilitate access to memories linked to those contextual threads, contextual shifts can inhibit access to memories associated with out-of-date contextual threads. Evidence from the *list-method directed forgetting paradigm* suggests that, when participants are asked to forget list *A* items prior to studying list *B*, they respond to this *forget cue* instruction by deliberately shifting their state of context, thereby making the list *A* items less accessible (for a review of the relevant evidence see Sahakyan et al., 2013).

Contextual matching may also be used to explain how people judge the order in which a series of events occurred. Specifically, if our current context is more similar to the context associated with event *y* than with event *x*, we may judge event *y* to have occurred more recently than event *x* (see discussion of Manns et al., 2007 in Section 3 for neural data relating recency judgments to contextual drift).

2.3 Electrophysiological evidence for slow drift and contextual reinstatement

Recent advances in neural recording and analysis methods have allowed researchers to test whether neural patterns during episodic memory experiments are consistent with retrieved context models. As described above, retrieved context models predict that recalling a studied item should lead to the reinstatement of a gradually evolving contextual representation. To test this prediction, Manning et al. (2011) recorded electrical signals from electrodes implanted throughout the brains of human neurosurgical patients as they participated in a delayed free recall experiment (Fig. 3A,B). The researchers first sought to isolate slowly drifting neural patterns that might be involved in representing context. Next, the researchers examined whether those (putative) context representations were reinstated as the participants recalled the studied words.

The researchers isolated candidate context representations by identifying *temporally autocorrelated* neural patterns (i.e. patterns that were more similar during the study of nearby words than temporally distant words) as the participants studied the words. After the patients had studied and recalled words from many lists, the researchers computed the similarity between the neural patterns recorded just prior to recalling a word and the neural patterns recorded at study.

If the neural patterns recorded as a participant recalls an item reflect only item information, then they should match the neural patterns recorded when the item was studied and also possibly afterwards (to the extent that the item representation persisted in the participant's brain). In contrast, if the neural pattern at retrieval reflects reinstated, slowly drifting context from the study phase, then it should match the patterns recorded *both* before and after that item was studied, with similarity falling off gradually in both directions. Figure 3C shows that the data matched this latter pattern, thereby supporting retrieved context models. Furthermore, the degree to which individual patients exhibited this neural signature of contextual reinstatement was correlated with the behavioral contiguity effect (i.e., participants' tendency to successively recall neighboring list items; Fig. 3D). In a related study, Howard et al. (2012) collected extracellular recordings from various medial temporal lobe regions in humans during a recognition memory test, and found a similar neural signature of contextual reinstatement.

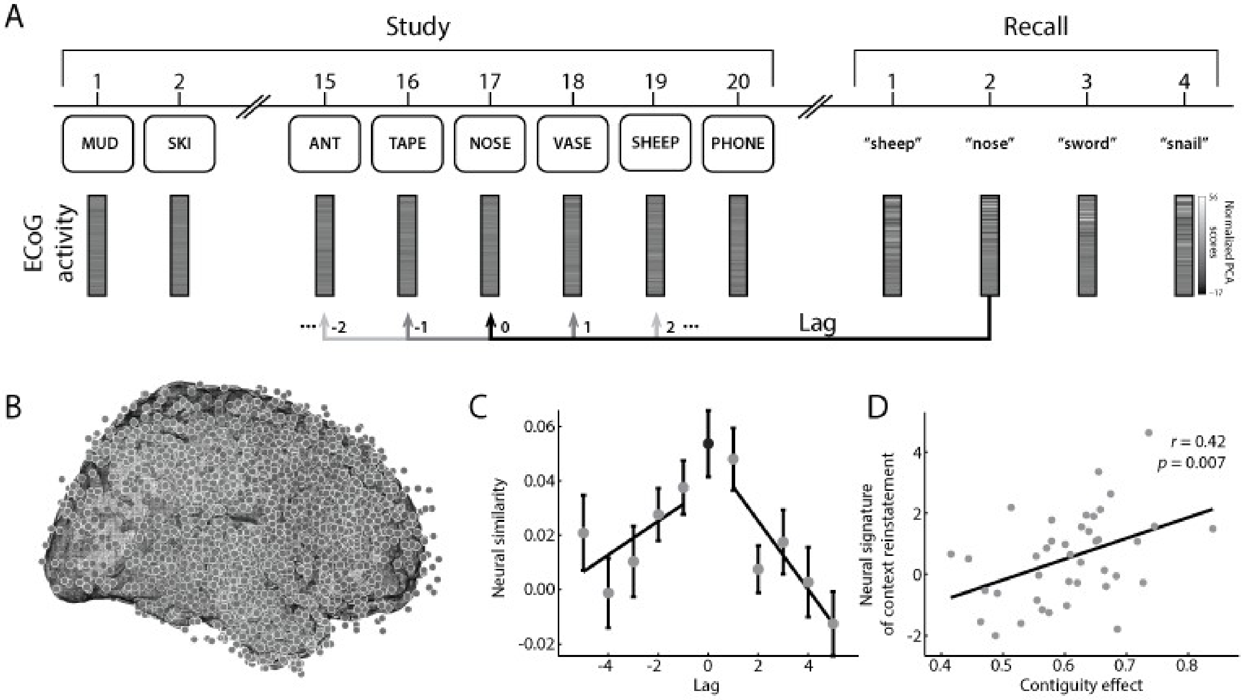

Figure 3. Neural evidence for contextual reinstatement in humans. A. After studying a list of 15 or 20 words and performing a brief distraction task, participants recalled as many words as they could remember, in any order. ECoG activity was recorded during each study and recall event. The similarity between the recorded patterns was computed as a function of lag. **B.** Each dot marks the location of a single electrode implanted in the brain of a neurosurgical patient. **C.** Similarity between the activity recorded during recall of a word from serial position i and study of a word from serial position *i* + lag (black dot denotes study and recall of the same word, i.e., lag = 0. **D.** Participants who exhibited stronger neural signatures of context reinstatement also exhibited more pronounced contiguity effects.

3 Brain systems involved in representing context

3.1 Prefrontal cortex

As reviewed by Polyn and Kahana (2008), prefrontal cortex (PFC) has several properties that make it an especially good candidate for representing contextual information. In particular, PFC can actively maintain patterns of neural firing in the face of distraction (e.g. Miller et al., 1996). This capacity for active maintenance should cause neural patterns in PFC to change more slowly than they would otherwise (see Ranganath, Hasselmo, and Stern, this volume). In keeping with this idea, a number of neural recording studies have found direct evidence that neural patterns in PFC drift slowly. For example, Hyman et al. (2012) recorded from dozens of neurons in the rodent medial prefrontal cortex (mPFC) as rats navigated in two environments. They found that the firing rates of mPFC neurons were temporally autocorrelated. Furthermore, Jenkins and Ranganath (2010) found that patterns of fMRI activity in the right lateral PFC drifted slowly while participants studied lists of pictures; the rate of neural drift predicted how accurately participants could remember when particular pictures were presented over the time course of the experiment. Converging evidence for the PFC's role in representing contextual information comes from studies of patients with frontal lobe damage. These studies have found that damage to frontal regions impairs memory performance on strongly contextually mediated tasks like free recall, whereas performance is relatively spared on tasks where context plays a lesser role (e.g. Shimamura, 1994; Wheeler et al., 1995).

3.2 MTL structures

Studies have found evidence for slowly drifting patterns of neural activity in several medial temporal lobe structures, including parahippocampal cortex (PHC) and hippocampus (e.g. Howard et al., 2012). We elaborate on the roles of each of these structures below (for additional discussion of MTL contributions to episodic memory, see Davachi and Preston, this volume).

3.2.1 Parahippocampal cortex and the posterior medial system

Several recent papers have argued that PHC represents a person's inference about the *situation* they are currently in (e.g. reading a book, listening to music, cooking dinner, etc). For example, Bar and Aminoff (2003) found that the PHC shows greater activation in response to objects that are strongly diagnostic of situational context (e.g. a roulette wheel or a beach chair) than to objects without a strong associated context (e.g. a cherry or a fly), although see Epstein and Ward (2010) for an alternative interpretation. Insofar as a person's representation of the situation they are in changes gradually over time (except at event boundaries; see Section 4), brain regions that represent situational information should exhibit gradually changing neural patterns. In keeping with this view, the slow drift property predicted by context models has been demonstrated in PHC using both electrophysiology (Howard et al., 2012) and fMRI (Turk-Browne et al., 2012).

Although PHC plays an important role in representing situational context, it is not the only such region. As reviewed by Ranganath and Ritchey (2012), PHC is part of a densely interconnected network of regions called the *posterior medial system* that includes retrosplenial cortex, the mamillary bodies, anterior thalamic nuclei, presubiculum, parasubiculum, posterior cingulate, precuneus, angular gyrus, and ventromedial PFC. Ranganath and Ritchey (2012) propose that structures in the posterior medial system work together to match incoming cues about the current context to internal *situation models* that specify the spatial, temporal, and causal relationships that define specific situations. For example, the situation model for going to a movie might describe the properties of movie theaters, the typical time sequence of events during a movie-going outing, movie theater etiquette, and so on (Zacks et al., 2007; Zwaan & Radvansky, 1998).

3.2.2 The hippocampus

The hippocampus is the key structure responsible for binding quickly drifting information (e.g. sensory representations of words on a study list) to more slowly drifting information, thereby making it possible for studied items to cue reinstatement of slowly drifting contextual information, and vice-versa (Cohen & Eichenbaum, 1993; Diana et al., 2007; O'Reilly & Rudy, 2001). To play this binding role, the hippocampus needs to receive inputs from areas representing slowly drifting information. This is accomplished via connections from areas like PFC and PHC that go through the entorhinal cortex (which itself exhibits gradually evolving neural patterns; Egorov et al., 2002) into the hippocampus.

A recent study by Manns et al. (2007) provides clear evidence that patterns of hippocampal firing evolve gradually over time, and that these gradually evolving patterns are behaviorally relevant. In their study, Manns et al. (2007) simultaneously recorded activity from multiple neurons in the CA1 hippocampal subregion as rats sequentially sampled a "list" of odors. After sampling the sequence of odors, the rat had to choose which of two odors in the sequence had been presented more recently. The authors found that patterns of neural firing in CA1 changed gradually as animals sampled the odors, and that the degree of neural drift over the course of the list predicted behavioral accuracy on a recency-discrimination test. The authors interpret this finding in terms of the idea that greater neural drift indicates greater contextual separation, which (in turn) makes it easier to temporally discriminate between items on the recency test. Notably, in addition to showing a within-list neural drift effect, Manns et al. (2007) also observed slow drift across lists of odors (for discussion, see Howard & Eichenbaum, 2013); this fits with the idea (mentioned in Section 2) that context drifts at multiple time scales (Shankar & Howard, 2012). For a related finding showing slow drift in CA1, see Mankin et al. (2012).

3.3 Temporal receptive windows

In addition to the aforementioned regions, how can we discover other areas involved in representing context? Naïvely, one could just look for regions exhibiting gradually evolving neural patterns using fMRI. The problem with this approach is that the fMRI signal is constrained to drift slowly due to the sluggishness of the blood flow response that it measures, regardless of the drift rates of the underlying neurocognitive processes. Therefore the mere presence of slow drift in the fMRI signal is not diagnostic of slow drift in the person's thoughts.

To address this problem, Hasson et al. (2008) devised a new technique for measuring a brain region's sensitivity to information at different time scales. Instead of directly measuring the drift rate of neural patterns, they measured the *history-dependence* of neural activity in a region. Specifically, they manipulated what came before a particular stimulus (by rearranging scenes in a movie), and asked whether the response of a region to a particular scene was altered by changing the scenes that came before it. For example, if a region's response to a scene is altered by changing what happened 5 minutes previously (but not 10 minutes ago), this indicates that the region retains information from 5 minutes ago (but not 10 minutes ago). Hasson et al. (2008) define the *temporal receptive window* (or TRW) for a particular region as the length of time, prior to the stimulus presentation, during which the presentation of other information may affect the neural response to the stimulus.

Using a variant of this approach, Lerner et al. (2011) found that the temporal parietal junction (including the angular gyrus and supramarginal gyrus) and precuneus were sensitive to the previous sentence in an auditory story, and medial PFC showed an even longer TRW (extending to the previous paragraph and possibly further). Hasson et al. (2008) identified a similar set of "long TRW" regions using a movie stimulus, including precuneus and the temporal parietal junction. Importantly, there is strong overlap between the set of long TRW regions and the posterior medial network regions identified by Ranganath and Ritchey (2012). We should emphasize that regions can show long TRWs for a variety of reasons – for example, a region might have a long TRW because it has intrinsic integrator properties (e.g. Arnsten et al., 2012), or because it can actively maintain specific patterns of activity (as in PFC), or because it is receiving information from other regions involved in memory storage (e.g. the hippocampus).

4 What drives contextual drift?

One of the main goals of theories of context and memory is understanding *variability* in contextual drift by explaining the circumstances that result in mental context changing more or less quickly. A key implication of the situation model view described in Section 3 is that mental context will change sharply when a person's (inferred) situation changes.

This view is supported by data showing that *event boundaries* (moments when participants infer a change in their situation; e.g. shifting from eating dinner to washing dishes) can cause forgetting. In the event processing literature, several behavioral studies have found that (controlling for elapsed time) participants are impaired at recalling details from the previous event compared to the current event (e.g. Radvansky & Copeland, 2006; Swallow et al., 2009). Using a long-term memory paradigm, Ezzyat and Davachi (2011) found that participants had difficulty recalling associations between adjacent sentences that spanned an event boundary, compared to sentences that were part of the same event. These behavioral findings are consistent with the idea that event boundaries induce a sharp discontinuity in context, resulting in decreased accessibility of details from the previous event and also decreased contiguity effects in long-term memory tests (for relevant neural evidence, see Swallow et al., 2011).

Explaining these event-boundary effects poses a major challenge for computational models of contextual drift. Computational models that update context via a simple integration process (i.e., by computing a running average of recently encountered stimuli) posit that, when the situation shifts, information about previously encountered stimuli will gradually fade out of context rather than exhibiting a rapid shift. As discussed by Polyn et al. (2009), this gradual fade is not enough to explain the sharp drop in recall observed at event boundaries. Polyn et al. (2009) created situational shifts at study by having participants switch (multiple times) between encoding tasks as they studied lists of words, and then had participants freely recall the studied items. To model the effects of these task switches on free recall, Polyn et al. (2009) had to incorporate an extra context disruption mechanism that was triggered whenever participants switched between encoding tasks.

While this context disruption mechanism helps to fit the data, it does not provide a clear mechanistic account of why context is disrupted at event boundaries (it just posits that it happens). Modeling work by Shankar et al. (2009) may provide some insight into this issue. Their model (which they call pTCM, for predictive TCM) modifies TCM such that, instead of updating context with item information, context is updated with a prediction of which items will be presented next. Insofar as event boundaries are marked by sharp changes in predictions (i.e., what you predict at the end of one event is very different from what you predict at the beginning of another), this model may be able to simulate the findings described above (also see Reynolds et al., 2007).

5 Concluding remarks

Context has long been the "dark matter" of memory theories. Researchers have found it necessary to posit a gradually evolving context representation in order to explain patterns of memory data from free recall and other tasks. This gradually evolving representation (and the idea that it can be reinstated during retrieval) is the glue that holds together most modern theories of memory retrieval. However, until recently, no one had been able to observe contextual drift or reinstatement directly. Instead, the role of context during memory encoding and retrieval had been something indirectly inferred through its effects on behavioral memory performance.

In this chapter, we have reviewed recent progress in the cognitive neuroscience of memory that has allowed us to start bringing the dark matter of context into the light. Neurophysiological and fMRI studies have given us a much better idea of which regions are most strongly involved in representing contextual information and, more importantly, they have given us the ability to track how neural activity drifts within those regions. In the coming years, the ability to track this drift and relate it to memory behavior will allow us to develop even more powerful models of how our brains time stamp our memories and how these time stamps allow us to retrieve information concerning the past.

References

Arnsten, A. F. T., Wang, M. J., and Paspalas, C. D. (2012). Neuromodulation of thought: flexibilities and vulnerabilities in prefrontal cortical network synapses. *Neuron,* 76:223 – 239.

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

Bar, M. and Aminoff, E. (2003). Cortical analysis of visual context. *Neuron,* 38:347 – 358. Bjork, R. A. and Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology,* 6:173 – 189.

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

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

Dennis, S. and Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review,* 108:452 – 478.

Diana, R. A., Yonelinas, A. P., and Ranganath, C. (2007). Imaging recollection and familiar- ity in the medial temporal lobe: a three-component model. *Trends in Cognitive Sciences,* doi:10.1016/j.tics.2007.08.001.

Egorov, A., Hamam, B., Fransen, E., Hasselmo, M., and Alonso, A. (2002). Graded persistent activity in entorhinal cortex neurons. *Nature,* 420(6912):173 – 8.

Epstein, R. and Ward, E. (2010). How reliable are visual context effects in the parahippocampal place area? *Cerebral Cortex,* 20(2):294.

Ezzyat, Y. and Davachi, L. (2011). What constitutes an episode in episodic memory? *Psychological Science,* 22(2):243 – 252.

Gershman, S. J., Schapiro, A. C., Hupbach, A., and Norman, K. A. (2013). Neural context reinstatement predicts memory misattribution. *Journal of Neuroscience,* 33(20):8590 – 8595.

Glenberg, A. M., Bradley, M. M., Kraus, T. A., and Renzaglia, G. J. (1983). Studies of the long-term recency effect: Support for a contextually guided retrieval theory. *Journal of Experimental Psychology: Learning, Memory & Cognition,* 12:413 – 418.

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., and Rubin, N. (2008). A hierarchy of temporal receptive windows in human cortex. *Journal of Neuroscience,* 28(10):2539 – 2550.

Honey, C. J., Thesen, T., Donner, T. H., Silbert, L. J., Carlson, C. E., Devinsky, O., Doyle, J. C., Rubin, N., Heeger, D. J., and Hasson, U. (2012). Slow cortical dynamics and the accumulation of information over long timescales. *Neuron,* 76:423 – 434.

Howard, M. W. and Eichenbaum, H. (2013). The hippocampus, time, and memory across scales. *Journal of Experimental Psychology: General.* In press.

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

Howard, M. W. and Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology,* 46:269 – 299.

Howard, M. W., Viskontas, I. V., Shankar, K. H., and Fried, I. (2012). Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus,* 22:1833 – 1847.

Howard, M. W., Youker, T. E., and Venkatadass, V. (2008). The persistence of memory: Contiguity effects across hundreds of seconds. *Psychonomic Bulletin & Review,* 15:58 – 63.

Hupbach, A., Gomez, R., Hardt, O., and Nadel, L. (2007). Reconsolidation of episodic memories: a subtle reminder triggers integration of new information. *Learning & Memory,* 14:47 – 53.

Hyman, J., Ma, L., Balaguer-Ballester, E., Durstewitz, D., and Seamans, J. (2012). Con- textual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences, USA,* 109(13):5086 – 5091.

Jenkins, L. J. and Ranganath, C. (2010). Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *Journal of Neuroscience,* In press.

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

Lerner, Y., Honey, C. J., Silbert, L. J., and Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *Journal of Neuroscience,* 31(8):2906 – 2915.

Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., and Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences, USA,* 109(47):19462 – 19467.

Manning, J. R., Polyn, S. M., Baltuch, G., Litt, B., and Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences, USA,* 108(31):12893 – 12897.

Manns, J. R., Howard, M. W., and Eichenbaum, H. (2007). Gradual changes in hippocampal activity support remembering the order of events. *Neuron,* 56(3):530 – 540.

Miller, E. K., Erickson, C. A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefontal cortex of the macaque. *Journal of Neuroscience,* 16:5154.

Moreton, B. J. and Ward, G. (2010). Time scale similarity and long-term memory for autobiographical events. *Psychonomic Bulletin and Review,* 17(4):510 – 515.

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

O'Reilly, R. C. and Rudy, J. W. (2001). Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. *Psychol Rev,* 108(2):311 – 345.

Polyn, S. M. and Kahana, M. J. (2008). Memory search and the neural representation of context. *Trends in Cognitive Sciences,* 12:24 – 30.

Polyn, S. M., Norman, K. A., and Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review,* 116(1):129 – 156.

Radvansky, G. A. and Copeland, D. E. (2006). Walking through doorways causes for- getting: situation models and experienced space. *Memory & Cognition,* 34(5):1150 – 1156.

Ranganath, C. and Ritchey, M. (2012). Two cortical systems for memory-guided behavior. *Nature Reviews Neuroscience,* 13:713 – 726.

Reynolds, J. R., Zacks, J. M., and Braver, T. S. (2007). A computational model of event segmentation from perceptual prediction. *Cognitive Science,* 31:613 – 643.

Sahakyan, L., Delaney, P. F., Foster, N. L., and Abushanab, B. (2013). List-method directed forgetting in cognitive and clinical research: a theoretical and methodological review. *Psychology of Learning and Motivation,* 59:131 – 189.

Schwartz, G., Howard, M. W., Jing, B., and Kahana, M. J. (2005). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychological Science,* 16:898 – 904.

Sederberg, P. B., Gershman, S. J., Polyn, S. M., and Norman, K. A. (2011). Human memory reconsolidation can be explained using the temporal context model. *Psychonomic Bulletin & Review,* 18(3):455 – 468.

Sederberg, P. B., Howard, M. W., and Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review,* 115(4):893 – 912.

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

Shankar, K. H., Jagadisan, U. K. K., and Howard, M. W. (2009). Sequential learning using temporal context. *Journal of Mathematical Psychology,* 53:474 – 485.

Shimamura, A. P. (1994). Memory and frontal lobe function. In Gazzaniga, M. S., editor, *The Cognitive Neurosciences,* pages 803 – 815. MIT Press, Cambridge, MA.

Smith, S. M. and Vela, E. (2001). Environmental context-dependent memory: a review and meta-analysis. *Psychonomic Bulletin and Review,* 8(2):203 – 220.

Swallow, K. M., Barch, D. M., Head, D., Maley, C. J., Holder, D., and Zacks, J. M. (2011). Changes in events alter how people remember recent information. *J Cogn Neurosci,* 23(5):1052 – 1064.

Swallow, K. M., Zacks, J. M., and Abrams, R. A. (2009). Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General,* 138(2):236 – 257.

Turk-Browne, N. B., Simon, M. G., and Sederberg, P. B. (2012). Scene representations in parahippocampal cortex depend on temporal context. *Journal of Neuroscience,* 32(21):7202 – 7207.

Wheeler, M. A., Stuss, D. T., and Tulving, E. (1995). Frontal lobe damage produces episodic memory impairment. *Journal of the International Neuropsychological Society,* 1:525 – 536.

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., and Reynolds, J. R. (2007). Event perception: a mind-brain perspective. *Psychol Bull,* 133:273 – 293.

Zwaan, R. A. and Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin,* 123(2):162 – 185.