Replay Comes of Age

David J. Foster

Department of Psychology and Helen Wills Neuroscience Institute, University of California, Berkeley, California 94720; email: davidfoster@berkeley.edu

ANNUAL REVIEWS **Further**

Click here to view this article's online features:

- Download figures as PPT slides
- Navigate linked references
- Download citations
- Explore related articles
- Search keywords

Annu. Rev. Neurosci. 2017. 40:581–602

The *Annual Review of Neuroscience* is online at neuro.annualreviews.org

https://doi.org/10.1146/annurev-neuro-072116-031538

Copyright © 2017 by Annual Reviews.
All rights reserved

Keywords

hippocampus, place cell, replay, memory

Abstract

Hippocampal place cells take part in sequenced patterns of reactivation after behavioral experience, known as replay. Since replay was first reported, nearly 20 years ago, many new results have been found, necessitating revision of the original interpretations. We review some of these results with a focus on the phenomenology of replay.

Contents

1. INTRODUCTION

Hippocampal replay is one of the best-known specimens in the menagerie of modern systems neuroscience phenomena. The story usually told is this: Neurons in the hippocampus replicate, during sleep, patterns of activity from previous wakeful experience in a temporally compressed form. This narrative was established about 20 years ago and suggested a mechanism for systems memory consolidation—the putative process by which memories, initially dependent upon the hippocampus, become independent of the hippocampus with time, supported instead by cortical circuits (Squire 1992, Squire et al. 2015). However, as more results have been gathered, it has become clear that many of these initially described properties of replay do not hold in general. The purpose of this review then is to provide an update of results in the field of hippocampal replay. Several recent reviews have covered similar ground, including some shorter reviews (Carr et al. 2011, O'Neill et al. 2010) and a very comprehensive review of related and overlapping phenomena in the local hippocampal electroencelphalogram (EEG) (Buzsaki 2015). However, the particular and narrower focus of this review is on the detailed phenomenology of hippocampal replay itself and how a properly updated picture changes the implications for how replay contributes to hippocampal function.

2. HISTORICAL REPLAY

The hippocampus is one of the most intensively studied areas in the brain, arguably with good reason. Deep in the medial temporal lobe, its connections to multiple cortical and subcortical areas suggest an anatomical hub, it has a distinct role in memory, it has a relatively well understood internal circuitry, its synapses are highly plastic, and its neurons are arranged conveniently for electrophysiological recordings (Andersen et al. 2007).

The study of replay is based on the responses of individual hippocampal neurons during free behavior that respond reliably to place (O'Keefe & Dostrovsky 1971, Wilson & McNaughton 1993). These responses have been interpreted as a cognitive map (O'Keefe & Nadel 1978), one component of a network of brain areas with a variety of forms of spatial representation that can support navigation (Moser et al. 2008). For example, grid cells in entorhinal cortex represent current location using spatial frequencies and spatial phase information, and cells in several areas represent current head direction in a manner resembling an internal compass. So the starting point for our review is an observation. All these responses, although highly suggestive of the processing of spatial information for the purposes of navigation, are limited in their application to navigation in the same way. Responses are bound to the features that define them; for example,

according to the classical description, a place cell may not fire outside its place field. The result is that exploration of the map is limited to the current location. A complementary observation can be made taking a mnemonic perspective. Whereas the form of memory most closely associated with the hippocampus, episodic memory, has been characterized as mental time travel (Tulving 2002), place cells and grid cells and the rest appear confined to the present. Thus, whether from the perspective of navigation or that of memory, single-unit responses in the hippocampus and related areas appear oddly ill-suited to their task.

The key question was, could hippocampal units be reactivated away from the original stimulus? An early study reported that those place cells active in an environment, which typically comprise 30–50% of all place cells (Guzowski et al. 1999, Wilson & McNaughton 1993), were more active in subsequent sleep than those that had not been active in the environment (Pavlides & Winson 1989). However, the key experimental development in uncovering hippocampal unit activity not bound by present stimuli or behavior was the following. By recording the responses of multiple hippocampal units simultaneously, the coordination between neurons became measurable, as a separate channel of information processing beyond responses to an external stimulus. The tetrode electrode design, comprised of four microelectrode wires twisted together, allowed extracellular signal separation superior to that of a single electrode and the separation through triangulation of signals from as many as 25 different units per tetrode. Lightweight microdrives were developed to allow the independent depth-adjustment of many tetrodes simultaneously. These methods enabled measurement of the coactivity of pairs of place cells, to find that cells with overlapping place fields during behavior had elevated coactivity during subsequent sleep (Wilson & McNaughton 1994). This result was very important in establishing a higher order of organization beyond the firing of individual cells and in revealing the central importance of temporal coincidence of activity. Several further studies elaborated the theme of pairwise reactivation (Gerrard et al. 2001, Kudrimoti et al. 1999, Skaggs & McNaughton 1996), but arguably the next breakthrough finding was that of sequenced activity across four or more cells matching behavioral sequences (Lee & Wilson 2002, Louie & Wilson 2001, Nadasdy et al. 1999). Together, these results established the classical picture of hippocampal replay: After multiple repetitions of an experience (the repeated running laps), the exact same sequential pattern of activity was repeated during subsequent sleep. The discrepancy between the short duration of replay events (50–150 ms) and the duration of running events (several seconds) was interpreted as temporal compression. However, these classical features of replay do not hold in general, as we see below.

3. REPLAY AWAKES

An early study based on pairwise reactivation raised the possibility that replay-related activity might occur during wakefulness (Kudrimoti et al. 1999), but it was a succession of later results from several different labs that established highly robust replay sequences across tens of simultaneously recorded cells occurring during the awake state (Csicsvari et al. 2007, Davidson et al. 2009, Diba & Buzsaki 2007, Foster & Wilson 2006, Gupta et al. 2010, Karlsson & Frank 2009). These reports established several new features. The immediacy of awake replay seconds after experience suggested new functional possibilities, such as a role in initial learning, that would not be open to a mechanism delayed by hours until sleep. Further, although the tasks involved the same sort of repetition of running laps as in previous sleep replay studies (e.g., Lee & Wilson 2002), awake replay could be observed after individual laps, when there were typically multiple replay events (Foster & Wilson 2006). Replays were even observed after the first lap on a novel track (Foster & Wilson 2006). This inverted the story from sleep: Instead of multiple repetitions of experience generating a replay, rather a single experience gave rise to multiple replays. The most striking difference was

also the most unexpected. Awake replay, although on a similar temporally compressed timescale as in sleep, was often in the reverse order to the order in which the fields were arranged along the track (Csicsvari et al. 2007, Diba & Buzsaki 2007, Foster & Wilson 2006, Gupta et al. 2010, Karlsson & Frank 2009). On the linear track, this reverse ordering is potentially confounded with the order of place fields along the immediately subsequent outgoing lap; however, because of the directional preference of place cells, these two kinds of sequences could be distinguished (Davidson et al. 2009, Foster & Wilson 2006). The reverse replay after a single lap was also important, as at that moment the animal had not yet run in the return direction, and this line of evidence was later strengthened by the use of a running task in which rats only ever ran in one direction (Gupta et al. 2010). The particular meaning of reverse ordering is considered further in Section 7, but for now we note that reverse replay represents an abstraction from the original experience and a challenge to the original conception of replay as recapitulation.

An interesting question is, why was awake replay not discovered before? One subtle difference from previous studies was the avoidance of overtraining. Rats must be trained to run along a track (unlike mice, for example), but with too much training, their behavior can become stereotyped, with very little pausing between laps. This maximizes the measurement of run-time spiking for a given period of time but reduces to nil the amount of time available for replay. The most critical difference, however, was an increase in the number of tetrodes aimed at the CA1 region of the hippocampus, from 12 to 18. With few cells, sleep replay is easier to observe because of the longer duration of a sleep session compared to the few minutes in total of stopping during a track session, thus increasing the amount of data and so its statistical power. However, increasing the numbers of cells recorded simultaneously led to a qualitative transition from a statistical demonstration of the probability of ordered sequences to the visual pop-out demonstration of large numbers of ordered replay events. An analogy might be made to a microscope, in which structures become visible only at a certain resolution. In the same way, successive increases in tetrode number up to 42 have revealed progressively more intricate replay-related phenomena (**Figure 1**).

Several further features of awake replay were established. Research using the same linear track task as previous studies used showed that both forward and reverse replay occurred during wakefulness (Diba & Buzsaki 2007). Thus, the conspicuous ordering difference with previously reported replay in sleep was not after all fundamental. Another result that brought awake and sleep replay closer together was the finding that awake replay does not always relate to the current environment but can be remote replay of a different environment (Karlsson & Frank 2009). By this definition, sleep replay is almost always remote replay, as it is usually recorded while the animal occupies a small and visually isolated holding area. Nevertheless, these remote replays are relatively uncommon in the awake state and may correlate inversely with the degree of engagement the animal has in the current task. The more common observation is that awake replay tends to reflect behavioral trajectories in the current environment and furthermore tends to start in the actual current location of the animal (Karlsson & Frank 2009, Pfeiffer & Foster 2013a).

A particularly fascinating study of awake replay (Davidson et al. 2009) used a 10 m–long linear track, which was much longer than the approximately 2 m–long linear tracks used in previous replay studies, albeit twisted around so as to fit within the reach of the standard length recording tethers. They reported that replays appeared to depict behavioral trajectories at a constant speed, that is, with duration proportional to distance depicted. Indeed, the authors reported durations much longer than any reported previously, up to 700 ms, matching the unusually long track length. There was a very tight relationship between replay duration and the length of the depicted trajectory; although for technical reasons, it is possible that some of this relationship might have resulted from incomplete measurement of replay, further support for at least a monotonic relationship between replay duration and represented distance came from another study in which complete

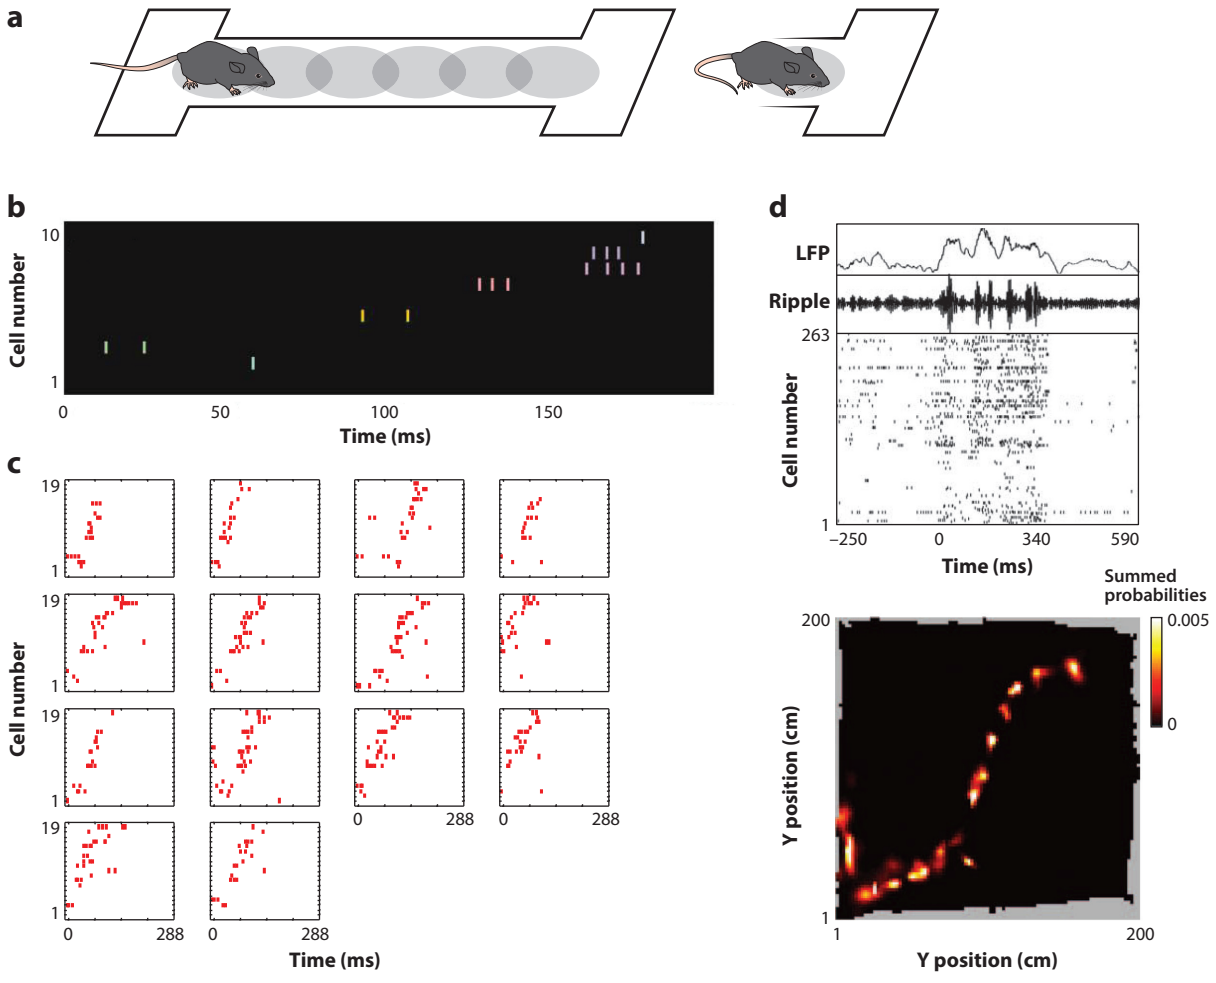

Figure 1

The evolution of replay data as a function of numbers of cells recorded simultaneously. (*a*) Schematic of a simple replay experiment: (*Left*) A rat runs along a linear track through the place fields of several hippocampal place cells (*ellipses*). (*Right*) Replay phenomena can be observed during subsequent rest periods, either on the track or in a separate sleep area. (*b*) An early example of replay, recorded during sleep. Ten simultaneously recorded cells were ordered by the positions of their place field peaks along a linear track, and events were detected in which at least four cells fired in the same order. Approximately six such events on average were found in sleep sessions lasting up to 2 h. Panel adapted from Lee & Wilson (2002) with permission. (*c*) Example replays observed across 19 simultaneously recorded place cells in awake rats. These 14 replays were observed within a single stopping period lasting less than a minute. Each individual event could be analyzed using a simple rank correlation between position and cell order. Panel adapted from Foster & Wilson (2006). (*d*) Greater place cell coverage has enabled measurement of replay during a spatial memory task in a 2-m × 2-m open arena. (*Top*) Raw local field potential (LFP) trace recorded from the hippocampus, ripple-filtered LFP trace, and spike rasters from 263 simultaneously recorded hippocampal place cells during a replay event. (*Bottom*) The event depicted a two-dimensional replay across the open-field environment. Posterior probabilities of position found using Bayesian decoding methods (see Section 4 of main text) were concentrated sufficiently to allow the whole event to be depicted spatially, collapsing across time. This format reveals gaps in the representation where the replay jumps from one hover point to another as part of its intrinsic dynamics. Panel adapted from Pfeiffer & Foster (2015).

replays of behavioral trajectories of different lengths were compared (Wu & Foster 2014). A picture emerges of replay representing not temporally compressed behavioral episodes but rather traversable distances and thus a model of the world as opposed to the veridical recording of experience. Also, by mapping distances into durations, replay might support the evaluation of alternative paths in terms of length, for navigation.

The idea that replay represents a model of the world rather than stored experiences has received support from several further studies. First, a direct study of whether replay would represent the most recent of different preceding experiences demonstrated conclusively that this was not the case (Gupta et al. 2010). Second, a study using a somewhat unusually shaped maze demonstrated that replay captured several aspects of the unique maze topology (Wu & Foster 2014). Finally, studies have now shown that replays can piece together previous experiences into whole trajectories that have not been previously experienced together (Gupta et al. 2010, Pfeiffer & Foster 2013a). The model of the world that informs replay is developed very rapidly, after only one to two experiences (Foster & Wilson 2006, Wu & Foster 2014).

A final study to mention in this section is a causal study in which the hippocampus was electrically stimulated to disrupt processing, either coincidentally or noncoincidentally with events in the local field potential known to co-occur with replay, while animals performed an alternation task on a W-shaped maze (Jadhav et al. 2012). This task required memory while on the central arm, where the sensory stimuli did not indicate which choice was correct, but not on the outer arms, where the choice was always to enter the central arm. Correspondingly, disruption of the replay-associated events impaired choice on the central arm but not the outer arms, revealing a role for awake replay in what is known as the working memory part of the task. The experimental design is vulnerable to the classic criticism of the single dissociation: that one of the tasks (working memory) may have been more demanding than the other (reference memory) (Shallice 1988). Nevertheless, the result is consistent with the idea that awake replay plays a role other than memory consolidation—which by definition affects only memory in the future—by supporting working memory function in the present. The representation of distances within replay might usefully contribute to such a function.

4. PREPLAY, MORE PREPLAY, AND STATISTICS

A fundamental assumption in the consideration of replay function is that replays should represent learned information—it is difficult to imagine a role in navigation or memory if this is not so. So a major challenge to this conception was a series of reports of sequenced activations exactly resembling replays that occurred prior to the experience being replayed (Dragoi & Tonegawa 2011, 2013a,b). To distinguish these events from replays, they were termed preplays, but in fact this obscures a difficulty. If any given sequence can be preplayed, it is unclear how a replay is any different, as it might also have occurred prior to experience, albeit unobserved during the short experimental sample of pre-experience time. In this way, the existence of preplay undermines the existence of replay.

The specific claim was that sequences of cells prior to experience predicted the relative position of the constituent cells' place fields during subsequent behavior. This implies an a priori ordering of cells into sequences, which governs the recruitment of cells to places during subsequent behavior. An analogy would be frames of movie tape, which preexist in a defined order and become associated with experiences—that is, have images deposited onto them—only subsequently during filming. So learning still occurs, in the associations between experiences (images) and place cells (frames), but the structure is heavily restricted, being constrained to a single timeline. The metaphor isn't

perfect: It should not be possible to predict how much tape is needed prior to an experience, whereas in the preplay papers, the preplay events just like replay events spanned the track entirely.

We addressed preplay recently in a paper examining sleep prior to the exploration of several linear tracks and presented three arguments against the existence of preplay (Silva et al. 2015). First, despite using many more tetrodes than the original reports, we found no evidence for preplay in our own data. Second, we identified certain methodological issues that led us to question the original conclusions. Specifically, the statistical approach to preplay is complicated by the existence of two possible future trajectories that could be preplayed, corresponding to the two running directions along the linear track (**Figure 2**). Although the preplay papers recognized this factor, it was not processed in the correct way: Each putative preplay event was apparently assessed against the best correlated running direction, but the same procedure was not used on the shuffled data, which were instead yoked to whatever running direction was best for the associated preplay event. The telltale for this unfortunate error was a unimodal, zero-mean distribution of correlations for the shuffles, which is inconsistent with a maximization step as described above.

The third argument is perhaps the most important for future developments in the field. Increasingly, probabilistic decoding methods are being used for measuring replay phenomena (Davidson et al. 2009, Zhang et al. 1998). Spiking activity within a given time bin is converted into a posterior probability distribution over some variable, such as position. A full discussion of the issues involved with this approach could fill an entire review; here we simply note that although the calculated quantity is position, the true value of this variable is rarely of interest to the experimenter, and indeed can be extracted directly from the overhead camera. Rather, the decoding procedure acts as a model of how the brain itself might process hippocampal activity patterns, and errors, such as replay, are often interesting. For a historical perspective, the approach is close in spirit to the work of Georgopoulos (Georgopoulos et al. 1989). One way to conceptualize replay in the context of decoding is as a covert variable (Johnson et al. 2009).

In practice, individual replays of linear trajectories can be graphed as a two-dimensional matrix of probabilities with time along one axis and position along the other (**Figure 3**). There are excellent reasons to analyze replay in this way. Place fields come in a variety of shapes and sizes, and the activity of place cells during replays appears to preserve these characteristics, so that a single rank order parameter cannot capture the cell's contribution. By contrast, probabilistic decoding takes the whole place field into account. Further, where place field coverage is sparse, decoding can effectively interpolate using the tails of place fields where coverage is good. Alternatively, when the number of recorded cells is very large, plotting every cell's spike raster becomes messy, whereas the corresponding decoded replays become pin-sharp. Each decoded replay event can be assigned a score using a variety of metrics, including amount of probability along a line representing an idealized trajectory (Davidson et al. 2009), or a generalization of correlation weighted by the posterior probabilities (Wu & Foster 2014). The problem comes in determining the statistical significance of this score using Monte Carlo shuffle methods. With an ordered set of cells, shuffling (i.e., permuting) cell order is the obvious choice. With decoded replays, it is no longer clear. Every conceivable shuffle has some weakness (see **Figure 4** and the sidebar titled The Problem with Replay Decoding: Shuffles). A second laboratory recently reported preplay, using decoded events and the weighted correlation method correlating position and time (Grosmark & Buzsaki 2016). It turns out that in our paper we had also examined this measure, to find that even random data generated by us yielded significance, therefore spuriously by definition (Silva et al. 2015). Naturally, spurious significance was also found in our data during sleep periods preceding experience. These forms of spurious significance can be eliminated by using stricter definitions of replay, but these were not applied in any of the preplay reports.

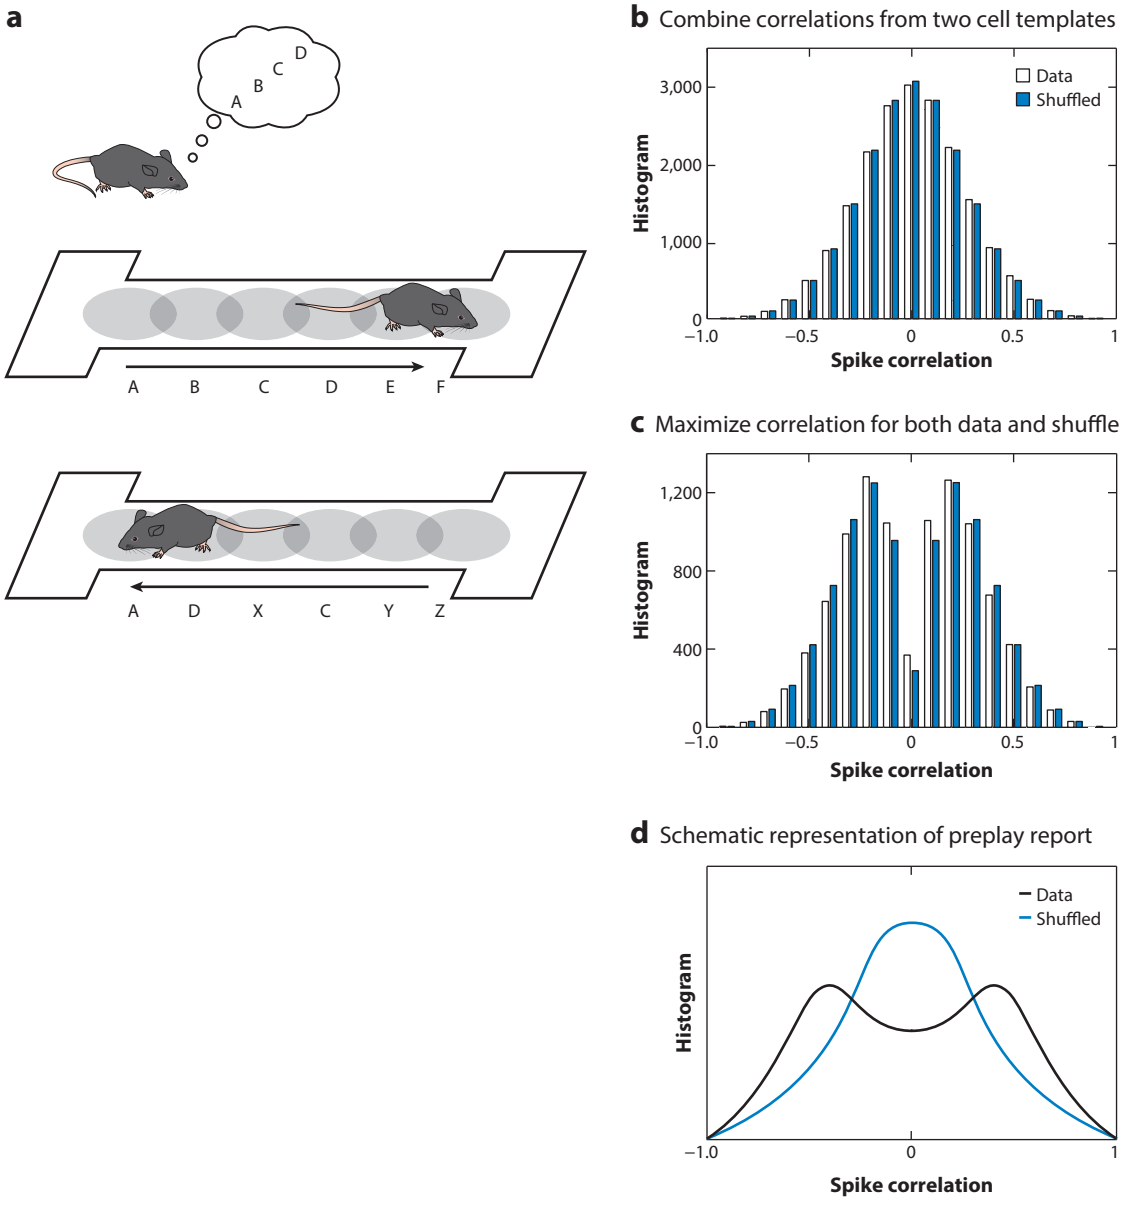

Figure 2

The measurement of preplay. (*a*) Schematic depicting hypothetical preplay prior to experience: There are two running directions that any given preplay could refer to that will not in general be ordered in the same way, and each gives rise to a template that events can be tested against. In the drawings, cells B, E, F, X, Y, and Z have place fields in only one direction, whereas cells A, C, and D have place fields in both directions. (*b*) When each potential preplay is assessed against both templates, there are twice as many comparisons as data, and shuffles are centered around zero correlation. In these data from Silva et al. (2015), the distribution of real data overlapped completely with shuffled data. (*c*) When each potential preplay is assessed against its best correlated template, there is only one comparison per event, and both shuffles and real data are pushed away from zero correlation, toward 1 or −1. Again, in these data from Silva et al. (2015), the two distributions were completely overlapping. (*d*) Schematic of the reported results from Dragoi & Tonegawa (2011). The number of comparisons was equal to the number of potential events, indicating that a selection of template was made for each potential event. However, the shuffles were centered on zero, whereas the real data were pushed towards −1 and 1. One erroneous way that this could have occurred is if the best template was chosen for each real event and the shuffles were yoked to this choice.

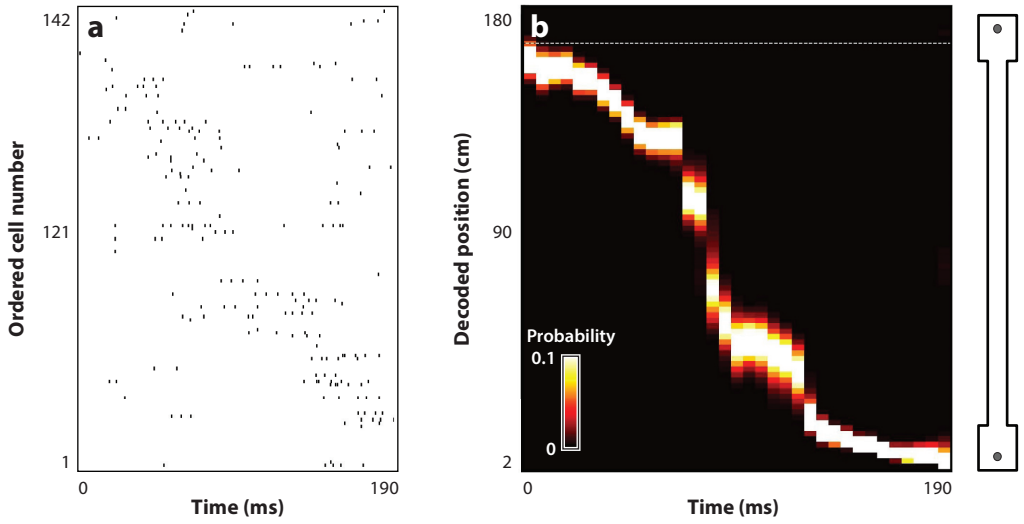

Figure 3

Bayesian decoding of replay. (*a*) A replay event across 142 place cells, ordered by the positions of their place field peaks along a linear track. Some degree of order to the spiking is clear, but the relationship is noisy. (*b*) The same event expressed using Bayesian decoding. Time and position are both binned (and arranged along the *x* and *y* axes, respectively), and a posterior probability is calculated for each joint position-time bin and displayed using a color scale. Although Bayesian decoding is often used as a framework for leveraging assumptions about cell spiking for more accurate estimation, in this case the only assumed knowledge for each cell is its place field. From this deliberate restriction follow a likelihood model with independence between cells and Poisson spike numbers, and a uniform prior. As shown, the posterior probabilities of position are extremely sharply defined during the replay event.

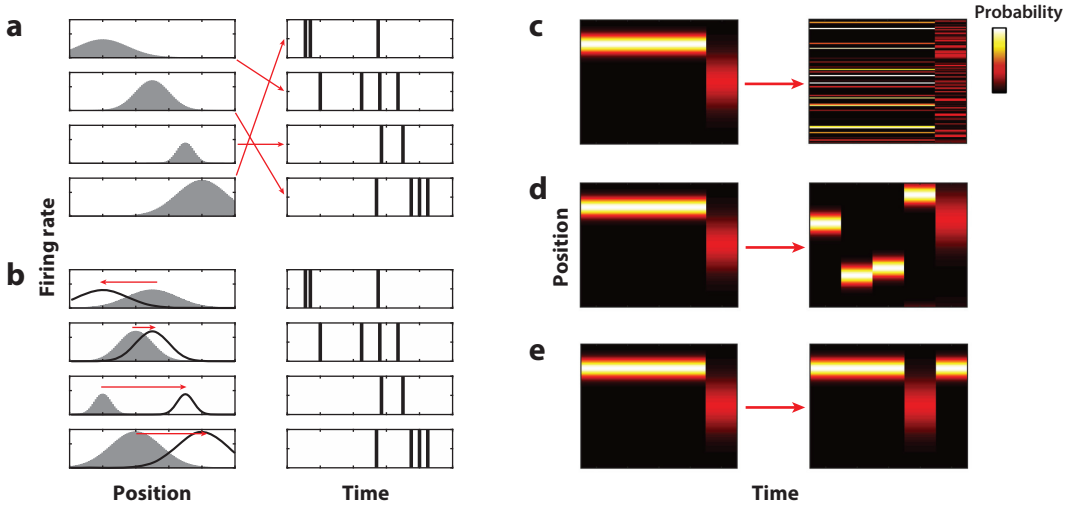

Figure 4

Schematics of different shuffling methods used for evaluating the significance of decoded replay (see the sidebar titled The Problem with Replay Decoding: Shuffles).

THE PROBLEM WITH REPLAY DECODING: SHUFFLES

Shuffling is a nonparametric statistical method for evaluating the significance of a data set, in which data labels are randomized to simulate chance data. Ideally, shuffles randomize the factor of interest (e.g., sequential order of cells) while leaving all other factors (e.g., temporal structure of spike trains) the same. In the case of decoded replay data, poor choices of shuffle can introduce randomness unrelated to the factor of interest, and any increase in the randomness of the shuffles is tantamount to an increase in the reported significance of the original data. We consider five example shuffle methods.

Pre probability methods apply shuffles to the raw data (i.e., place fields and spike trains) (**Figure 4*a*,*b***). The shuffled data are subsequently transformed into probabilities, making the effects of shuffling hard to interpret.

- Cell identity shuffle (**Figure 4*a***): The relationship between fields and spike trains is randomly shuffled. This produces combinations of spike train and place field that did not (and perhaps could not) exist in the real data, adding excessive noise to the shuffles.
- Place field shift (**Figure 4*b***): This method, explicitly devised to overcome the above problem with the cell identity shuffle, shifts the position of place fields independently for each cell. This is closer to a jitter than a shuffle, as positions can be generated that are not in the original data, with the result that excessive noise is introduced. Further, edge effects of the shift introduce further noise.

Post probability methods act on decoded posterior probabilities. We illustrate these methods with a correlated event (**Figure 4*c*–*e***) that could have been generated by a binary transition between two activity states (e.g., transition into or out of a sharp wave) and therefore is not necessarily indicative of a trajectory-depicting replay event.

- Position shuffle (**Figure 4*c***): This method randomly permutes rows (positions), with the result that smoothness in the representation of nearby positions is disrupted. Even the decoding from a single place cell's activity could lead to significance with this shuffle.
- Circular shift of position (**Figure 4*d***): This method preserves local smoothness in position by circularly shifting probabilities independently for each time bin. Like the place field shift, this can generate position representations (and edge effects) that do not exist in the original data, which adds excessive noise to the shuffles.
- Time shuffle (**Figure 4*e***): This method, a true shuffle, is unlikely to disrupt smoothness in the temporal domain because of the relatively large time bins usually used. Nevertheless, even this method will lead to reported significance in our test case when using, for example, the weighted correlation measure, because correlation does not by itself imply an extended sequence of activation across many neurons. The method is also underpowered, as with only a handful of time bins, the number of shuffles that can be generated is small.

We recommend the strategy of increasing the number of constraints to be satisfied jointly by true trajectory-depicting events, for example, high correlation combined with a restricted step size between successively represented positions. These constraints are arbitrary but can function as a threshold that, given sufficiently powered recordings, will be exceeded only by real instances of sequenced activity.

5. LEARNING TO PLAY

If replays are after all learned, the next question is how. NMDA receptor activation is a well-established mechanism of synaptic plasticity associated with memory formation (Morris 2013, Morris et al. 1986), and several studies have investigated its role in hippocampal reactivation. These findings converge on the result that, whether assessed as pairwise coordination or the presence of full replay sequences, NMDA receptor activation and associated downstream processes are necessary for the formation of the coordinated activity (Dupret et al. 2010, Gerrard et al. 2008,

McHugh et al. 1996, Silva et al. 2015, Suh et al. 2013). This dependence matches that of behavioral memory, in that it is specific to the encoding of replay and not retrieval (Silva et al. 2015). These results are in contrast to the effects that have been reported for the run-time place field responses of the cells. The literature is not completely settled: Reports vary from no effect of NMDA receptor antagonism on place field formation in a novel environment for at least an hour (Kentros et al. 1998) to reduced spatial definition of fields (McHugh et al. 1996, Silva et al. 2015). But across all reports, the effects on fields are much less striking than the effects on replay-related coordinated spiking. It is worth considering the differences between the two objects of measurement. Replay spikes represent an instantaneous sample from the whole population. Place fields are arguably more artificial, formed by averaging firing over multiple running laps while at the same time ignoring population effects. They are presumably less sensitive to small changes and of course not dependent on the number of cells recorded. Nevertheless, the location represented by a place cell may be, at least initially, solely a function of its inputs [which can be spatially tuned (e.g., combinations of grid cells) (Solstad et al. 2006)] without requiring learning, whereas synaptic plasticity may be required to stitch together the representations of places within the environment to reflect the spatial contingencies between them.

How does a neural network play a sequence? One class of models of sequence generation that does not even require learning uses transient disinhibition to expose successively lower levels of subthreshold drive onto place cells (Foster & Wilson 2006). This is similar to a well-known model of sequence generation during theta cycles (Mehta et al. 2002). The drive could come from subthreshold fields that span the entire environment smoothly (monotonically on each side of the peak) or from an activity trace (Buzsaki 1989). However, although intracellular recordings of place cells do in some cases show larger subthreshold fields than spiking fields, the data do not support monotonically peaking fields that span entire environments (Epsztein et al. 2011, Harvey et al. 2009). The activity trace model is unlikely, at least in its simplest form, because of the long periods, up to several minutes, during which awake replays can be observed after experience (Foster & Wilson 2006) and because replays do not in general replay the most recent experience (Gupta et al. 2010). Moreover, inhibition increases during replay events (Klausberger & Somogyi 2008), and the controlled and steady recruitment of spikes through a replay event does not suggest progressive disinhibition (e.g., Pfeiffer & Foster 2015).

The next simplest idea to account for neuronal sequences is something like an activity chain—the place cells representing one place connect to and drive the place cells representing the next place along. This could be accomplished within a single circuit containing recurrent connectivity, with CA3 being the obvious candidate. However, there are reasons to doubt this model. The first awake replay study showed that replays became stronger with each lap of experience on a novel track, recruiting more cells and more spikes per cell; however, the propagation speed of replay progressively slowed down across laps, with replay events on later laps taking approximately twice as long as early laps to depict the same trajectory (Foster & Wilson 2006). This is the opposite of what a simple chain mechanism would do, as increases in spiking, perhaps driven by increases in synaptic strength, ought to lead to faster synaptic integration and hence a faster moving sequence. Indeed, this logic is reflected in a recent prediction that replays will slow down eventually as the strength of replay memory decays (Schwindel & McNaughton 2011).

A more complex mechanism for replay sequence generation was proposed recently, based in part on recordings from a relatively large number of neurons as rats ran in the classic linear track task (Pfeiffer & Foster 2015). Although the behavioral trajectories taken by the animals progressed smoothly from one end of the track to the other, the increased cell resolution revealed that awake replays were strikingly discontinuous, alternately hovering at, and jumping between, several locations along the length of the track. One should naturally suspect a measurement artifact due to

insufficient place field coverage, but several analyses confirmed that given the actual place cells and fields, a smooth trajectory could have been decoded if the underlying represented trajectory had been smooth. It might have been the case that the hovering locations were of special significance to the animal, although nothing in the animals' behavior or in the environment indicated that this was so. The more disruptive interpretation is that the hippocampus was attempting to depict a smooth running trajectory, but this was the best approximation that it could produce. One possible explanation for this is as follows. Given the highly distributed nature of the hippocampal place code, with 30–50% of all hippocampal place cells recruited to each environment, out of presumably thousands of potential environments, the problem of creating a simple chain between place cells becomes analogous to that of storing sequences of distributed patterns in a recurrent neural network such as a Hopfield network (Hopfield 1982). As Hopfield noted, this architecture does not store sequences well, in contrast to the way in which the network is guaranteed to relax into a fixed-point attractor. A solution is to use rapid Hopfield dynamics to recover attractors representing discrete stored memory items, but embed this within a slower heteroassociative network that can recover the next item in sequence (Kleinfeld 1986, Sompolinsky & Kanter 1986). Interestingly, this scheme predicts jumpy sequences. The important point mechanistically is that this scheme requires two sets of synapses, fast and slow, which can be satisfied by a network architecture that combines multiple brain areas, such as CA3 for fast autoassociative sharpening of individual items, and a wider hippocampal loop for slower heteroassociative jumping between items (Lisman 1999, Lisman et al. 2005). Finally, the alternation between hovering and jumping was weakly locked to the slow gamma rhythm in the local hippocampal EEG, which is consistent with a prior report that ripple amplitude is also weakly locked to slow gamma (Carr et al. 2012). Ongoing work in my laboratory suggests that the slowing of replay with experience results from the strengthening of attractors (Feng et al. 2015).

A further question of interest is how reverse replay sequences can result from forward experience. The simplest models of sequence formation from spike timing–dependent plasticity (STDP) (Bi & Poo 1998) suggested that this is problematic (e.g., Mehta et al. 2002), as they were exquisitely sensitive to the order of firing of cells during encoding, but recent results in the STDP literature ameliorate this concern. In visual cortical slices (Seol et al. 2007) and cultured hippocampal neurons (Zhang et al. 2009), researchers have shown that the neuromodulatory environment can render the STDP rule temporally symmetric; for example, dopamine promotes symmetric potentiation in the hippocampus. Consider hippocampal place cells firing in two adjacent place fields, say, locations A and B. There will be thousands of A cells and thousands of B cells, both in CA3 and in CA1. Without specifying precisely the circuits in which A cells connect to B cells (given the uncertainty about this that was noted above), we can state unequivocally that some A cells will be presynaptic to some B cells, and some A cells will be postsynaptic to some B cells. Under a symmetric STDP rule, both sets of synapses may be potentiated by the experience of traveling from A to B (**Figure 5**). With symmetry breaking accomplished by diverse mechanisms such as short-term synaptic depression (York & van Rossum 2009) or spike adaptation (Ben-Yishai et al. 1997), such weights can generate place cell sequences in either direction (Romani & Tsodyks 2015). One important point to note is that, for any given environment, effective bidirectional connections between places represent only a subset of all possible connections and in this way capture the spatial topology. For example, A cells might be connected symmetrically with B cells, and B cells with C cells, but not A cells with C cells. Further, neuromodulatory control of plasticity might also be a mechanism by which novel transitions between locations are encoded but not familiar ones, so that replay captures a model of transition probabilities rather than being dominated by recency or familiarity.

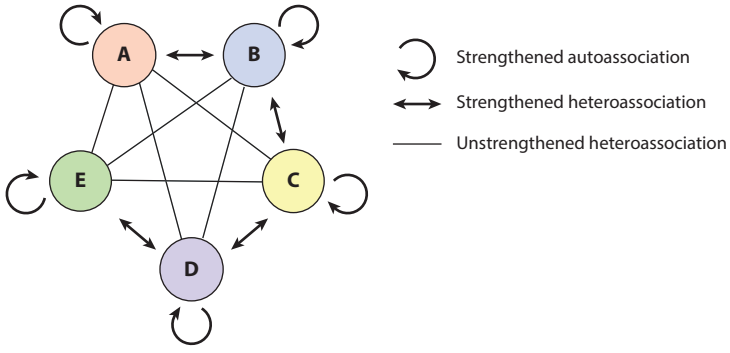

Figure 5

A schematic depiction of possible mechanisms underlying replay. A–E represent not individual place cells but rather populations of thousands of cells representing each of five locations A–E. During repeated experience, autoassociative connections between the cells representing individual locations become strengthened, supporting error correction of the representations of these locations during replay and also forcing replay to hover in these locations. Heteroassociative connections are also strengthened between connected locations and in both directions. These connections drive the movement of replay between locations and are strong in both forward and reverse directions, to support both forward and reverse replay (with some momentum supplied by adaptation mechanisms to prevent jiggling back and forth). However, not all pairs of locations are connected, and thus this pattern of connections defines an environmental topology.

6. REPLAY WITH A PURPOSE

At this point, several important features of awake replay have been established. It is learned, rapidly, but can express the learned information in an entirely different form, such as in reverse. The learning requires NMDA receptor activation. The resulting replays appear to reflect a model of which paths are traversable rather than recordings of specific experiences, and sampling from this model may be necessary for choice in spatial tasks. However, the picture remains of a rather spontaneous and undirected phenomenon, like echoes of the original experience. In this section, we find that this picture is misleading and that there is striking selectivity in terms of which replays occur and at what times.

Using a creative experimental design that switched between tasks each day, Gupta et al. (2010) demonstrated that the less frequently experienced trajectory out of two was the more frequently replayed. This pattern is consistent with a long-standing theoretical proposal that when training neural networks on new patterns, retraining on the old patterns must be interleaved so as to avoid catastrophic interference (McClelland et al. 1995), although the window over which this retraining can occur remains to be determined, with little direct evidence for retrieval of experiences from previous days. For the present discussion, it is noteworthy that the selection of replay appears selective rather than random. However, one of the many other interesting observations in this study was that during an alternation task, the selection of replay did not predict behavior. Similarly, in an earlier study from the same group, activity during the theta state was reported to represent places ahead of an animal paused at a choice point but also did not predict which of two trajectories would be taken (Johnson & Redish 2007). Theta activity is not replay but is organized similarly into sequences across different neurons. This tendency to replay both of two future choices with equal probability was also reported by others (Singer et al. 2013). This last example is particularly interesting because the occurrence of replay was found to be predictive of correct performance on the task. It is of course the case that a single, random sample of future outcome in a binary task

informs the decision fully. So although some evidence was offered for selection in replay, much of the passive, neutral character remained.

With these previous results in mind, we set out to test for replay selectivity by designing a spatial memory task in which the goal location was one of 36 similar locations in a large arena (Pfeiffer & Foster 2013a). The task used alternating phases of random foraging and goal-directed navigation to meet the twin demands of sampling of place fields from the entire environment and drawing on hippocampal memory. Given the 10–20 place cells needed to observe linear track replay, it was expected that roughly the square of this would be required in the square arena; this was accomplished by miniaturization of the microdrive, allowing an increase in the number of tetrodes deployed simultaneously, to 42, and a corresponding increase in cell yield of up to 250 neurons that were active in the environment. The major findings were that (*a*) two-dimensional replay sequences that depicted realistic behavioral trajectories occurred whenever the rat stopped; (*b*) replays just prior to goal-directed navigation tended to move toward and terminate at the remembered goal location and predicted the actual trajectory that would be taken; and (*c*) predictive replays occurred mostly during the phase of the task when they might be needed, that is, before goal-directed navigation rather than before random foraging. Note that these results require a modification to the models of sequence learning discussed in the previous section, as it is not obvious how goal selectivity can or should bias sequences. A curiosity in the data was that often during a stopping period, more than one trajectory would be replayed, but both could be predictive. For example, during a stop preceding goal-directed movement, the trajectory to the goal, but also a different trajectory, might be replayed. Sometimes, this other trajectory was relevant to where the animal explored immediately after visiting the goal, as if reflecting a multistep planning process. This offers a potential explanation for the failure to discern behaviorally predictive replay using a maze with only two choices: Replay often generates multiple alternatives while still being targeted toward future choices. Such selectivity is hidden when there are too few options.

An interesting question is, how much further can the problem-solving phenomenology of replay be taken? One is reminded of the work of behavioral psychologists in the early part of the twentieth century studying the problem-solving behavior of maze-running rats—for example, Maier (Maier 1929, Maier & Schneirla 1935), whose three-table experiment anticipated the shortcutting phenomenon of Gupta et al. (2010), and of course Tolman (Tolman 1949, Tolman & Honzik 1930), whose notion of inference inspired the changing goal aspect of the Pfeiffer study. The intriguing possibility exists of identifying Tolmanian inference in the activity of hippocampal neurons, for example, in more navigationally complex environments, but without behavioral confounds, as animals are typically stationary during replay events. Moreover, such activity is likely generated by a small number of neural circuits, offering a reduced biological substrate for studying high-level cognition. This at least is the promise. In the next section, more precise theoretical ideas are presented for the roles these various features of replay may have.

7. THEORETICAL PERSPECTIVES

Ideally, we would like to consider the functional relevance of replay to behavior in a precise way. What would help is a general and mathematical framework for how internal representations such as replay promote reasonable behavior. Fortunately, such a framework exists in the form of reinforcement learning (RL) (Sutton & Barto 1998). We begin by considering the appropriateness of RL for understanding replay and then turn to specific ways of formalizing some of the more recent findings in replay that are described in this review.

Replay expresses relationships between places, and whether the animal can travel between the places appears to be much more important for replay than whether the places are physically close

together in space. This is particularly striking in the Davidson et al. (2009) study, in which replay can take 700 ms to wend its way around a bendy track without ever seeming to jump between physically close (but uncrossable) points, but this has been just as true in every experiment we are aware of. This contrasts with the responses of grid cells, head direction cells, and the other recently discovered spatially responsive cells associated with the cognitive map (Moser et al. 2008), all of which have a more geometrical character. RL methods are often general learning methods that are capable of learning arbitrary relationships between elements and thus fit well with this characteristic of hippocampal replay.

The key issue for RL in sequential decision problems such as navigation is that planning by forward search of possible action sequences is prohibitively inefficient, owing to a combinatorial explosion in the number of possible future sequences. Many sequential decision problems can be expressed as Markovian decision problems or MDPs, in which states can be defined such that the optimal action choice in each state is independent of state history. For navigation to a fixed goal, locations are the states of an MDP, and so place cells have been modeled as providing states in an RL network for navigational learning (Foster et al. 2000). Location can become non-Markovian if other dependencies are introduced, such as dependency on the previous trial in an alternation task. A solution is to use states that capture both location and trial memory, and evidence suggests hippocampal cells develop similar representations (Foster & Knierim 2012, Frank et al. 2000, Wood et al. 2000). An efficient solution to MDPs is to run planning backward, starting from the goal and moving back along incoming trajectories. The reason why this is efficient is obvious: Unlike in the forward direction, the horizon in the reverse direction is always only a single step because of the Markovian property. What this means in practical terms is that an algorithm proceeding in reverse can learn in a single sweep about the multiple navigation problems corresponding to each of the states along the trajectory as a potential starting point.

These considerations suggested immediately that reverse replay could provide a mechanism for associating places with values, hypothetically through connections between place cells and downstream value-representing neurons such as those of ventral striatum (Roesch et al. 2009) or prefrontal cortex (Kennerley et al. 2011). A simple model was proposed in which reverse replays co-occur with a fast onset, slowly decaying reward-related signal, such as the phasic release of dopamine, with both signals converging in a value-representing area (Foster & Wilson 2006). Indeed, both nucleus accumbens (Lansink et al. 2009, Pennartz et al. 2004) and prefrontal cortex (Jadhav et al. 2016) exhibit activity that is temporally coordinated with hippocampal replay, as is the activity of dopamine neurons in the ventral tegmental area (Gomperts et al. 2015), in line with this model. The result of pairing reverse replay and reward signaling in the model is that place cells drive a gradient of reward expectation during subsequent navigation, which has been reported in subsequent studies (van der Meer & Redish 2009, 2011). Curiously, actual reverse replay takes a different form from those proposed previously in the computational literature. As noted above, reverse replay is not a buffer of recent experience, as within the same stopping period, multiple reverse trajectories can be replayed. However, reverse replay also does not resemble an algorithm such as value iteration (Sutton & Barto 1998) in which values spread out along multiple incoming trajectories simultaneously. This is particularly clear in two-dimensional replay, for which the focus on a single location is clearly preserved throughout the entirety of the replay event (Pfeiffer & Foster 2013a). Nevertheless, the actual form of reverse replay observed is perfectly compatible with the value learning function proposed.

The idea that reverse replay is for learning values from rewards predicts a special relationship between reward and reverse replay. This is interesting because phenomenologically, forward and reverse replay appear similar (Diba & Buzsaki 2007). This idea was tested recently by varying reward magnitude at certain locations in a very simple linear track running task (Ambrose et al.

2016), extending earlier work examining the effect of reward on replay-related events in the local hippocampal field potential (Singer & Frank 2009). Changing reward magnitude, either by increasing or removing it, increased or decreased the rate of reverse replay correspondingly but not the rate of forward replay, which remained constant. This was especially striking because both reverse and forward replays occurred intermingled within the same stopping periods and while animals were engaged in the same reward consumption behavior. Moreover, the key variable predicting the rate changes was not the absolute magnitude of reward but the change in reward from the previous experimental period and the relative amount compared to the total reward available. Both these effects resemble the pattern of dopamine signaling in response to rewards [prediction error effects (Schultz et al. 1997) and menu effects (Tobler et al. 2005), respectively], suggesting that dopamine might regulate replay, and reverse replay specifically. Recent results have shown that stimulation of dopamine neurons can promote hippocampal reactivation, although it is not known whether replay, and in particular reverse replay, is involved (McNamara et al. 2014).

If reverse replay is for learning values, a natural suggestion is that forward replay in contrast is for retrieving values, for the purposes of decision making, particularly at choice points. This suggestion has been made explicitly (Carr et al. 2011, Jadhav et al. 2012), and a similar role has been suggested for place cell sequences during theta, which sweep ahead of an animal's position along different future paths while the animal is paused at a choice point (Johnson & Redish 2007, Redish 2016). However, from an RL perspective, this intuitive forward replay idea is, ironically, the problematic one. As mentioned above, forward search is inefficient, and the notion that forward replay searches agnostically along all possible paths is already disproved by the data when the number of potential future choices is made very large (Pfeiffer & Foster 2013a). However, even if we posit that forward replay allows consideration of an animal's top two or three choices prior to movement, there is a problem. Tolman's (1948) cognitive map idea was famously criticized for leaving his rats buried in thought; similarly, we may consider a simple-minded but vivid example of the cost to a rat or mouse of replaying even two alternative future trajectories in succession. Thirty-six km/h is a fairly easy running speed for a cat, or 1 cm/ms. The time required to sit still and replay (or theta sweep) two different trajectories is at least several hundred milliseconds, as each event may take about 100 ms and there is often a gap of at least 100–200 ms between them. The cat will cover several meters before the rat decides on an escape route. Clearly the predictive nature of forward sequences demands an explanation, as does the necessary role of replay-related activity in immediate choice (Jadhav et al. 2012), but any explanation must pass the test of computational reasonableness.

The specific functions hypothesized for reverse and forward replay have different implications for their downstream circuits. As discussed above, reverse replay is likely to engage circuits related to reward and value. Forward replay may also engage circuits related to value; however, if forward replay is related to planning immediate behavior, several different strategies might be deployed, with different implications. The simplest strategy is what we have argued on computational grounds is the least likely, which would be that forward replay is used to evaluate and plan the immediate action based on retrieved information about the reward available at the end of the trajectory. Some short-term memory would be required across two or three successive replays to remember the retrieved outcomes of each and choose the best associated initial action. Speaking to this, we have not found evidence that the last trajectory to be replayed corresponds to the action chosen, as might be expected from an efficient deployment of this strategy. A different strategy is if extended forward replays reflect a planned sequence of actions, sometimes called options in the RL literature (Sutton et al. 1999). The use of these extended plans can reduce the dimensionality of a large state space by constraining exploration, which can be especially useful for generalization between similar tasks (Parr & Russell 1997, Singh 1992), as in our open-field navigation task

(Pfeiffer & Foster 2013a). The extended plans can be evaluated using the same RL networks as for ordinary actions (Sutton et al. 1999). However, under this interpretation, the forward replay represents an extended plan of action that the animal must commit to and evaluate upon its conclusion. It is not clear exactly how this commitment can occur, but it would involve a short-term memory, in this case lasting for the entire duration of the planned behavior and through to the outcome. It may be relevant that replays did in fact predict future behavior over at least 10 s in the open-field task, as if the behavior did tend to be constrained by the replay (Pfeiffer & Foster 2013a). A third very different possibility is that forward sequences have no immediate behavioral consequence at all but are part of a separate prediction system existing outside behavioral control, one that generates predictions of future behaviors and compares these with the actual future behaviors and the outcomes of those behaviors to learn a model of the animal's own behavior, although studies demonstrating the necessity and correlation of replay to performance in choice tasks would seem to contradict this (Jadhav et al. 2012, Singer et al. 2013). Determining which of these schemes is correct will require the precise manipulation of hippocampal replay contingent on trajectory and of the downstream circuits hypothetically involved.

8. CONCLUSIONS

We have presented many new results in the replay field that revise considerably the classical picture of replay. To summarize the most salient revisions: Replay occurs not only in sleep but also at any time that an awake rat stops running. It does not repeat activity patterns: It may, for example, ignore recent patterns or create new patterns never experienced before. It does not temporally compress experience, but rather draws speeded-up samples from a model of the spatial contingencies available to the moving animal, a cognitive map in the sense of Tolman (1948). It can be selective for the current goal location, and for current task phase, and predict the immediate future behavior of the animal. Indulgently speaking, it looks less like dreaming and more like thought.

Many omissions in this presentation demand explanation. For one, the treatment of replay has not explicitly contrasted awake and sleep replay. Many of the newer studies have been of awake replay only. However, we have deliberately avoided emphasizing this distinction because the fact is that precious little evidence separates the phenomenology of awake and sleep replay. Where we have examined sleep replay in our studies, activity patterns appear strongly conserved. For example, we found exactly the same sort of behavior-like trajectory events in sleep after a spatial memory task as occurred during task performance (Pfeiffer & Foster 2013b). Where differences exist, they tend to be more quantitative; for example, there is certainly more remote replay in sleep than in the awake state, but it can nevertheless occur in either state. Similarly, the presentation has not carefully demarcated replay results according to hippocampal subfield, for example, CA1 versus CA3 (there have not been comparable studies of replay in dentate gyrus or CA2 at this point). Different studies have in fact examined one or the other area, or both, but again, no phenomenological differences have yet been reported, and one study to show examples of replay in CA1 and CA3 recorded simultaneously is striking for how undifferentiated the activity in the two areas is (Karlsson & Frank 2009). So although hypotheses have been presented for the role of CA3 compared to other subregions, the distinction has not been made at the level of data.

A large and increasing literature has focused on the interaction between the hippocampus and other areas of the brain. Although some of this work was referred to, much has been omitted. Again, this has been deliberate, as the focus of this review has been to describe the phenomenology of hippocampal replay. Understanding how hippocampal activity patterns respond to, and influence, activity outside the hippocampus is clearly fundamental to future progress, but it is also worth proceeding cautiously. For example, interactions between hippocampal and cortical activity have

been interpreted as evidence in favor of a consolidation account of hippocampal replay (Squire et al. 2015), but if the hippocampus is not actually repeating patterns of previously experienced activity, then no amount of communication with cortex is going to support a consolidation mechanism, at least according to the simple models of consolidation that have been proposed.

Another large body of work that is potentially relevant but has not been much discussed here concerns the phenomenon of theta sequences. A huge amount of work has been published on theta-related activity in the past decade, including sequences. In many ways, activity during theta can resemble replay. However, there are conspicuous differences, such as the shortness of the sequences and, most importantly, the way that during theta there is marked interplay between incoming stimuli and internal dynamics. Clearly, replay also has some of that interplay, but it seems safe to suppose that replay is more largely representative of internal dynamics in the hippocampus than theta sequences are, and so the specific focus of this review on characterizing replay phenomenology can be seen to have an integral whole without the inclusion of theta sequences.

A final area of omission has been events in the hippocampal local field potential and, in particular, sharp-wave ripples. Expansion of the scope to include these events would easily increase the size of this review by an order of magnitude. A recent review gives a comprehensive account of work in this larger field (Buzsaki 2015) and in so doing highlights an interesting possible relationship to neuropsychiatric disease. Indeed, it is fascinating to consider that there might be diseases of replay. Many neuropsychiatric diseases have been linked with abnormal rest-state activity assessed by functional neuroimaging (Buckner et al. 2008), and so an effect on replay is plausible. In my own work, we found that an animal model of schizophrenia exhibited a remarkable pattern of hippocampal activity: Place cell firing was completely normal during running, but the moment the animal stopped, overall excitability increased sixfold, entirely because of overactive and overabundant sharp-wave ripple events during which replay was completely swamped by noise (Suh et al. 2013). In a follow-up study with a second, completely genetically distinct animal model of schizophrenia, we found a similar pattern of results, which we suggested might model how a disease with a heterogeneous etiology such as schizophrenia could exhibit convergence in information processing abnormalities that might account for common symptoms (Altimus et al. 2015). Given our revised understanding of hippocampal replay, it is interesting to note that psychosis has been proposed to result from an overdominant internal model (Dolan & Dayan 2013).

In summary, hippocampal replay is an extremely active area of investigation in many laboratories. Our understanding of the phenomenon has changed much in the past 20 years, leading necessarily to changes in interpretation and new insight into how the hippocampus supports memory. Additionally, hippocampal replay offers a window on a degree of complexity that may exist throughout the cortex, but which in most other brain areas is not yet amenable to large-scale simultaneous recording. This in the end may be its most important contribution.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The author would like to acknowledge the contribution of past and present members of the laboratory and funding from the National Institute of Mental Health, the Alfred P. Sloan Foundation, the Whitehall Foundation, the Brain & Behavior Research Foundation, the Brain Research

Foundation, the LuMind Research Down Syndrome Foundation, the McKnight Endowment Fund for Neuroscience, and the Johns Hopkins Brain Science Institute and Science of Learning Institute.

LITERATURE CITED

Altimus C, Harrold J, Jaaro-Peled H, Sawa A, Foster DJ. 2015. Disordered ripples are a common feature of genetically distinct mouse models relevant to schizophrenia. *Mol. Neuropsychiatry* 1:52–59

Ambrose RE, Pfeiffer BE, Foster DJ. 2016. Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron* 91:1124–36

Andersen P, Morris R, Amaral D, Bliss T, O'Keefe J, eds. 2007. *The Hippocampus Book*. Oxford, UK: Oxford Univ. Press

Ben-Yishai R, Hansel D, Sompolinsky H. 1997. Traveling waves and the processing of weakly tuned inputs in a cortical network module. *J. Comput. Neurosci.* 4:57–77

Bi GQ, Poo MM. 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18:10464–72

Buckner RL, Andrews-Hanna JR, Schacter DL. 2008. The brain's default network: anatomy, function, and relevance to disease. *Ann. N. Y. Acad. Sci.* 1124:1–38

Buzsaki G. 1989. Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience* 31:551–70

Buzsaki G. 2015. Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. *Hippocampus* 25:1073–188

Carr MF, Jadhav SP, Frank LM. 2011. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* 14:147–53

Carr MF, Karlsson MP, Frank LM. 2012. Transient slow gamma synchrony underlies hippocampal memory replay. *Neuron* 75:700–13

Csicsvari J, O'Neill J, Allen K, Senior T. 2007. Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. *Eur. J. Neurosci.* 26:704–16

Davidson TJ, Kloosterman F, Wilson MA. 2009. Hippocampal replay of extended experience. *Neuron* 63:497–507

Diba K, Buzsaki G. 2007. Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* 10:1241–42

Dolan RJ, Dayan P. 2013. Goals and habits in the brain. *Neuron* 80:312–25

Dragoi G, Tonegawa S. 2011. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* 469:397–401

Dragoi G, Tonegawa S. 2013a. Development of schemas revealed by prior experience and NMDA receptor knock-out. *eLife* 2:e01326

Dragoi G, Tonegawa S. 2013b. Distinct preplay of multiple novel spatial experiences in the rat. *PNAS* 110:9100–5

Dupret D, O'Neill J, Pleydell-Bouverie B, Csicsvari J. 2010. The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* 13:995–1002

Epsztein J, Brecht M, Lee AK. 2011. Intracellular determinants of hippocampal CA1 place and silent cell activity in a novel environment. *Neuron* 70:109–20

Feng T, Silva D, Foster DJ. 2015. *Mechanisms contributing to experience-dependent changes in the structure of hippocampal replay sequences*. Presented at Soc. Neurosci., Oct. 17–21, Chicago

Foster DJ, Knierim JJ. 2012. Sequence learning and the role of the hippocampus in rodent navigation. *Curr. Opin. Neurobiol.* 22:294–300

Foster DJ, Morris RGM, Dayan P. 2000. A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus* 10:1–16

Foster DJ, Wilson MA. 2006. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680–83

Frank LM, Brown EN, Wilson M. 2000. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* 27:169–78

Georgopoulos AP, Lurito JT, Petrides M, Schwartz AB, Massey JT. 1989. Mental rotation of the neuronal population vector. *Science* 243:234–36

Gerrard JL, Burke SN, McNaughton BL, Barnes CA. 2008. Sequence reactivation in the hippocampus is impaired in aged rats. *J. Neurosci.* 28:7883–90

Gerrard JL, Kudrimoti H, McNaughton BL, Barnes CA. 2001. Reactivation of hippocampal ensemble activity patterns in the aging rat. *Behav. Neurosci.* 115:1180–92

Gomperts SN, Kloosterman F, Wilson MA. 2015. VTA neurons coordinate with the hippocampal reactivation of spatial experience. *eLife* 4:e05360

Grosmark AD, Buzsaki G. 2016. Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science* 351:1440–43

Gupta AS, van der Meer MAA, Touretzky DS, Redish AD. 2010. Hippocampal replay is not a simple function of experience. *Neuron* 65:695–705

Guzowski JF, McNaughton BL, Barnes CA, Worley PF. 1999. Environment-specific expression of the immediate-early gene *Arc* in hippocampal neuronal ensembles. *Nat. Neurosci.* 2:1120–24

Harvey CD, Collman F, Dombeck DA, Tank DW. 2009. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461:941–46

Hopfield JJ. 1982. Neural networks and physical systems with emergent collective computational abilities. *PNAS* 79:2554–58

Jadhav SP, Kemere C, German PW, Frank LM. 2012. Awake hippocampal sharp-wave ripples support spatial memory. *Science* 336:1454–58

Jadhav SP, Rothschild G, Roumis DK, Frank LM. 2016. Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron* 90:113–27

Johnson A, Fenton AA, Kentros C, Redish AD. 2009. Looking for cognition in the structure within the noise. *Trends Cogn. Sci.* 13:55–64

Johnson A, Redish AD. 2007. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* 27:12176–89

Karlsson MP, Frank LM. 2009. Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* 12:913–18

Kennerley SW, Behrens TE, Wallis JD. 2011. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat. Neurosci.* 14:1581–89

Kentros C, Hargreaves E, Hawkins RD, Kandel ER, Shapiro M, Muller RV. 1998. Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280:2121–26

Klausberger T, Somogyi P. 2008. Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science* 321:53–57

Kleinfeld D. 1986. Sequential state generation by model neural networks. *PNAS* 83:9469–73

Kudrimoti HS, Barnes CA, McNaughton BL. 1999. Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J. Neurosci.* 19:4090–101

Lansink CS, Goltstein PM, Lankelma JV, McNaughton BL, Pennartz CM. 2009. Hippocampus leads ventral striatum in replay of place-reward information. *PLOS Biol.* 7:e1000173

Lee AK, Wilson MA. 2002. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36:1183–94

Lisman JE. 1999. Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. *Neuron* 22:233–42

Lisman JE, Talamini LM, Raffone A. 2005. Recall of memory sequences by interaction of the dentate and CA3: a revised model of the phase precession. *Neural Netw.* 18:1191–201

Louie K, Wilson MA. 2001. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29:145–56

Maier NRF. 1929. Reasoning in white rats. *Comp. Psychol. Monogr.* 6:1–93

Maier NRF, Schneirla TC. 1935. *Principles of Animal Psychology*. New York: McGraw-Hill

McClelland JL, McNaughton BL, O'Reilly RC. 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. *Psychol. Rev.* 102:419–57
McHugh TJ, Blum KI, Tsien JZ, Tonegawa S, Wilson MA. 1996. Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell* 87:1339–49
McNamara CG, Tejero-Cantero A, Trouche S, Campo-Urriza N, Dupret D. 2014. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nat. Neurosci.* 17:1658–60
Mehta MR, Lee AK, Wilson MA. 2002. Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417:741–46
Morris RGM. 2013. NMDA receptors and memory encoding. *Neuropharmacology* 74:32–40
Morris RGM, Anderson E, Lynch GS, Baudry M. 1986. Selective impairment of learning and blockade of long-term potentiation by an *N*-methyl-D-aspartate receptor antagonist, AP5. *Nature* 319:774–76
Moser EI, Kropff E, Moser MB. 2008. Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* 31:69–89
Nadasdy Z, Hirase H, Czurko A, Csicsvari J, Buzsaki G. 1999. Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* 19:9497–507
O'Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34:171–75
O'Keefe J, Nadel L. 1978. *The Hippocampus as a Cognitive Map*. London: Clarendon
O'Neill J, Pleydell-Bouverie B, Dupret D, Csicsvari J. 2010. Play it again: reactivation of waking experience and memory. *Trends Neurosci*. 33:220–29
Parr R, Russell S. 1997. Reinforcement learning with hierarchies of machines. *Adv. Neural Inf. Process. Syst.* 10:1043–49
Pavlides C, Winson J. 1989. Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.* 9:2907–18
Pennartz CM, Lee E, Verheul J, Lipa P, Barnes CA, McNaughton BL. 2004. The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. *J. Neurosci.* 24:6446–56
Pfeiffer BE, Foster DJ. 2013a. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497:74–79
Pfeiffer BE, Foster DJ. 2013b. *Place-cell sequences depict behaviorally relevant trajectories during sleep*. Presented at Soc. Neurosci., Nov. 9–13, San Diego
Pfeiffer BE, Foster DJ. 2015. Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science* 349:180–83
Redish AD. 2016. Vicarious trial and error. *Nat. Rev. Neurosci.* 17:147–59
Roesch MR, Singh T, Brown PL, Mullins SE, Schoenbaum G. 2009. Ventral striatal neurons encode the value of the chosen action in rats deciding between differently delayed or sized rewards. *J. Neurosci.* 29:13365–76
Romani S, Tsodyks M. 2015. Short-term plasticity based network model of place cells dynamics. *Hippocampus* 25:94–105
Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science* 275:1593–99
Schwindel CD, McNaughton BL. 2011. Hippocampal-cortical interactions and the dynamics of memory trace reactivation. *Prog. Brain Res.* 193:163–77
Seol GH, Ziburkus J, Huang S, Song L, Kim IT, et al. 2007. Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. *Neuron* 55:919–29
Shallice T. 1988. *From Neuropsychology to Mental Structure*. Cambridge, UK: Cambridge Univ. Press
Silva D, Feng T, Foster DJ. 2015. Trajectory events across hippocampal place cells require previous experience. *Nat. Neurosci.* 18:1772–79
Singer AC, Carr MF, Karlsson MP, Frank LM. 2013. Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* 77:1163–73
Singer AC, Frank LM. 2009. Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron* 64:910–21
Singh SP. 1992. Reinforcement learning with a hierarchy of abstract models. *Proc. Tenth Natl. Conf. Artif. Intell., San Jose, CA, July 12–16*, pp. 202–7. Menlo Park, CA: AAAI Press

Skaggs WE, McNaughton BL. 1996. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271:1870–73

Solstad T, Moser EI, Einevoll GT. 2006. From grid cells to place cells: a mathematical model. *Hippocampus* 16:1026–31

Sompolinsky H, Kanter I. 1986. Temporal association in asymmetric neural networks. *Phys. Rev. Lett.* 57:2861–64

Squire LR. 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99:195–231

Squire LR, Genzel L, Wixted JT, Morris RG. 2015. Memory consolidation. *Cold Spring Harb. Perspect. Biol.* 7:a021766

Suh J, Foster DJ, Davoudi H, Wilson MA, Tonegawa S. 2013. Impaired hippocampal ripple-associated replay in a mouse model of schizophrenia. *Neuron* 80:484–93

Sutton RS, Barto AG. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press

Sutton RS, Precup D, Singh S. 1999. Between MDPs and semi-MDPs: a framework for temporal abstraction in reinforcement learning. *Artif. Intell.* 112:181–211

Tobler PN, Fiorillo CD, Schultz W. 2005. Adaptive coding of reward value by dopamine neurons. *Science* 307:1642–45

Tolman EC. 1948. Cognitive maps in rats and men. *Psychol. Rev.* 55:189–208

Tolman EC. 1949. *Purposive Behavior in Animals and Men*. Berkeley: Univ. Calif. Press

Tolman EC, Honzik CH. 1930. "Insight" in rats. *Univ. Calif. Publ. Psychol.* 4:215–32

Tulving E. 2002. Episodic memory: from mind to brain. *Annu. Rev. Psychol.* 53:1–25

van der Meer MAA, Redish AD. 2009. Covert expectation-of-reward in rat ventral striatum at decision points. *Front. Integr. Neurosci.* 3:1

van der Meer MAA, Redish AD. 2011. Theta phase precession in rat ventral striatum links place and reward information. *J. Neurosci.* 31:2843–54

Wilson MA, McNaughton BL. 1993. Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–58

Wilson MA, McNaughton BL. 1994. Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–79

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. 2000. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* 27:623–33

Wu X, Foster DJ. 2014. Hippocampal replay captures the unique topological structure of a novel environment. *J. Neurosci.* 34:6459–69

York LC, van Rossum MC. 2009. Recurrent networks with short term synaptic depression. *J. Comput. Neurosci.* 27:607–20

Zhang JC, Lau PM, Bi GQ. 2009. Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. *PNAS* 106:13028–33

Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ. 1998. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* 79:1017–44