

Progress and issues in second-order analysis of hippocampal replay

Matthijs A. A. van der Meer^{1*}, Caleb Kemere^{2*}, Kamran Diba^{3*}

¹Dept. of Psychological and Brain Sciences, Dartmouth College, ²Dept. of Electrical and Computer Engineering, Rice University, ³Dept. of Anesthesiology, University of Michigan.

*All authors contributed equally. Correspondence: mvdM@dartmouth.edu, caleb.kemere@rice.edu, kdiba@umich.edu

Keywords

Replay, reactivation, sharp-wave ripple, decoding, sequence analysis

Funding

HFSP RGY0088/2014 (MvdM & CK), NSF CAREER IOS-1844935 (MvdM), NINDS 1R01NS115233 (CK & KD), NIMH 1R01MH109170 (KD)

Abstract (1459/2000 characters)

Patterns of neural activity that occur spontaneously during sharp wave-ripple events (SWRs) in the hippocampus are thought to play an important role in memory formation, consolidation and retrieval. Typical studies examining the content of SWRs seek to determine whether the identity and/or temporal order of cell firing is different from chance. Such “first-order” analyses are focused on a single time point and template (map), and have been used to show, for instance, the existence of preplay. The major methodological challenge in first-order analyses is the construction and interpretation of different chance distributions. In contrast, “second-order” analyses involve a comparison of SWR content between different time points, and/or between different templates. Typical second-order questions include rigorous tests of experience-dependence (replay) that compare SWR content before and after experience, and comparisons or replay between different arms of a maze. Such questions entail additional methodological challenges that can lead to biases in results and associated interpretations. We provide an inventory of analysis challenges for second-order questions about SWR content, and suggest ways of preventing, identifying, and addressing possible analysis biases. Given evolving interest in understanding SWR content in more complex experimental scenarios and across different timescales, we expect these issues to become increasingly pervasive.

Introduction

The hippocampus spontaneously generates spike sequences whose firing order corresponds to the order observed during behavior. These ordered spike sequences occur during specific time windows identified by sharp waves in the stratum radiatum of the CA1 region, along with fast ripple oscillations in the CA1 pyramidal layer (hence known as sharp wave-ripples, SWRs; Buzsaki 2015). SWRs and their temporally ordered activity are not only a strikingly beautiful phenomenon, but also provide an important neural systems-level access point into understanding higher-order cognitive and mnemonic processes such as memory encoding, consolidation and planning. Experimental studies have shown impairments in learning and performance of various memory tasks when SWRs are disrupted (Girardeau et al. 2009; Ego-Stengel and Wilson 2010; Jadhav et al. 2012; Michon and Kloosterman 2019). In parallel, computational models illustrate how SWRs may contribute to the learning and performance of such tasks (Sutton 1990; Johnson and Redish 2005; Mattar and Daw 2018) motivating further work that increasingly relies on identifying SWR *content* under various conditions. Here, by SWR *content* we mean not simply aggregate properties such as the number or duration of SWRs, but “what is being replayed”: the structured spiking patterns during SWRs, such as the activation of specific ensembles and temporal orderings. These patterns have been associated with particular events, trajectories, and experiences, motivating a growing body of work that seeks to identify and decode SWR content.

Early studies of SWR content were mostly concerned with demonstrating the statistical robustness of particular kinds of non-randomness in SWR activity (see Tingley and Peyrache, this issue, for an excellent review). Although this seemingly straightforward issue is far from trivial (as we will discuss below; see also Foster, 2017), it is now established beyond doubt that SWR activity is structured in ways that deviate robustly from chance. However, the status quo is largely facing different questions: to what extent is SWR activity shaped by experience? Does SWR activity transform or prioritize particular experiences to serve specific cognitive or network benefits, e.g. by preferentially replaying salient experiences, trajectories less or more travelled, and/or conversely forgetting or suppressing other experiences? Does the presence of some specific set of factors dramatically change SWR content? These questions come with additional complexity in the data analysis, in that it is no longer sufficient to simply demonstrate that SWR activity is more structured relative to some notion of chance.

As experimental and theoretical interest in probing SWR content continues to evolve, we see a corresponding need for analysis methods that are appropriate for increasingly complex and subtle questions. The objective of this paper is to facilitate the analysis of SWR content by organizing questions about SWR content into a rudimentary taxonomy. A key feature of this taxonomy is the distinction between *first-order* questions, which rely on determining whether and how SWR content is different from chance, and *second-order* questions, which seek to establish whether SWR content differs between two or more conditions, such as time points or arms of a maze. Then, we will provide an inventory of challenges in data analysis specific to second-order questions, and point to some possible ways to diagnose, prevent, and address them.

First-order vs. second-order replay analysis

Figure 1a diagrams the distinction between first-order and second-order SWR content questions. First-order questions (cyan rectangles) are concerned with a single set of SWR events and a single *template* (depending on the question, a template may be a list of cells or ensembles, a set of tuning curves, or a specific ordering of cells); no comparison between multiple templates or between different sets of SWRs is required. First-order questions have sought to establish, for instance, that there is sequential order in SWR content similar to an environment that has yet to be explored (“preplay”, Dragoi and Tonegawa 2011, 2013; Farooq et al. 2019); that sequential activity can be in forward and reverse order compared to the place cell order experienced during behavior (Foster and Wilson 2006, Diba and Buzsaki 2007); that SWR content can be of a remote environment (i.e. one that the animal is not currently in, Karlsson and Frank 2009). These first-order content and order detection questions remain important, requiring careful consideration of assumptions built into the null hypotheses, particularly the shuffle/resampling methods, to determine whether the observed sequence activity is unexpected by chance. Such first-order issues have been recently discussed in several papers (Silva et al. 2015; Foster 2017; Farooq et al. 2019) and we will review them briefly in the next section.




Figure 1: Schematic illustration of two conceptual distinctions in replay analysis: first-order vs. second-order questions about replay content (A) and ensemble reactivation vs. sequence replay (B). **A:** *First-order* questions (cyan boxes) test whether a given activity pattern (e.g. which cells, what order) is distinct from chance. First-order questions are concerned with a single set of SWR events and a single template. In contrast, *second-order* questions compare replay content between multiple data sets (e.g. time 1 and time 2, purple boxes), templates (e.g. track A and track B, red boxes) or both. Typical second-order questions include determining whether post-experience SWR content better resembles activity during behavior compared to pre-experience content, and whether the left arm of a maze is more frequently (re)activated than the right arm. **B:** Both first-order and second-order questions about SWR content can be categorized as focusing on *which cells* participate, while ignoring temporal order (*ensemble reactivation*, left), and/or as focusing on *what order* these cells are active (*sequence replay*, right). Typical ensemble reactivation questions include determining if a given pair or ensemble of cells is more co-

active than expected by chance (a first-order question) or which of multiple possible ensembles is more active (e.g. “blue” vs “green” cells, a second-order question); typical sequence replay questions include determining if firing order is different from chance (a first-order question), and determining which of two firing orders is more prevalent (a second-order question).

Second-order questions include comparisons between different experimental conditions (red rectangles in Figure 1a) and/or time points (purple rectangles). For example, when viewed this way, establishing whether SWR content reflects *replay* of prior experience is a second-order question that requires comparing SWR content at two different time points: prior to, and following experience (Wilson and McNaughton 1994). Other typical second-order questions include asking whether the left or a right arm of a T-maze is preferentially expressed during SWRs (Gupta et al. 2010; Singer and Frank 2013; Carey et al. 2019), whether there is an increase in replay in the presence or absence of reward (Singer and Frank 2010; Ambrose et al. 2016); whether pharmacological and/or genetic manipulations affect SWR content (Suh et al. 2013; Middleton et al. 2018), and so on. These second order questions face additional scrutiny that requires accounting for an additional set of potential confounds. A major goal of this paper is to identify challenges in detecting and interpreting second-order patterns in replay and to suggest some practical advice.

First-order replay analysis

The main question at stake in a first-order SWR analysis is determination of whether activity in a given SWR (or set of SWRs, taken together) is different from a chance distribution. Activity may differ from chance in a number of ways, such as in *which cells* are active (for instance, particular pairs, ensembles, etc. may be more likely to be co-active than by chance), and in whether the *temporal ordering* of some set of cells is different from chance (Figure 1b). The choice of which chance distribution(s) to compare the observed data to (typically “shuffles” in a resampling procedure) directly determines the interpretation and level of specificity of the conclusion that can be drawn from a first-order analysis. In the literature, the term “replay” does not have a consistent technical definition. Some studies consider the (re)activation of specific cells, pairs or ensembles without any specific temporal ordering as replay, whereas others use the term *reactivation* to distinguish it from temporally ordered (sequential) activity. In this review, we try to be as explicit as possible in distinguishing *ensemble reactivation*, which can occur without any temporal order but does not preclude it, and *sequence replay*, which has a temporal order, but may additionally involve activity in a specific ensemble. In places, however, we will use the term “replay” nonspecifically to include either or both of these types of SWR structure.

This distinction between ensemble reactivation (“which cells”) and sequence replay (“temporal order”) illustrated in Figure 1b is important from several interrelated but distinct perspectives: neural, psychological, and analytical. Neurally, different kinds of structure in SWR activity will likely be read out differently depending on the downstream circuitry. Several studies have shown relationships between SWR content in the hippocampus and spiking activity in putative populations of readout neurons (ventral striatum: Lansink et al. 2009, Sosa et al. 2019; prefrontal

cortex: Jadhav et al. 2016; Shin et al. 2019; Todorova and Zugaro 2019; entorhinal cortex: Olafsdottir et al. 2016, but see Trimper et al. 2017; lateral septum: Tingley and Buzsaki 2019; auditory cortex: Rothschild et al. 2016; ventral tegmental area: Gomperts et al. 2015). However, these correspondences generally take the form of a general statistical model that does not reveal what specific features of hippocampal SWR activity are most important for downstream neurons. For instance, it is currently not known how accurately postsynaptic neurons distinguish between forward and reverse SWR sequences (but see Gutig and Sompolinsky 2006 for a proposed mechanism for how single neurons may do so). Alternatively, some structures, such as the lateral septum (Tingley and Buzsaki 2019), may only care about the strength and size of the SWR-associated population activity. In general, the problem of determining what features of SWRs matter is an important overall challenge in the replay field, and will need to be confronted with complementary lines of inquiry (see Box 1 for some promising directions). When considering different analysis issues and methods, such as those discussed in this review, it is important to keep in mind that the choice of analysis method should ultimately be informed by the features of SWR activity that are neurophysiologically or behaviorally meaningful.

Box 1: Ground Truth for replay?

Spiking activity during sharp-wave ripples (SWRs) exhibits a diverse array of structured patterns, including the systematic co-activation of particular subsets of cells (pairs and ensembles, “ensemble reactivation”), and temporal ordering that matches that observed during behavior (“sequence replay”). Analysis of SWR activity aim to determine whether such patterns are different from what could be expected by chance; however, they cannot, by themselves, decide which patterns are physiologically or behaviorally relevant. “Grounding” SWR content is an important but difficult problem that may be pursued by, among others, the following approaches:

- *Approach 1: determine what SWR content features are read out by downstream single neurons and brain structures.* For extra-hippocampal neurons that show a statistically reliable change in activity following SWRs, a number of analysis strategies may be used to determine that neuron’s preferred SWR pattern, or more generally, its tuning to SWR activity. Ideally, candidate readout neurons would be positively identified as receiving hippocampal input (e.g. with opto-tagging), and show temporal relationships with SWRs that reflect a genuine readout rather than merely correlation. Since SWR activity is potentially dynamic and high-dimensional, a range of appropriate dimensionality reduction techniques may need to be employed to characterize readout. Different neurons and brain structures may be “tuned” to different SWR features, applying different projections or decision boundaries to the input. For instance, lateral septum neurons may respond preferentially to an “overall activity” dimension (Tingley & Buzsaki 2019), whereas ventral striatal neurons may respond preferentially to activity associated with reward (Lansink et al. 2008).
- *Approach 2: trigger disruption or stimulation on specific SWR content features and observe the neural/behavioral consequences.* Real-time detection of specific SWR features can be used for a variety of interventions that causally test the importance of such features. For instance, selectively disrupting forward or reverse replay may reveal distinct behavioral impairments associated with each. Recording in putative readout areas while artificially triggering SWRs, and inhibiting synaptic terminals in triggered on SWRs occurring, can provide complementary evidence that a given neuron or area in fact performs a readout.

A second reason to care about different kinds of structure in SWR activity is that such differences may be associated with distinct cognitive and/or mnemonic processes. For example, temporally structured sequence replay may be better suited to signaling multimodal, temporally organized episodic events for which encoding and initial retrieval are supported by the hippocampus (we note, however, that it is currently unclear whether the content of replay is in fact episodic; the relationship between SWRs and the subjective experience of retrieval or mental time travel remains to be established.) Conversely, ensemble reactivation that is not explicitly ordered in time, but does involve a consistently co-active set of cells may be appropriate for retrieval of specific cues, spatial contexts, and semantic information that lacks a clear temporal component. Sequence replay and ensemble reactivation are in principle not mutually exclusive, either: for instance, the phenomenon of remapping suggests that distinct contexts are associated with specific ensembles (“maps”), in which multiple distinct temporally organized experiences can occur (Nadel 2008; Colgin et al. 2008; Kubie and Fenton 2019). Thus, a given SWR may involve both the reactivation of a specific ensemble (indicating a specific context) and a specific ordering of cells (indicating a specific trajectory or experience). Plausibly, the neural mechanisms underlying both these components are distinct to some degree, and certain experimental manipulations may have effects on only one, but not the other.

Finally, in terms of analysis, the distinction between the reactivation of cells and their temporal order is important because the appropriate analysis method depends on the phenomenon of interest. Cell reactivation (e.g. Pavlides and Winson 1989; Singer & Frank 2009), pairwise co-activation methods (e.g. Wilson and McNaughton 1994; Kudrimoti et al. 1999; Cheng & Frank, 2008; Singer et al. 2013; McNamara et al. 2014) and PCA/ICA-based ensemble reactivation studies (e.g. Peyrache et al. 2009; Benchenane et al. 2010; van den Ven et al. 2016) are agnostic towards temporal order. These methods also typically apply a fixed time window to measure co-activity in cells and ensembles, neglecting temporal compression across different brain states. While asymmetries in pairwise cross-correlograms can capture temporal relationships (Skaggs & McNaughton 1996; Euston et al.; Giri et al. 2019), as the template used to evaluate the order of firing is expanded to three (Nadasdy et al. 1999) or more neurons (Lee and Wilson 2002; Louie and Wilson 2001; Foster and Wilson 2006), the methods become increasingly sensitive to temporal structure. All of these approaches require careful consideration of the null hypothesis (shuffles) to compare the results to, and may give false positives based on factors such as firing rate differences (discussed in more detail below); but if the analysis is carried out appropriately then results can be interpreted cleanly in terms of either “which cells” or “what order” cause the non-random structure detected in the data.

How do current analysis methods detect reactivation and/or sequential replay? In this issue, Tingley & Peyrache (2020) review the various methods used to detect replay in ensembles and sequences and the statistics associated with each. In the present review, we focus our attention particularly on sequence replay methods that capture the order of firing across populations of neurons.

There are several popular metrics for quantifying sequence replay: the rank order (Spearman) correlation (Foster and Wilson 2006; Diba and Buzsaki 2007; Wang et al. 2015), the replay score

(“radon transform”) introduced by Davidson et al. (2009), and the linear weighted correlation (Wu and Foster 2014; Grosmark et al. 2016; Farooq et al. 2019) -- though other methods continue to undergo development (Maboudi et al. 2018; Makevicius et al. 2018; Rubin et al. 2019). These metrics ostensibly capture sequential activity, but can also be affected by factors unrelated to temporal order. This is because for any of these metrics, the key question becomes: what is the appropriate null distribution against which we should compare the data? Each shuffling method is based on assumptions about what random data ought to look like. If the data deviates from these assumptions then the null hypothesis can be rejected. However, rejection of the null hypothesis may be due to aspects of these assumptions that are not directly related to the sequential firing of cells during replay. We will next highlight several instances of this general issue.

For example, the *spike-id shuffle* randomizes the cell identity for each spike in a candidate SWR event (e.g. Wang et al. 2014), creating a surrogate dataset in which the overall firing patterns during the event are preserved, but each neuron is equally likely to fire any given spike from the onset to offset of this event. Therefore this null hypothesis can potentially be rejected if the distribution of inter-spike intervals of neurons deviates from uniformity, as is known to be the case (e.g. Ranck, 1973).

The *cell-id shuffle* (randomizing the cell identity for each train of spikes in a candidate event; used in Diba and Buzsaki 2007; Foster and Wilson 2006) further controls for the statistics of neurons’ observed firing patterns (e.g. burstiness, inhibition of return), but otherwise assumes that neurons fire independently of each other during ripple events. By randomizing the order of neurons, which also randomizes which neurons are co-firing in a given time window, this null hypothesis may be rejected by non-uniform co-activity, as well as non-random sequential activity. Thus, it tests against ensemble reactivation and sequence replay simultaneously, which may or may not be desirable. Moreover, it does not readily lend itself to Bayesian decoding analysis because, similar to the spike-id shuffle, it does not preserve non-uniformity in the firing rates of neurons (e.g. Mizuseki & Buzsaki, 2013).

The *circular place-field shuffle* (circularly shifting each tuning curve by a different random amount; used by Grosmark et al. 2016) also allows that neurons’ spike-train statistics should be preserved but assumes that their preferred firing locations are randomly dispersed. Importantly, this shuffle does not assume continuity in place-fields; surrogate place-fields can represent locations that start at the end of a track and reprise on the other side of the track. Therefore, this null hypothesis can be rejected if there is an uneven distribution of place-fields across the track, or the shape of place-fields observes specific relationships to the maze, as has been reported in several studies (Mehta et al. 1997; Hollup et al. 2001; Dupret et al. 2010).

Other surrogate methods first separate candidate events into distinct time bins, prior to shuffling. After binning, the *circular column cycle shuffle* (used by Davidson et al. 2009; Silva et al. 2015) shifts the decoded (posterior) positions in candidate bins by random amounts, thus controlling for the limited variance of decoded position in each bin, but assumes that all maze positions are equally likely to be reprised in any given time bin. Similar to the place-field shuffle, it allows that

decoded locations can be discontinuous across the end/beginning of tracks in the null distribution, which does not happen in real data. Therefore, this null hypothesis can be rejected if the decoded positions in candidate events are not uniformly distributed along the track and/or have specific skewed relationships to the end/start of tracks.

The *pooled time-swap shuffle* (also introduced by Davidson et al. 2009) randomizes the decoded position across all the candidate replay events in the dataset. Thus, it assumes that the decoded position in a given bin is random, but maintains the overall distribution of decoded positions across all observed bins. The *within-event time-swap* (used by Farooq et al. 2019) shuffles the decoded position bins in each candidate event separately (rather than across events), thus more conservatively preserving the overall statistics of decoded positions in each event. Yet, both of these time-swap shuffles assume that any decoded position bin is equally likely to follow any other decoded position in the null case—an assumption that is violated when neurons fire continuously in bursts that span multiple time bins (Ranck, 1973).

For each of these shuffles, there are further experimenter decisions and criteria that can affect the data and null distributions. These include, but are not limited to, the percentage and/or number of active cells, duration of events, the choice of time bin, the average or maximum jumps between bins allowed, and how empty bins are handled (e.g. is the posterior probability uniform or zero across all positions in such bins and are they ignored in shuffles if they fall in the first or last bins), a point which is often not explicitly noted in methods sections. Another critical point to keep in mind is that if multiple templates are being independently evaluated (e.g. for forward vs. reverse, and for each of several maze arms), there is a greater chance of false positives due to multiple-hypothesis testing (see also Silva et al. 2015). It is therefore important that the surrogate data following shuffles is treated similarly to real data and tested against each template to determine null hypothesis replay scores. Additionally, further properties of the data, aside from the factors already mentioned, can yield deviations from null distribution due to unintended reasons. For example, serial-position effects, such as biases for SWRs to be initiated by place-cells encoding the current and/or rewarded locations (as opposed to random locations along the track), could produce deviations from shuffle distributions that randomize decoded locations. Likewise, stationary events (SWRs with content that is fixed in one location), would qualify as “replay” under some metrics (e.g. “replay score”) but not others (e.g. “weighted correlation”). One reasonable solution to manage these issues is to consider combinations of measures (e.g. weighted correlation, and mean jump distance) at different thresholds (Foster 2017; Silva et al. 2015).

In summary, the main message from the above considerations is that many replay analysis methods can be sensitive to both which cells are active, and to their temporal order. Depending on the choice of shuffle(s), data may deviate from the null distribution due to factors that are not related to the sequential firing patterns during SWRs. In these cases, it is important to avoid interpreting the results as necessarily being due to one particular factor, such as sequential order. In the next section, we consider “second-order” questions that compare replay content across different conditions, such as “is there more replay in sleep after experience than in sleep before?” These questions inherit all the above issues related to the choice of shuffle and associated interpretations, and additionally entail another layer of analysis issues, which we will discuss.

Second-order replay analysis

A typical second-order sequence replay question is of the form: is sequence A replayed more frequently and/or more strongly than sequence B? Second-order questions go beyond the existence proof required of first-order questions (does reverse replay/remote replay/etc. exist?) to comparing the relative prevalence or strength of sequential structure across different conditions. Thus, the question “does reverse replay exist?” is a first-order question, whereas comparing the relative strengths of forward and reverse replay (Diba and Buzsaki 2007; Ambrose et al. 2016) is a second-order question. Other common second-order questions include comparing replay of different segments of a maze environment, or of replay after versus before an experience.

Forward vs. reverse sequences. Such comparisons began almost immediately after the observation of awake replay (Foster and Wilson, 2006). Diba and Buzsaki (2007) compared the relative prevalence of forward (positively-correlated) and reverse (negatively-correlated) sequences both before and after runs across the track. Forward sequences were found to be of upcoming trajectories prior to a run, suggestive of planning, while reverse sequences were observed after consumption of reward at the end of the run, suggestive of reward processing. Extending this work, Ambrose et al. (2016) directly compared templates for outgoing and incoming trajectories under high-reward, low-reward, and no-reward conditions on a linear track. They found that increasing the amount of reward on every trial produced an increase in the relative amount of reverse vs. forward sequences, supporting the proposed role of reverse replay in reward processing (Foster and Wilson, 2006). Olafsdottir et al. (2017) examined the relative prevalence of different trajectories during short pauses in the task versus long pauses that likely reflect disengagement from the task, and found that forward sequences were more prevalent in short pauses when the animal was engaged in the task. On the other hand, Davidson et al (2009), found that when the animal stopped in the middle of a track and could resume running in either of two directions, there was no correlation between sequence replays and the trajectories taken by the animal. Similarly, Shin et al. (2019) reported a strong correspondence between forward sequences and upcoming trajectories, and reverse sequences and completed paths--except at choice points. In a remarkable study, Xu et al. (2019) demonstrated that while forward sequences corresponded to the upcoming trajectory of the animal from its current position, reverse sequences frequently originated at remote locations and propagate towards reward sites, highlighting the special role of reward processing in reverse sequences, and planning in forward sequences. Pfeiffer and Foster (2013) examined these questions in open field 2D trajectories, where place-fields are omnidirectional (so that forward vs. reverse sequences are not distinguishable) and found that SWR sequences more closely matched upcoming trajectories rather than replays of past trajectories, consistent with a proposed role for replay in planning (Diba and Buzsaki 2007). Interestingly, Stella et al. (2019) found that in the absence of a reliable goal, SWRs instead reflect random-walk trajectories through the 2D open field.

In summary, while many of these studies differed regarding the relative prevalence of forward vs. reverse sequences across all trajectory sequences, there appears to be an emerging consensus that forward sequences benefit immediate planning while reverse sequences are modulated by

recent reward. Nevertheless, it is important to keep in mind that these forward and reverse replays are not mutually exclusive (Davidson et al. 2009). Indeed, Wu and Foster (2014) found that forward, reverse, and mixed replays frequently stitched together to form spatially consistent trajectories, and the “shortcut” sequences in Gupta et al. (2010) may result from the juxtaposition of spatially contiguous forward and reverse events. Ultimately, both types of sequences likely aid the animal in evaluating the reward structure of an environment to benefit future spatial decisions.

The main question for these analyses involved comparing whether a temporal event was better matched to one template (e.g. run from left-to-right) or to a second template (run from right-to-left) in either forward or reverse. Though there are several different ways to perform such a comparison (Diba and Buzsaki 2007, Davidson et al. 2009, Wu and Foster 2014), evaluation in these studies has generally been straightforward because (1) the templates compared (forward and reverse) are balanced -- that is, subjects run through each template trajectory an equal number of times -- and (2) they are applied to the same SWR data. Yet, even this relatively simple case is not immune to possible biases due to different numbers or different properties of cells in each of the templates (discussed in the next section).

Sequences at time-1 vs time-2. A different second-order question is a comparison of sequences that occur at different time points. A salient example of this type of question is preplay vs. replay, which requires comparing sequences before experience with sequences following experience. Such a comparison between two time points circumvents potential confounds from shuffling methods alone, although the quantification will still depend on the choice of shuffle. As a result, all investigators agree that replay in post-task sleep is stronger than preplay in sleep before the task, although by exactly how much is still under investigation and debate (Silva et al. 2016; Farooq et al. 2019).

Further examples of this kind of question focus on the conditions under which replay at one time is enhanced relative to replay at other times. For example, Singer et al. (2009) observed increased SWR reactivation of place-cells following rewarded trials, compared to unrewarded trials, indicating an important role for reward processing during these events. Note, however, that because reverse and forward sequences were not detected or separated in this study, it is unknown if this role is exclusive to reverse replay (Foster and Wilson 2006; Ambrose et al. 2016). In an alternation task on a W-shaped maze, Singer et al. (2013) reported increased SWR reactivation of place-cells prior to correct trials vs. incorrect trials. This provided support for the role of replay in planning and effective execution of upcoming routes, though this study similarly did not examine forward and reverse sequences, to determine whether forward sequences alone could account for the planning component (Diba and Buzsaki, 2007). In a recent study, Ting et al (2019) found that reverse replay decreases in frequency with time on the maze, whereas forward replay increases in frequency during this same period, indicating that SWR content can indeed change dynamically with experience.

Time-1 vs time-2 comparisons are similar to forward vs. reverse questions because they use the same templates for analysis, but are distinct because the templates are applied to the analysis of

different data (e.g. pre vs. post-task). While this approach has clear benefits, it does also introduce some possible issues, which we discuss in the section below.

Replay of track-1 vs. track-2. The more complicated version of second-order sequence analysis involves comparing replay of different tracks or mazes, on which the animals potentially have different amounts of experience, different behavior, and different numbers of place fields. For example, Olafsdottir et al. (2015) observed greater sequence preplay of an arm of the T-maze that the animal was cued to enter, compared to the uncued arm, indicating anticipation of the future path during SWRs. Wu et al. (2017) examined different regions of a linear track after delivering electric shocks when the animal entered one of the segments. They found more replay of the shock zone that the animal actively avoided, in comparison to a control region at the opposite end of the track that the animal also avoided, but that represented less danger. Xu et al. (2019) compared replays of different arms of a multi-arm maze. Remarkably, the arm that was replayed varied according to the cognitive demands of the task. In reference memory at the choice point, where the animal needed to remember the rewarded arms, forward replays tended to predict the upcoming path. On the other hand, during working memory on the same maze, when the animal needed to remember previously visited arms, the previously-visited arm was replayed in reverse. The observations in these studies provide strong support for the flexibility of hippocampal sequences in supporting cognition and decision making (see also Mattar and Daw 2018).

However, other second-order examinations present findings that are more challenging to reconcile with this simple picture. In the first study of its kind to question the relationship between experience and replay, Gupta et al. (2010) found that when rats were rewarded on only the left arm of a continuous T-maze, they replayed the opposite (right) arm more often than if they alternated between left and right. This is in conflict with an experience-driven account of replay, which would predict that the more frequently chosen arm should be replayed more. In a free-choice variant of this task, where the animal could choose whether to run for a water reward if thirsty, or a food reward if hungry, Carey et al. (2019) remarkably found that replay was biased toward the arm less visited, even in rest before the actual task. In their recent study, Xu et al. (2019) saw that when the animal was at the reward site and simply had to return back to the choice point, replays tended propagate starting remotely from the choice point over to the animal's current location, and in reverse. These types of trajectory sequences do not fit readily within the planning/reward framework and present a challenge to simple models of the cognitive or computational benefits of sequence replay.

Overall, the above types of second-order comparisons are more complex than the previous forward/reverse comparisons because they require different templates for each of the arms. This introduces possible biases due to differential electrode sampling (different numbers of cells recorded, cells with different firing rates, etc.) and behavioral sampling of the environment (when the animal's behavior involves an experiential bias towards a subset of the environment, e.g. by spending more time, running through it faster, etc). We discuss such biases in the next section.

Framework for issues in second-order sequence analysis

A number of potential analysis and interpretation issues arise when addressing seemingly straightforward A vs. B or time-1 vs time-2 questions comparing replay under different conditions. In general, under these conditions the detection and scoring of replay is susceptible to biases that are unrelated to underlying sequential content, but may nevertheless result in differences between observed replay of A vs. B. These false positives can lead to incorrect interpretations of results. Fortunately, these biases can be prevented or minimized by experimental design and careful analysis. In addition, some diagnostics are available to determine if these biases are occurring so that the interpretation can be modified accordingly. We conceptualize these biases within an overall conceptual framework for replay generation (Figure 2):




Figure 2: Framework for identifying issues in second order sequence analysis. Underlying “true” sequences (top, shown here as linear changes in position over time, though non-linear changes can be similarly represented) occur on either of two maps (A, green shading; B, blue shading). These are unobservable “ground truth” sequences that the data analysis seeks to recover. True sequences are postulated to generate observable spiking activity (shown in the rasterplots) through *sequence tuning curves*, which relate an internal position signal to firing rate, and are also not directly observable. True sequence content may differ between time points ($t = 1$ shown in the left column vs. $t = 2$ shown on the right) and between maps (A vs. B). Such analyses are inherently restricted because they sample neural activity from a finite number of cells (*electrode sampling*, left), and sample *behavioral tuning curves* (firing rate as a function of position, used to determine the sequence template) based on finite behavioral data (*behavioral sampling*, illustrated by the noisy tuning curves on the right). Sequence tuning curves,

electrode sampling, and behavioral sampling may all differ between time points and/or maps, potentially resulting in biases in resulting sequence scores.

In this framework, the hippocampus maps any given environment with a set of **true behavioral tuning curves**, which describe the relationship between locations in that environment and the firing rates of neurons. These are not directly observable for two reasons. First, any given experiment records from only a limited number of neurons (**cell sampling**). Second, tuning curve estimates are inherently based on finite and variable **behavioral sampling** of locations in the environment, and are therefore susceptible to various covariates and confounds. For example, animals may spend more time at some locations rather than others. Their running speeds and acceleration can also vary, not only according to the spatial configuration in the maze environment, but because of behavioral factors such as familiarity, anxiety, or expected reward. Furthermore, there are nonstationarities in the true tuning curves (e.g. emerging, drifting, and expanding place fields), yet the experimenter must average over multiple trials to obtain tuning curve estimates. Thus the observed tuning curves are inherently imperfect approximations of true behavioral tuning curves.

Likewise, in our framework observed sequences of spikes (rasterplots in Figure 2) during candidate replays arise from **true sequences** generated by the brain from a set of **true sequence tuning curves**, which are also not directly observable. The true sequence tuning curves are presumably translated by the brain from the true behavioral tuning curves and are therefore related but not necessarily identical to them; for instance, instantaneous firing rates during SWRs can differ from those during behavior (Csicsvari et al. 1999). Moreover, the sequence tuning curves presumably represent “mental” time travel through space, as opposed to physical travel during experience. As such, they may substantially differ from those obtained from behavioral sampling. The spikes during SWRs are nevertheless analyzed by comparison with the **observed behavioral tuning curves**, resulting in sequence scores. We note that because of uncertainty regarding the transformation between behavioral and sequential tuning curves, initial studies often used non-linear metrics for quantifying replay (e.g. Spearman rank-order correlation coefficients), whereas studies using Bayesian decoding methods assume identical sequence and behavioral tuning curves with interchangeable firing rates (Davidson et al. 2009 and others). In our framework, we remain agnostic concerning the nature of this transform.

This conceptual framework allows us to highlight a number of possible biases that are particularly relevant when the analyses involve second-order questions about sequences, which we discuss in turn below:

- (1) Electrode sampling from limited numbers of place cells
- (2) Imperfect estimates of behavioral tuning curves, which can arise from (2a) non-uniform behavioral sampling of the environment, and (2b) drift in place cell tuning over time
- (3) Changes in the sequence tuning curves




Figure 3: Schematic illustration of bias in second order sequence analysis due to unequal electrode sampling. Following the general flow introduced in Figure 2, unobservable underlying sequences (top) generate sequences of spiking activity (rasterplots) across two different hippocampal maps (equivalently, different arms of a maze). Differences in the number of recorded neurons between the two maps (here: five cells for map A, green; three cells for map B, blue) will result in differences in the sequence scores for the two maps, even though there is no underlying difference in replay content.

Biases in cell sampling. In the scenario illustrated in Figure 3, limited sampling of neurons results in an unequal distribution of place cells recorded for two maps (A and B; these could be different arms of a maze or two different environments entirely; 5 cells for map A, 3 cells for map B in this example). It may be that the different cell counts arise from random chance (shot noise), or that units representing one map are more noisy and/or less isolated. As the number of recorded cells increases, the magnitude of this bias will decrease correspondingly. However, the non-linear nature of Bayesian decoding methods, where posterior probabilities from spiking neurons are multiplied together, means that a small difference in the number of recorded cells can lead to large differences in decoded output, even when there is no underlying difference in replay content.

Alternatively, it may be that different cell counts result not from finite sampling but from inherent differences in the number of cells representing maps A and B. Such differences could be important (i.e. larger rewards associated with maze A, resulting in increased density of place fields (e.g.

Dupret et al. 2010), or they may be trivial (e.g. maze B may be smaller in size). Increased replay as a consequence of such underlying differences in place field density may or may not be functionally important, but is conceptually distinct from increases in replay that result from other factors not associated with place field density, such as memory prioritization and recall bias.

Regardless of how a difference in number of recorded place fields comes about, in the null case when the true replay distribution has equal sequence frequency and strength for both A and B, the observed spike trains will appear unequal: the 5 ordered cells for A result in higher replay scores than the 3 ordered cells for B (right panel). Taken in isolation, this difference in replay scores may be (incorrectly) interpreted as replay content favoring A over B.

Biases in behavioral sampling. In addition to unequal sampling of neurons, differences in the animal's behavior under two conditions introduces concerns that require careful consideration. Even if the true neuronal representation is similar in two mazes, if the animal spends more time in A, the experimenter's reconstructed template for B will be more noisy simply because less data was available, which in turn can lead to corresponding bias in replay analysis. Other potential sources of unequal behavior between different conditions include differences in running speed or acceleration through maze locations and the degree of path stereotypy.




Figure 4: Schematic illustration of bias in second order sequence analysis due to unequal behavioral sampling. Following the general flow introduced in Figure 2, unobservable underlying sequences (top) generate sequences of spiking activity (rasterplots) across two different hippocampal maps (equivalently, different arms of a maze). Differences in the data used to estimate tuning curves for the two maps, such as the number of trials and/or the amount of trial-to-trial variability in behavior, can result in “clean” tuning curves for one map (green tuning curves) and “noisy” tuning curves for another map (blue tuning curves). Note that for the noisy blue map tuning curves the sequence of tuning curve peaks is now in a different order (red dots), resulting in differences in the sequence scores for the two maps, even though there is no underlying difference in replay content.

As an example of the above possibilities, under the null scenario depicted in Figure 4, A and B maps are represented by an equal number of cells, but the *quality* of their observed tuning curves is different: messy, multi-peaked fields for B, clean, unimodal fields for A. The true behavioral tuning curves are not systematically different for A and B (left panel); rather, the behavioral tuning curves for B can deviate from the true tuning curves due to factors such as differences in behavioral sampling (e.g. 50 trials for A, 5 trials for B), differences in the animal’s behavior (e.g. more consistent runs in A, compared to stoppages or changing running speeds in B, etc.), and

differences in nonspatial cues influencing place cell activity between the arms. In a similar way to Figure 3 above, this difference in the observed tuning curves between the A and B maps will lead to a difference in replay scores for A and B because the noisier “B” template is now out of order (note red dots). This difference may be incorrectly interpreted as a difference in replay content. Furthermore, for questions that require trial-unique analysis of replay content (e.g. in a reward devaluation or replanning scenarios where the key manipulation can only be done once) behavioral sampling bias remains a particular concern.

Non-stationary tuning curves. A different source of bias in replay content that arises from imperfect behavioral tuning curve estimates relates to temporal drift, i.e. the known tendency for at least some cells to change their place tuning over time (Mankin et al. 2012, Ziv et al. 2013, Rubin et al. 2015), particularly in new environments (Frank et al. 2004). This problem is further compounded by instabilities in unit recordings, during which electrode drift or cell attrition can affect the isolation distance and cluster quality of a unit over time. Imagine that behavioral tuning curves are estimated from an initial task epoch and first applied to sequences that occur shortly after (Figure 5, left; “time 1”). In this case, the behavioral tuning curves are a good match for the true tuning curves. However, if the same behavioral tuning curves are used for analyzing sequences that occur some time later (“time 2”) the true tuning curves may have shifted (Figure 5, right) causing a difference in observed replay score even in the absence of a difference in true underlying sequences.

If this scenario seems far-fetched, consider that van der Meer et al. (2017) showed clear differences in decoding accuracy when decoding trials that were 1 vs. 10 trials apart. Additionally, observed tuning curves are known to change over time as a result of several factors, such as experience-dependent place field expansion (Mehta et al. 1997), look-ahead at decision points (Johnson and Redish 2007), rapid switching between multiple maps (Fenton et al. 1998; Jackson et al. 2007; Dvorak et al. 2018) and the presence of different gamma rhythms (Bieri et al. 2014).

On the other hand, for instances in which “time 1” and “time 2” are closer together (e.g. rewarded versus unrewarded trials, or correct versus error trials), examining the reactivation or co-activation probabilities of place cells or ensembles (i.e. “what cells”) can potentially provide the desired information without the need to evaluate their temporal sequences per se (e.g. Singer and Frank 2009; Singer and Frank 2013).




Figure 5: Schematic illustration of bias in second order sequence analysis due to non-stationary tuning curves. Following the general schematic introduced in Figure 2, unobservable underlying sequences (top) generate sequences of spiking activity (rasterplots), illustrated here for a single map. Behavioral tuning curves are estimated from an initial task epoch and first applied to sequences that occur shortly after (“time 1”). In this case, the behavioral tuning curves are a good match for the true tuning curves. However, if the same behavioral tuning curves are used for analyzing sequences that occur some time later (“time 2”) the true tuning curves may have shifted (note the difference in the true tuning curves and the behaviorally observed tuning curves, highlighted in yellow). Thus, outdated tuning curves are used for replay scoring, leading to a reduction in observed replay score even in the absence of a difference in true underlying sequences.

Changes in sequence generation tuning curves. A final source of bias can arise from differences or changes over time in the true sequence tuning curves. In our generative framework, sequence tuning curves underlie the instantaneous firing rates that occur during SWRs. Figure 6 illustrates the possibility that for map A (green), the sequence generation tuning curves are more precise than for map B (blue), resulting in larger sequence scores for map A. This could occur, for instance, because of differences related to deep vs. superficial areas of the CA1 pyramidal cell layer (Mizuseki et al. 2011). However, it presents an even more significant concern when comparing replay across time points that might be accompanied by changes in the excitability of cells. For example, neuromodulatory changes between awake and sleep states, or across the sleep/circadian cycle, can potentially affect the likelihood of cells to fire during SWRs, and their

true sequence tuning curves will vary accordingly across time points. Furthermore, as noted above, the decoded posterior probabilities during SWRs are sensitive to the number of co-active cells. Thus changes in SWR excitability alone can also impact replay scores.




Figure 6: Schematic illustration of bias in second order sequence analysis due to non-stationary sequence tuning curves. Unobservable underlying sequences (top) generate sequences of spiking activity (rasterplots) across two different hippocampal maps (equivalently, different arms of a maze). The tuning curves for the generation of map A sequences (green), may be more precise than those for the generation of map B sequences (blue; note wider tuning curves whose peaks are now out of order, red dots). As a consequence, the generated sequences for map B are now out of order compared to the behaviorally observed tuning curves, resulting in lower sequence scores for map B.

Relevance of these biases for ensemble reactivation measures. In some instances, investigators may simply be interested in second-order questions involving cells or cell ensembles, rather than temporal sequences. While we have highlighted these issues and potential biases with regards to sequence analysis, it is important to recognize that they also

factor into reactivation analyses using pairwise or PCA/ICA ensemble analyses for second-order questions. For example, less sampled regions of a maze, or fewer sampled neurons inherently result in poorer estimates for the co-activation patterns of neurons. For ensemble analyses in particular, because these methods typically involve a predetermined time-bin, temporal compression between behavior and replay presents an additional significant issue. Since many ensembles are most activated during SWRs, the common approach is to actually use SWR activity during immobility and sleep, rather than activity during behavior alone, to define the ensembles, and then to examine their instatement post-hoc during the behavior and/or sleep periods. That neurons contribute both positive and negative weights, and that reactivation strength can also take positive and negative values introduce other points worthy of careful consideration. Finally, it should also be kept in mind that correlation methods measure coordinated changes in firing rates and are therefore affected not only by global activity but also by global silences, such as during DOWN states (Mochol et al. 2015), LOW states (Miyawaki et al. 2016) or infra-slow oscillations (Watson, 2018). If these silent or low-activity periods vary across the periods under comparison, they will inevitably produce further confounds.

Strategies for diagnosis and prevention of second-order sequence analysis issues

We believe that the above biases resulting from unequal sampling of place cells and unequal estimates of tuning curves are pervasive. It is rare that animal behavior is exactly equal across an environment. So, what is a replay researcher to do? We suggest two main categories: **bias minimization**, which can occur by experimental design and analysis, and **bias diagnostics**, to determine what biases exist so that the interpretations can be appropriately qualified. Strategies for bias minimization include:

- In experimental design, take steps to promote equal behavioral sampling of the environment wherever possible. This can include thinking not only about ensuring that the animal samples different locations as equally as possible, but also about when and how (under what behavioral conditions) the animal samples these locations.
- In analysis, subsample numbers of trials to equalize them for the different locations of interest. For instance, if there are 50 A trials and 10 B trials, subsample the A trials to make the A-B comparison more equitable. Below a certain number, however, decoding accuracy will likely drop; for instance, van der Meer et al. (2017) found that decoding accuracy dropped substantially when less than 5 trials were used to construct tuning curves for decoding.
- If uneven numbers of cells are recorded in the conditions to be considered, a similar-sized section of the environment (e.g. length of track) should be compared against each other. This approach will therefore avoid inequalities arising from cell sampling, while preserving more important inequalities that may be due to differences in the true behavioral tuning curves, such as a higher density of place-fields in one condition versus the other.
- Future replay scoring methods should be able to take into account not only the mean firing rate in tuning curves, but also the uncertainty in that measure. This would help mitigate

biases due to unequal sampling that underlies mean firing rate estimations for different locations. The commonly used Bayesian decoding framework for replay analysis (Zhang et al. 1998; Brown et al. 1998) uses a Poisson process limited to a mean firing rate (λ) for each location, discarding information about the amount of data this mean estimate was derived from (more uncertain for low sampling, less uncertain for high sampling). Ghanbari et al. (2019) show that incorporating uncertainty in this sense improves decoding accuracy of reaching movements based on motor cortex data; a similar approach will likely be fruitful for hippocampal data under limited sampling conditions. Unit isolation and stability on an event-by-event basis could likewise be incorporated into a more comprehensive probabilistic framework.

- Future replay scoring methods should also aim to avoid the conflation between how strong a replay is (which should be independent of the number of cells) and how much evidence/uncertainty there is about that estimate (which does depend on the number of cells). Current methods such as z-scoring relative to a shuffled distribution mix these two together.
- Alternate methods based on a probabilistic framework such as Hidden Markov Models (Maboudi et al. 2018; Chen et al. 2016) can be valuable for incorporating uncertainties regarding the factors mentioned above, since in these methods, observation states and transition matrices are constructed from the distributions observed in the data in a template free manner. These methods can also be helpful when constructing a template from place cells is not desirable, for example, during one-trial learning or non-spatial behaviors.

Diagnostics for detection of replay content bias include:

- Compute a cross-validated decoding error on behavioral data where a true correct answer, such as the subject's position in space, is available. Importantly, the decoding error on data not included in the training set (=the data from which the decoder is obtained) is sensitive to all biases discussed above (van der Meer et al. 2017). Thus, this error can be used to obtain a null hypothesis for replay content given that the true underlying replay distribution is uniform, i.e. it can reveal the replay content differences that would be expected from experimental biases unrelated to true replay content.
- Generative models of replay can generate spike sequences from a specific model of ground truth. Such synthetic data can then be used to quantify the expected bias due to the various factors discussed above (e.g. limited behavioral sampling of true underlying tuning curves; different numbers and/or firing rates of place cells). Of course, the relevance of the results from such a model is proportional to how accurately the synthetic data captures the properties of real data; ongoing development of such generative models of replay is likely to be a fruitful area of research for multiple reasons.

Conclusion

The temporally compressed sequential firing patterns of hippocampal neurons during sharp-wave ripples have rightfully sparked widespread fascination and captured the interest and attention of

researchers with a broad range of scientific interests. This is due to the potential role of hippocampal sequences not only in learning and memory, but also in complex decision making, planning and imagination. Furthermore, recent years have seen (1) the emergence of a productive interaction with reinforcement learning and artificial intelligence systems that incorporate replay (Mnih et al. 2015; Mattar and Daw 2018), (2) the development of behavioral paradigms and fMRI/MEG analysis methods to study not only reinstatement, but also sequential replay in humans (Kurth-Nelson et al. 2017; Liu et al. 2019), and (3) the development of real-time causal interventions targeting SWRs and SWR content (Girardeau et al. 2009; Ego-Stengel and Wilson 2010; Jadhav et al. 2012; Ciliberti et al. 2017; Fernandez-Ruiz et al. 2019).

These different lines of work drive current questions about SWR content in increasingly complex experimental designs, such as those involving larger and more richly structured spatial environments with multiple possible trajectories, representation of not only the self but other agents, and multi-faceted non-spatial behavior. The continuing evolution of wireless and recording technologies further enable the exploration of larger environments, and longer continuous recordings. As these developments push the boundaries of replay questions, potential issues, such as unequal behavioral sampling in large complex environments, and temporal drift in longer duration recordings, become increasing concerns. These are likely to be combined with closed-loop approaches that will modify SWRs and their underlying sequence tuning curves. Yet, as we have argued here, these concerns are already prevalent today, even in seemingly innocuous settings such as comparing two arms of a maze or two different time points. There is ample potential for biases in replay analysis that can lead to erroneous inferences. While the potential for confounds and biases are pronounced, by careful experimental design, and consideration of the assumptions underlying null hypotheses, these issues can be understood and reasonably managed.

References

- Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron*, 91(5), 1124–1136. doi:10.1016/j.neuron.2016.07.047
- Bendor, D., & Wilson, M. A. (2012). Biasing the content of hippocampal replay during sleep. *Nature neuroscience*, 15(10), 1439–44. doi:10.1038/nn.3203
- Buzsaki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10), 1073–188. doi:10.1002/hipo.22488
- Carey, A. A., Tanaka, Y., & van der Meer, M. A. A. (2019). Reward revaluation biases hippocampal replay content away from the preferred outcome. *Nature Neuroscience*, 22(9), 1450–1459. doi:10.1038/s41593-019-0464-6

- Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2), 147–53. doi:10.1038/nn.2732
- Cheng, S., & Frank, L. M. (2008). New experiences enhance coordinated neural activity in the hippocampus. *Neuron*, 57(2), 303–13. doi:10.1016/j.neuron.2007.11.035
- Ciliberti, D., & Kloosterman, F. (2017). Falcon: a highly flexible open-source software for closed-loop neuroscience. *Journal of Neural Engineering*, 14(4), 045004.
- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, 63(4), 497–507. doi:10.1016/j.neuron.2009.07.027
- Diba, K., & Buzsaki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10), 1241–2. doi:10.1038/nn1961
- Dragoi, G., & Tonegawa, S. (2010). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330), 397–401. Retrieved from <http://dx.doi.org/10.1038/nature09633>
- Dragoi, G., & Tonegawa, S. (2013). Distinct preplay of multiple novel spatial experiences in the rat. *Proceedings of the National Academy of Sciences of the United States of America*, 110(22), 9100–5. doi:10.1073/pnas.1306031110
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, 13(8), 995–1002. doi:10.1038/nn.2599
- Dvorak, D., Radwan, B., Sparks, F. T., Talbot, Z. N., & Fenton, A. A. (2018). Control of recollection by slow gamma dominating mid-frequency gamma in hippocampus CA1. *PLoS biology*, 16(1), e2003354.
- Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20(1), 1–10. doi:10.1002/hipo.20707
- Euston, D. R., Tatsuno, M., & McNaughton, B. L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science*, 318(5853), 1147–50. doi:10.1126/science.1148979
- Farooq, U., Sibille, J., Liu, K., & Dragoi, G. (2019). Strengthened Temporal Coordination within Pre-existing Sequential Cell Assemblies Supports Trajectory Replay. *Neuron*, 103(4), 719–733.e7. doi:10.1016/J.NEURON.2019.05.040

- Fernandez-Ruiz, A., Oliva, A., de Oliveira, E. F., Rocha-Almeida, F., Tingley, D., & Buzsaki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. *Science*, 364(6445), 1082–1086. doi:10.1126/SCIENCE.AAX0758
- Foster, D., & Wilson, M. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440, 680–683.
- Foster, D. J. (2017). Replay Comes of Age. *Annual Review of Neuroscience*, 40(1), 581–602. doi:10.1146/annurev-neuro-072116-031538
- Ghanbari, A., Lee, C. M., Read, H. L., & Stevenson, I. H. (2019). Modeling stimulus-dependent variability improves decoding of population neural responses. *bioRxiv*, 146415.
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsaki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10), 1222–3. doi:10.1038/nn.2384
- Gomperts, S. N., Kloosterman, F., & Wilson, M. A. (2015). VTA neurons coordinate with the hippocampal reactivation of spatial experience. *eLife*, 4, 321–352. doi:10.7554/eLife.05360
- Grosmark, A. D., & Buzsaki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, 351(6280), 1440–1443. doi:10.1126/science.aad1935. arXiv: arXiv:1011.1669v3
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, 65(5), 695–705. doi:10.1016/j.neuron.2010.01.034
- Gutig, R., & Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing–based decisions. *Nature Neuroscience*, 9(3), 420–428. doi:10.1038/nn1643
- Hollup, S. A., Molden, S., Donnett, J. G., Moser, M. B., & Moser, E. I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J Neurosci*, 21(5), 1635–1644.
- Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087), 1454–8. doi:10.1126/science.1217230
- Johnson, A., & Redish, A. D. (2005). Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. *Neural Networks*, 18(9), 1163–71. doi:10.1016/j.neunet.2005.08.009

- Joo, H. R., & Frank, L. M. (2018). The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nature Reviews Neuroscience*, 19(12), 744–757. doi:10.1038/s41583-018-0077-1
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), 913–8. doi:10.1038/nn.2344
- Kloosterman, F., Layton, S. P., Chen, Z., & Wilson, M. A. (2014). Bayesian decoding using unsorted spikes in the rat hippocampus. *Journal of neurophysiology*, 111(1), 217–27. doi:10.1152/jn.01046.2012
- Kubie, J. L., Levy, E. R. J., & Fenton, A. A. (2019). Is hippocampal remapping the physiological basis for context? *Hippocampus*.
- Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6), 1183–1194. doi:10.1016/S0896-6273(02)01096-6
- Leutgeb, S., Leutgeb, J., Barnes, C., & Moser, E. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, 309, 619–623. Retrieved from <http://www.sciencemag.org/content/309/5734/619.short>
- Louie, K., & Wilson, M. A. (2001). Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep. *Neuron*, 29(1), 145–156. doi:10.1016/S0896-6273(01)00186-6
- Maboudi, K., Ackermann, E., de Jong, L. W., Pfeiffer, B. E., Foster, D., Diba, K., & Kemere, C. (2018). Uncovering temporal structure in hippocampal output patterns. *eLife*, 7, e34467.
- Mackevicius, E. L., Bahle, A. H., Williams, A. H., Gu, S., Denisenko, N. I., Goldman, M. S., & Fee, M. S. (2019). Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *eLife*, 8, e38471.
- Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, 109(47), 19462–7. doi:10.1073/pnas.1214107109
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11), 1609–1617. doi:10.1038/s41593-018-0232-z
- McNamara, C. G., Tejero-Cantero, A., Trouche, S., Campo-Urriza, N., & Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature neuroscience*, 17(12), 1658.

- Mehta, M. R., Barnes, C. A., & McNaughton, B. L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences of the United States of America*, 94(16), 8918–21.
- Michon, F., Sun, J.-J., Kim, C. Y., Ciliberti, D., & Kloosterman, F. (2019). Post-learning Hippocampal Replay Selectively Reinforces Spatial Memory for Highly Rewarded Locations. *Current Biology*, 29(9), 1436–1444.
- Middleton, S. J., Kneller, E. M., Chen, S., Ogiwara, I., Montal, M., Yamakawa, K., & McHugh, T. J. (2018). Altered hippocampal replay is associated with memory impairment in mice heterozygous for the *Scn2a* gene. *Nature neuroscience*, 1.
- Mizuseki, K., & Buzsaki, G. (2013). Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell Reports*, 4(5), 1010–1021.
- Mizuseki, K., Diba, K., Pastalkova, E., & Buzsaki, G. (2011). Hippocampal CA1 pyramidal cells form functionally distinct sublayers. *Nature neuroscience*, 14(9), 1174–81. doi:10.1038/nn.2894
- 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(21), 9497–9507.
- Nadel, L. (2008). The hippocampus and context revisited. In *Hippocampal place fields* (pp. 3–15).
- O'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Csicsvari, J. (2010). Play it again: reactivation of waking experience and memory. *Trends in neurosciences*, 33(5), 220–9. doi:10.1016/j.tins.2010.01.006
- Olafsdottir, H. F., Barry, C., Saleem, A. B., Hassabis, D., & Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *eLife*, 4, e06063. doi:10.7554/eLife.06063
- Olafsdottir, H. F., Carpenter, F., & Barry, C. (2016). Coordinated grid and place cell replay during rest. *Nature Neuroscience*, 19(6), 792–794. doi:10.1038/nn.4291
- Olafsdottir, H. F., Carpenter, F., & Barry, C. (2017). Task Demands Predict a Dynamic Switch in the Content of Awake Hippocampal Replay. *bioRxiv*. Retrieved from <http://www.biorxiv.org/content/early/2017/08/04/172098>
- Pavlidis, 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(8), 2907–2918.
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., & Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature neuroscience*, 12(7), 919–26. doi:10.1038/nn.2337

- Rothschild, G., Eban, E., & Frank, L. M. (2017). A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nature neuroscience*, 20(2), 251.
- Rubin, A., Geva, N., Sheintuch, L., & Ziv, Y. (2015). Hippocampal ensemble dynamics timestamp events in long-term memory. *eLife*, 4, 723–727. doi:10.7554/eLife.12247
- Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., & Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. *Nature communications*, 10(1), 1–14.
- Shin, J. D., Tang, W., & Jadhav, S. P. (2019). Dynamics of Awake Hippocampal-Prefrontal Replay for Spatial Learning and Memory-Guided Decision Making. *Neuron*.
- Silva, D., Feng, T., & Foster, D. J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nature Neuroscience*, 18(12), 1772–1779. doi:10.1038/nn.4151
- Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron*, 77(6), 1163–73. doi:10.1016/j.neuron.2013.01.027
- Suh, J., Foster, D. J., Davoudi, H., Wilson, M. A., & Tonegawa, S. (2013). Impaired Hippocampal Ripple-Associated Replay in a Mouse Model of Schizophrenia. *Neuron*, 80(2), 484–493. doi:10.1016/j.neuron.2013.09.014
- Sutton, R. S. (1990). First results with Dyna, an integrated architecture for learning, planning and reacting. *Neural Networks for Control*, 179.
- Tingley, D., & Buzsaki, G. (2019). Routing of Hippocampal Ripples to Subcortical Structures via the Lateral Septum. *Neuron*.
- Todorova, R., & Zugaro, M. (2019). Isolated cortical computations during delta waves support memory consolidation. *Science*, 366(6463), 377–381.
- Trimper, J. B., Trettel, S. G., Hwaun, E., & Colgin, L. L. (2017). Methodological caveats in the detection of coordinated replay between place cells and grid cells. *Frontiers in systems neuroscience*, 11, 57.
- van der Meer, M. A. A., Carey, A. A., & Tanaka, Y. (2017). Optimizing for generalization in the decoding of internally generated activity in the hippocampus. *Hippocampus*, 27(5), 580–595. doi:10.1101/066670

Wang, Y., Romani, S., Lustig, B., Leonardo, A., & Pastalkova, E. (2014). Theta sequences are essential for internally generated hippocampal firing fields. *Nature Neuroscience*, 18(2), 282–288. doi:10.1038/nn.3904

Wilson, M. A. & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172), 676–9. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8036517>

Wu, C.-T., Haggerty, D., Kemere, C., & Ji, D. (2017). Hippocampal awake replay in fear memory retrieval. *Nature neuroscience*, 20(4), 571.

Xu, H., Baracskey, P., O'Neill, J., & Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze. *Neuron*, 101(1), 119–132.e4.

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., . . . Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, advance on. doi:10.1038/nn.3329

3. Future perspectives

- [Caleb] - Talk about the idea of self consistency as a way of distinguishing between noise and not-noise.

-

We have been discussing time-compressed sequences in the context of a particular conceptual model of how memories might be formed and recalled by the hippocampus. Specifically, our assumption has been that the template that is developed during behavior — during activity of place cells — is the primary definition for the model or cognitive map encoded by the hippocampus. Activity during sharp-wave ripple bursts are then evaluated based on their consistency with this map. Learning or memory related dynamics are defined relative to this map. Specifically, learning might be expressed as a change in the expression of place cell activity or dynamics in the structure of activity during sharp wave ripples *measured via decoding using the place cell map*. This consistency across hippocampal modes is not the only way that experience or learning might impact neural activity within the hippocampus. In particular, decoding-based techniques are incapable of evaluating changes in the model or cognitive map expressed by the hippocampus during sharp wave ripples without having measured place cell activity. Recently, alternative approaches based on latent variable models and unsupervised learning have demonstrated that it is possible to construct a model of hippocampal activity

Place cell map change

Ripple map change

Consistency between replay events

In the process of establishing the existence of replay, it became clear that some sharp wave ripples involved patterns that seemed very disordered relative to the structure of ensemble firing in the context(s) in which neural activity had been observed. As a consequence of the presence of noise events, subsequent papers which have carried out second-order analyses of the replay phenomenon have often limited their analysis to a subset of events which were deemed to be “true” replay events.

- Statistical modeling of replay content: given activity at time $t-1$, 2, etc.. how well can we predict activity at time t ? Can then ask how that model changes with experience, etc
- [Matt] Generative model of replay: useful to synthesize ideas in a way that produces measurable benchmarks (% prediction accuracy) and useful testbeds for identifying biases (how many cells do we need to detect replay consistently, etc.)
- Is there a way of doing the first-order replay analysis based on generative model comparison? For example, can we ask whether there is an optimal monotonic transformation of the place field map that explains a large fraction of the events during SWR better than chance firing rates?

- Tuning curve instability / multidimensional tuning (could be a box? Not sure how this fits in.)
- Closed-loop manipulations.

Key terms and definitions

SWR

MUA

Encoding

Decoding

Tuning curve

...

Outtakes (NOT to be included in submitted manuscript!)

For instance, the “jump-distance” method first identifies replay events if over some minimum interval of time, decoded location does not jump more than a certain distance. If this method is applied without further selection, such as needing to travel a certain minimum distance overall, replay detections could result from decoding that stays in one place (same ensemble being active over time; no temporal order required). Other events may be detected because they do have some temporal order, such that simply reporting a number of detected events will contain contributions from both underlying phenomena. In such cases,

Conceptually, such “mixed” methods may classify a given replay event or data set as significantly different from random, but leave the underlying reason open.

Replay definition:

In practice, the influence of experience on SWR content is often implicitly assumed, justifying the use of the term “replay” to mean “any form of SWR activity that bears a non-random relationship to experience”, and we will adopt this usage in the rest of the paper.

State of the art:

- Rodent replay studies are generally event-based, because replay is associated with punctate SWR events and we potentially care about single event content (don’t just want to average; note this is a contrast with typical human replay studies, which for the time being don’t take an event-based approach).
- Event-based analysis depends on a set of experimenter choices and parameters that are important and ongoing foci of methods development (see **Box 2**) but we don’t focus on them here

Box 2: Detecting replay events

Two overall approaches: either (1) detect candidate events based on LFP properties (+MUA) and then only analyze those, or (2) analyze entire epochs with e.g. moving window and optionally apply further selection based on LFP

Issues with approach (1):

How to deal with thresholds

Use one LFP or multiple

Spiking content doesn't always align with SWR envelope in LFP

Issues with approach (2):

Typically more parameter choices. e.g. jump distance

- Tingley & Peyrache review the many possible methods of detecting whether replay is different from chance, which regardless of the actual method used (e.g. template, decoding, pairwise) generally relies on comparison with resampled/shuffled data.
- Shuffles are a crucial, necessary starting point for many studies, but we want to go beyond saying simply “shuffle passed”. We discuss 2 instances: (1) distinguishing between which cells are active vs. in what order they are active, and (2) comparing replay across multiple places (times)

First order issues

First-order issues in replay include:

- Which cells are co-active vs. sequential order
- What are the right shuffles
- Non-binary sequence scoring/showing distributions
- Showing absence of replay


Categorical replay decisions

2c: changes in replay in tail vs. mode/mean of distribution

- Hypothetically, experimental manipulations such as reward & reward prediction errors, motivational shifts may affect the full distribution of replay events, or only the tails. (I don't fully understand this, but I think the converse question would be - if we assumed that some small subset of the maze was replayed more faithfully, how would we be able to detect that increased replay scores were limited to a specific subset of space?)

2d: showing absence of replay

- Drift in tuning curves could underlie failure to detect replay
- Mention in Box 1?
- (Easier to compare replay during A with B).



(Figure XXb). Detecting such structure requires a different analysis approach that is sensitive to temporal order [such as rank order correlation between time of activation and place field location, or linear regression fits of decoder output over time]. Of course, experimentally observed replay may be a combination of these possibilities, i.e. single replay events may involve both specific subsets of co-active neurons that are also ordered in time.


Ensemble reactivation vs Sequence replay

Neural activity during replay can be structured in multiple, distinct ways. A prominent distinction is between **which ensembles** are active on the one hand, and **whether their firing order** is preserved. Figure XX below illustrates these scenarios. In the left panel, there are two distinct groups of neurons that tend to be co-active: in any single SWR event, one or the other group, but not both, is active. In this case, observed replay activity clearly deviates from a null hypothesis based on single-cell SWR participation probability. Similarly, metrics such as correlating the activity of all cells in a given event with an activity template provides a measure of replay content that is not sensitive to temporal order. (as is common in fMRI multi-voxel pattern analyses but also used in other methods such as calcium imaging.) These “non-temporal” measures have in common that they treat the activity during a replay event as a single time window (bin) and thus are insensitive to temporal patterns within events.

In contrast, activity during replay is not contingent on whether the same cells reactivate across events, or if participation is random, but only requires that the firing pattern is structured in time

In considering scenarios that result in apparently different

An important, totally different way in which replay content differences between A and B may be observed is not because of the experimental biases above, but because the true underlying distribution of tuning curves is actually different (Figure XXd):



Such non-uniform tuning curve distributions commonly occur near reward sites, around doorways, and in other situations (CITE). The resulting difference in the density/number of tuning curves will bias replay detection, but the interpretation is now importantly different from the situation where replay content bias resulted from unequal experimental sampling of tuning curves. Specifically, if replay content is systematically biased (e.g. towards reward sites) in a way that is explained by an underlying difference in tuning curves, then replay cannot be said to have transformed ongoing experience (**behavior?**); its content is faithfully reflecting the already-biased encoding of that experience. In contrast, claiming that a replay content bias results from a transformation of experience (e.g. replaying rewarded experiences more often than non-rewarded experience) requires showing that this content bias cannot be explained by differences in tuning curves.