Hippocampal replay of experience at real-world speeds

Eric L. Denovellis[1, 2, 3], **Anna K. Gillespie**[2, 3], **Michael E. Coulter**[2, 3], **Marielena Sosa**[4], **Jason E. Chung**[5], **Uri T. Eden**[6], **Loren M. Frank**[1, 2, 3]

***For correspondence:** loren@phy.ucsf.edu (LMF)

[1]Howard Hughes Medical Institute, University of California, San Francisco, San Francisco, California; [2]Departments of Physiology and Psychiatry, University of California, San Francisco, San Francisco, California; [3]Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, California; [4]Department of Neurobiology, Stanford University School of Medicine, Stanford, CA; [5]Department of Neurological Surgery, University of California, San Francisco, California; [6]Department of Mathematics and Statistics, Boston University, Boston, Massachusetts

Abstract Representations related to past experiences play a critical role in memory and decision-making processes. The rat hippocampus expresses these types of representations during sharp-wave ripple (SWR) events, and previous work identified a minority of SWRs that contain "replay" of spatial trajectories at ~20x the movement speed of the animal. Efforts to understand replay typically make multiple assumptions about which events to examine and what sorts of representations constitute replay. We therefore lack a clear understanding of both the prevalence and the range of representational dynamics associated with replay. Here we develop a state space model that uses a combination of movement dynamics of different speeds to capture the spatial content and time evolution of replay during SWRs. Using this model, we find that the large majority of replay events contain spatially coherent, interpretable content. Furthermore, many events progress at real-world, rather than accelerated, movement speeds, consistent with actual experiences.

Introduction

The brain has the remarkable ability to store and retrieve representations of past events, enabling memories of the past to influence future behavior. These memory processes depend critically on the hippocampus, where rapid plasticity during an experience is thought to drive the initial encoding of events (***Eichenbaum and Cohen, 2004***). Subsequently, hippocampal "replay" of stored representations during both slow-wave sleep and periods of immobility during waking is thought to contribute to the longer term storage and updating of event memories in distributed hippocampal-neocortical circuits (***Frankland and Bontempi, 2005***; ***Joo and Frank, 2018***; ***Carr et al., 2011***).

The canonical example of "replay" is seen in the rodent hippocampus, where cells preferentially fire at specific locations in an environment, and thus ensembles of cells fire in sequence as the animal moves through a series of locations. When the animal is asleep or immobile, hippocampal cells can be reactivated during a "sharp-wave ripple" (SWR) event. A subset of SWRs contain sequential firing similar to that seen during a previous experience, and thus are thought to "replay" these previous experiences. Importantly, previous work has reported that these sequential firing events proceed at an average speed of ~10 meters per second, about 20x faster than the animal's usual movement speed (***Nádasdy et al., 1999***; ***Lee and Wilson, 2002***; ***Davidson et al., 2009***; ***Karlsson***

and Frank, 2009).

While the existence of these sequential events is well established, the current consensus is that only a minority (~5-45%) of hippocampal SWRs contain statistically identifiable, sequential replay (***Michon et al., 2019***; ***Kaefer et al., 2020***; ***Shin et al., 2019***; ***Davidson et al., 2009***; ***Tingley and Peyrache, 2020***). One might therefore conclude that all memories that are stored or updated by SWRs correspond to sequences of locations. Our subjective experience, however, suggests that we can retrieve memories of individual locations without having to mentally traverse a long distance. Such memories would seem to be useful, in that they could encode the stimuli and values associated with a given place irrespective of the path used to get there. If the rodent hippocampus is capable of storing those sorts of memories, then one might expect to see SWRs where the neural activity corresponds not to a rapid trajectory through space, but instead to either a stable pattern associated with a single location or perhaps a brief pattern more similar to one that occurs during a real experience.

Interestingly, there is evidence for SWRs where the spiking corresponds to a single location. Specifically, some SWRs contain spiking of neurons associated with single locations where animals are immobile (***Yu et al., 2017***), although another report suggested that events that represent a single location are only seen in young animals (***Stella et al., 2019***). Thus, while it would seem useful to replay memories associated with single locations or a small region of the environment, the prevalence of that type of replay in awake animals remains unclear.

Our uncertainty stems in part from the dominant approaches used to identify the content of replay events. These approaches typically involve multiple steps and assumptions about the nature of replay, which are most commonly characterized using the standard "Bayesian" decoder. First, an encoding model is constructed based on spiking activity during movement, most often using only spikes that have been clustered as single units (putative single neurons). Then, a subset of SWRs or events with high activity levels are selected based on a threshold for event size chosen by the experimenter (***Foster and Wilson, 2006***; ***Diba and Buzsáki, 2007***; ***Karlsson and Frank, 2009***; ***Stella et al., 2019***). A decoding algorithm is then applied to the spikes within these events, yielding a set of position probability distributions for each time bin. Current approaches use either overlapping or non-overlapping time bins whose size is also chosen by the experimenter. Finally, the most commonly used approaches for detecting sequential replay involve fitting a line to the resulting set of probability distributions, which relies on the assumption that the representations progress at a constant speed (***Foster and Wilson, 2006***; ***Diba and Buzsáki, 2007***; ***Karlsson and Frank, 2009***; ***Tingley and Buzsáki, 2020***; ***Bhattarai et al., 2020***; ***Shin et al., 2019***; ***Drieu et al., 2018***; ***Ólafsdóttir et al., 2017***; ***Tang et al., 2017***; ***Carr et al., 2012***; ***Davidson et al., 2009***; ***Karlsson and Frank, 2009***). A statistical test is then used to determine whether the line fit to the data is better than a line fit to shuffled versions of the data, where the shuffled version of the data represent an implicit definition of a "random" sequence.

While the standard approach identifies constant speed events, it does not consider events that are rejected by the statistical test. This is problematic because it has the potential to mischaracterize real events that do not move at constant speeds, such as those that change direction or are discontinuous (***Liu et al., 2018***). Furthermore, the use of large, fixed-size temporal bins acts as a boxcar smoother that limits the potential movement speeds of the representation. For example, with 20 ms time bins and 3 cm position bins, an event can only move in 1.5 m/s increments (1 or more bins in 20 ms) between time steps. The linear fit is also highly dependent on the estimation of event boundaries, such as the start and end times of the SWR, because the fit depends on all the data over the course of the event. Approaches that focus on the order of cell activity within each event (***Lee and Wilson, 2002***; ***Gupta et al., 2010***) allow for a relaxation of that linear assumption, but replace it with a loss of statistical power due to either ignoring all but the first spike from each cell or using an arbitrarily smoothed version of the spike train. These approaches also do not provide information about the dynamics of the underlying spatial representation. Moreover, these approaches often exclude events that have stationary representations of a single location (***Yu et al.,***

2017; ***Farooq and Dragoi, 2019***).

Recognizing the problems with the linear fit, several studies have moved away from the constant velocity assumption, using the most probable location at each time bin (***Kaefer et al., 2020***; ***Carey et al., 2019***; ***Grosmark and Buzsáki, 2016***; ***Wu and Foster, 2014***; ***Pfeiffer and Foster, 2013***). For example, using this approach, ***Pfeiffer and Foster*** (***2015***) found awake replay events containing alternation between representations of a single location and sequential spatial trajectories. On the other hand, ***Stella et al.*** (***2019***) reported that replays during sleep are spatially continuous and follow Brownian diffusion dynamics. However, both methods still used large time bins and neither took into account the uncertainty of the decoded estimates, making it hard to identify the source of the different conclusions.

An ideal approach to identifying and quantifying the dynamics of replay would circumvent these problems. It would use all of the available spiking data to yield the most accurate decoded positions. It would provide a moment-by-moment estimate of position and dynamics that would not be dependent on the estimation of SWR start and end times and could rapidly adjust to changes in dynamics. It would use very small temporal bins (1 or 2 ms) to allow for very rapid representational movement and would provide information about the certainty of the decoded estimates. It would be able to capture a range of movement dynamics from stationary or unchanging representations to trajectories that progress through space at constant or variable speeds to disorganized, spatially incoherent events. It would provide a robust statistical assessment of confidence for each dynamic. Finally, where assumptions are made, it would provide well-defined parameters whose values could be explored systematically to understand their influence on the results.

We therefore developed a state space model that achieves all of those goals. State space models are a well-understood, well-known statistical solution to the problems described above. By mathematically modeling the relationship between the data and latent dynamics, state space models make the assumptions of the model explicit and interpretable. Our model goes beyond previous approaches (***Maboudi et al., 2018***; ***Deng et al., 2016***) by characterizing the spatial representations during SWRs as a mixture of three underlying patterns of movement dynamics: stationary trajectories, continuous trajectories that can progress at many times the typical speed of the animal, and spatially fragmented trajectories. We show how this model can take advantage of clusterless decoding—which relates multiunit spike waveform features to position without spike sorting—giving us more information about the population spiking activity. We apply this model to spiking data during SWR events from 10 rats, enabling a direct comparison to previous work.

We find that the large majority of SWRs contain spatially coherent content; that is, trajectories that are spatially concentrated at each moment in time and have no large discontinuities in position. Surprisingly, while the expected high-speed, sequential replay events were identified, the most common category of events expressed representations that moved at slower speeds, more consistent with real-world experiences. These findings illustrate the power of state space models and provide a new understanding of the nature of hippocampal replay.

Results

Overview of the model

We begin with an application of the state space model to simulated data, both to validate the model and to provide intuition (***Figure 1***). We simulate 19 Poisson spiking cells with Gaussian place fields on a 180 cm virtual linear track. Each place field has a 36 cm variance and a 15 Hz peak firing rate, and fields are spaced every 10 cm along the virtual track. We then construct a 280 ms spiking sequence (***Figure 1***A) and apply our model to the sequence. For the first 60 ms of this sequence, a single place cell fires repeatedly, resulting in the extended representation of a single location. For the second 190 ms of the sequence, the cells fire in sequential spatial order, representing a fast moving trajectory across the virtual linear track. For the last 30 ms of the sequence, the cells fire in an incoherent spatial order. These three firing patterns represent three different types of

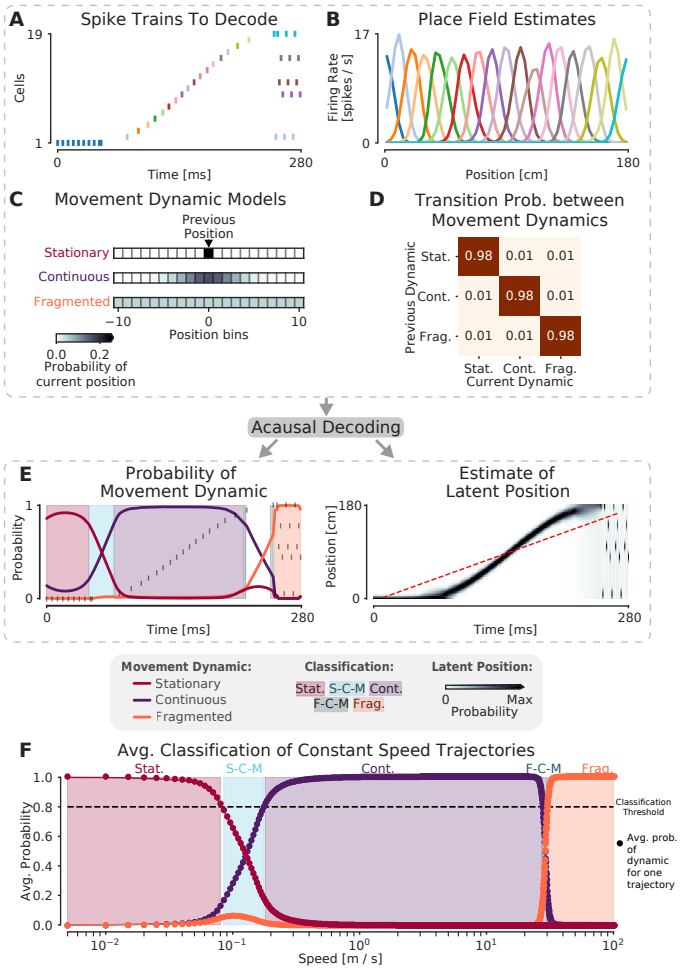

Figure 1. The model can capture different sequence dynamics on simulated data. **(A)** We construct a firing sequence of 19 simulated place cells that exhibits three different movement dynamics. For the first 60 ms, one cell fires repeatedly, representing one stationary location. For next 190 ms, the cells fire in sequence, representing a rapid continuous trajectory along the virtual track. For the last 30 ms, cells fire randomly, out of spatial order, representing an fragmented spatial sequence. **(B)** Like the standard decoder, the state space model uses estimates of cells' place fields from when the animal is moving and combines them with the observed spikes in (A) to compute the likelihood of position for each time step. **(C)** The prediction from the neural data is then combined with an explicit model of each movement dynamic, which determines how latent position can change based on the position in the previous time step. We show the probability of the next position bin for each movement dynamic model (color scale). Zero here represents the previous position. **(D)** The probability of remaining in a particular movement dynamic versus switching to another dynamic is modeled as having a high probability of remaining in a particular dynamic with a small probability of switching to one of the other dynamics at each time step. **(E)** The model uses the components in A-D over all time to decode the joint posterior probability of latent position and dynamic. This can be summarized by marginalizing over latent position (left panel) to get the probability of each dynamic over time. The shaded colors indicate the category of the speed of the trajectory at that time (Stat. = Stationary, S-C-M = Stationary-Continuous-Mixture, Cont. = Continuous, F-C-M = Fragmented-Continuous-Mixture, Frag. = Fragmented), which is determined from the probability. Marginalizing the posterior across dynamics also provides an estimate of latent position over time (right panel). Red dotted line in the right panel is the best fit line from the standard decoder using the Radon transform. **(F)** The probability of each dynamic depends heavily on the speed of the trajectory, as we show using a range of simulated spiking sequences each moving at a constant speed. Each dot corresponds to the average probability of that dynamic for a given constant speed sequence. We use a 0.80 threshold (dotted line) to classify each sequence based on the dynamic or dynamics which contribute maximally to the posterior (shaded colors).

Figure 1–Figure supplement 1. The model is robust to change of the probability of persisting in the same dynamic for a wide range of plausible expected durations (25 ms to 150 ms)

movement dynamics that could be expressed during SWRs, which we call stationary, continuous, and fragmented, respectively. The goal of our model is to characterize SWRs in terms of a mixture of these three dynamics at every time point.

Decoding the spiking sequence requires specifying two elements: the data model—how the spikes relate to position—and the movement dynamic models—how the position can change over time in each movement dynamic. For the data model, our decoder is the same as the standard ("Bayesian") decoder (***Davidson et al., 2009***; ***Pfeiffer and Foster, 2015***; ***Stella et al., 2019***). We compute an estimate of how each cell's firing rate varies over position during movement (i.e. the place field, ***Figure 1***B). This is used during decoding to compute the Poisson likelihood of position over time given the spiking sequence of interest. In contrast to the standard decoding approaches, we can use small time bins (in our case 2 ms vs. 20 ms or more) because we are able to take advantage of the prior placed on the dynamics by the state space model. This allows us to detect changes on smaller time scales than would be possible with the standard decoder and incorporate information about times when there is no spiking (***Deng et al., 2015***). Further, because place estimates from spikes within a single bin are combined, our small time bins allow us to measure the spatial information conveyed by single spikes, rather than assuming that a downstream neuron would integrate spikes from multiple neurons on the timescale of 20 ms or more.

Next, we specify movement dynamic models that describe a variety of ways that the latent position—the "mental" position of the animal represented by the cells— could evolve over time. (***Figure 1***C). We do this by defining a state transition matrix that defines how the latent position can move from the previous time step (in our case, 2 ms). Previous findings suggest that replay may exhibit at least three distinct types of movement dynamics: stationary (***Yu et al., 2017***; ***Farooq and Dragoi, 2019***), continuous (***Davidson et al., 2009***), and fragmented, which could correspond to both extended spatially incoherent representations and representations that jump from one position to another in a short time period (***Pfeiffer and Foster, 2015***). We therefore define movement dynamic models to capture each of these possibilities.

In the stationary movement dynamic, the latent position does not change between time steps. The state transition matrix can thus be defined as an identity matrix, which predicts that the next position will be the same as the last position (at the resolution of the position bin). In the continuous movement dynamic, the latent position is most likely to be "spatially close" to the position in the previous time step, so we use a Gaussian random walk state transition matrix. This means that, for a given latent position, the probability of moving to another position is modeled by a Gaussian centered at that position and "spatially close" is defined by the variance of the Gaussian. In our case, since replay has been shown to move at speeds much faster than the animal's speed (***Davidson et al., 2009***; ***Pfeiffer and Foster, 2015***), we set the variance of the Gaussian to 6.0 cm. This ensures that with a 2 ms time step, the latent position is 95% likely to be within 4.90 cm of the previous latent position (or equivalently, this means that latent speeds from 0 to ~25 m/s are most likely). Last, in the fragmented movement dynamic, the latent position can move to any available position instantaneously. We model this using a uniform state transition matrix, which makes transitions to each position equally likely. Importantly, the fragmented dynamic provides a well-specified definition of what we mean by random, or unstructured activity.

Finally, we specify how likely each movement dynamic is to persist in time versus change to another dynamic via another state transition matrix between the movement dynamics (***Figure 1***D). In order to be conservative with respect to switching between dynamics, we assume that each movement dynamic is likely to dominate for ~100 ms on average, which is approximately the duration of a SWR event. There is, however, a small probability of switching to one of the other movement dynamics. Accordingly, we set the probability of staying in a dynamic to 0.98 for each 2 ms time step, which corresponds to an expected duration of 100 ms for staying in a particular dynamic (the Markov assumption of the model means that the probability of staying in a dynamic follows a geometric distribution). Importantly, the data drives the estimated dynamic, so even if the probability of staying in a particular movement dynamic is high, data providing clear evidence for

a change in dynamic can, within a time frame of ~10 ms, drive a change to a different estimated dynamic, as in our simulated example. In line with this, we show below that our results are relatively insensitive to the value of this parameter.

Once we have specified the data and the movement dynamic models, we have fully specified the state space model. We use acausal decoding, meaning that we use all information from both past and future spikes, to estimate the joint posterior probability of position and dynamic (see Methods). With this, we can summarize the resulting posterior probability with two quantities: the probability of each movement dynamic over time (by integrating out position; ***Figure 1***E, left panel) and the probability of latent position over time, irrespective of movement dynamic (by summing over the movement dynamics; ***Figure 1***E, right panel).

An examination of the two summaries shown in ***Figure 1***E reveals that that the model successfully captures the dynamics of the population spiking activity in ***Figure 1***A. The stable firing of the one active neuron indicates a stationary representation, and accordingly, the probability of the stationary movement dynamic is high at the beginning of the simulated replay. A change in the data then drives a rapid transition to the continuous movement dynamic, reflecting the trajectory-like spatially sequential spiking from the population of cells. Subsequently, as the population activity becomes spatially incoherent, the fragmented movement dynamic dominates for the last 30 ms of the simulated event.

This illustrates two key features of the model. First, as mentioned above, the fragmented dynamic gives us a way to directly identify times when the position representation is spatially incoherent, a higher resolution and less assumption-dependent alternative to the more commonly used non-parametric shuffle. Second, this approach allows the model to to capture a wide range of movement speeds for the latent position in contrast to the standard decoder line fit (red dashed line in ***Figure 1***E, using the Radon transform). The model is defined in terms of a mixture of movement dynamics, as summarized by the probability of each movement dynamic, and which dynamic or dynamics are dominant at a given moment is related to the temporal evolution of the underlying position representation. To demonstrate this, we applied the model to 10,000 simulated trajectories, each trajectory proceeding at a constant speed (***Figure 1***F) from 1 cm/s to 10,000 cm/s. From this, we can see that not only are there regions of speed that correspond to each of our three movement dynamics being highly probable (where we define highly probable to be greater than or equal to 0.80 probability), but there are also intermediate speeds where two of the dynamics exhibit relatively high probability; and where the sum of two of the dynamics' probabilities exceeds 0.80. In this manuscript, we will refer to these intermediate speeds as mixture dynamics. For example, when the stationary dynamic has a probability of 0.6 and the continuous has a probability of 0.4, we call this a stationary-continuous mixture (light blue, ***Figure 1***F) and this indicates that the trajectory is moving slowly. Correspondingly, if the continuous dynamic has a probability of 0.5 and the fragmented dynamic has a probability of 0.4, then we would call this a fragmented-continuous-mixture and this indicates the trajectory is moving very quickly, but not as quickly as the fragmented dynamic dictates. In summary, we can characterize the speed or set of speeds that occur within a SWR based on the probability of each of the three movement dynamics over time. We further classify the probability of each movement dynamic as being part of one of five speed categories: stationary, stationary-continuous-mixtures, continuous, fragmented-continuous mixtures, and fragmented.

We note here that the choice of any particular threshold for classifying the movement dynamic of a SWR is arbitrary, and that the power of our approach lies in part on the ability to assign a probability for each dynamic or combinations of dynamics to each moment in time. Our goal in choosing 0.80 was to use a threshold that corresponds to highly probable events and to explicitly avoid the arbitrary choice of 0.95, as the 0.95 threshold often leads to the conclusion that only events that are statistically significant at the 0.05 level are meaningful in the brain. Nonetheless, we also verify that our central results hold with a higher threshold of 0.95.

Finally, as mentioned above, we wanted to test the robustness of the model to the choice of probability of staying in a dynamic, because our choice of 0.98 or an expected duration of

100 ms is only based on the expected duration of a SWR. To investigate this we decoded the spiking sequence in ***Figure 1***A with different probabilities of staying in the same dynamic versus switching to another dynamic (***Figure 1–Figure Supplement 1***). We found that for a large range of plausible probabilities of staying in one of the dynamics (between 0.96 and 0.993, corresponding to an expected duration between 25 and 150 ms), the model identified the movement dynamics consistently, with high probability (***Figure 1–Figure Supplement 1***A), demonstrating that data itself is the most influential element in the model. Furthermore, the most probable latent positions remained relatively consistent across the range of these probabilities as well (***Figure 1–Figure Supplement 1***B).

Identification of continuous dynamics in a SWR with sorted or clusterless spikes

We next sought to validate the model on real hippocampal data recorded from awake, behaving rats. To ensure that we could capture rapid trajectories many times the speed of the animal—as described by most studies of hippocampal replay—we first decoded a single SWR with sequential population activity (***Figure 2***A, top panel). Here the rat was performing a spatial alternation task on a W-shaped track (see Methods for details). As expected, we observed that the probability of the continuous dynamic is high throughout this event, but the probability of the stationary dynamic was also noticeable at the beginning and end of the SWR (***Figure 2***A, middle panel). Using our speed classification scheme, this indicates that the speed of the replay is initially slower—a mixture of continuous and stationary dynamics—and then speeds up and slows down again. This is also evident in the posterior probability of latent linear position over time, which shows that the replay most likely travels down the center arm and up the right arm (***Figure 2***A, bottom panel). We can also see this when we project the maximum of the posterior of this trajectory (the most probable "mental" position) to 2D to better see the spatial trajectory on the maze (***Figure 2***C). Importantly, when we apply the same model using the 2D position of the animal we get a similar result (***Figure 2–Figure Supplement 1***A).

One of our criteria for a more optimal method is for it to use all of the available spiking data. Using only clustered spikes discards any spike events that cannot be uniquely assigned to a putative single neuron, substantially reducing the amount of information that the resultant decoding can use. Additionally, spike sorting is not necessary to recover the underlying neural population dynamics (***Trautmann et al., 2019***) and often requires inclusion decisions that can vary widely across experimenters. We therefore adopted a "clusterless" approach which directly uses multiunit spikes and their spike waveform features to decode position without spike sorting (see Methods). Clusterless decoding has previously been used to successfully identify theta sequences and replay sequences in the hippocampus (***Kloosterman et al., 2014***; ***Chen et al., 2012b***; ***Deng et al., 2016***; ***Kay et al., 2020***). Applying a clusterless decoder to the same SWR event, we get similar classification of the sequence (***Figure 2***B, D), both with 1D linearized position and 2D position (***Figure 2–Figure Supplement 1***B). As predicted, the spatial extent of the event is longer and the estimate of the posterior probability of latent position is narrower for the clusterless model. This reflects the clusterless model's access to a larger pool of data that provides more information about the extent of the event and more certainty in the latent position and the dynamic (***Figure 2***D vs C).

Most replays contain coherent spatial content

After testing our model on a single SWR event, we applied our clusterless decoding algorithm to hippocampal recordings from 10 rats performing the W-track spatial alternation task (# tetrodes range: [10, 24], brain areas = [CA1, CA2, CA3]; some data previously used in ***Karlsson and Frank (2009)***; ***Kay et al. (2020)***; ***Carr et al. (2012)***; position linearized to 1D). We first confirmed that our cell population for each animal could closely track the position of the animal during behavior by comparing the most probable decoded position to the position of the animal. Even including theta sequences and replay sequences, which can represent positions other than where the animal is, we found that the average median difference between actual and decoded location is small (4 cm

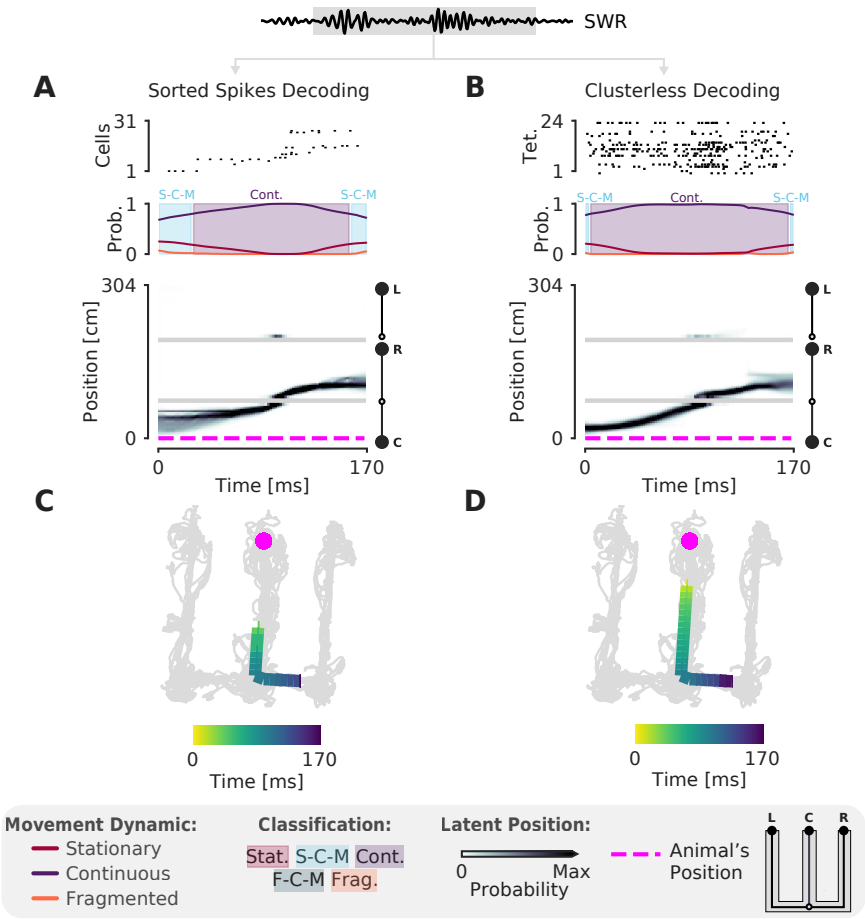

Figure 2. The model can decode hippocampal replay trajectories using either sorted and clusterless spikes from the same SWR event. **(A)** Decoding using sorted spikes. The top panel shows 31 cells on a W-track ordered according to linearized position by their place field peaks. The middle panel shows the probability of each dynamic over time as in ***Figure 1***E, left panel. Shaded regions correspond to the speed classifications as in ***Figure 1***F. The bottom panel shows the estimated probability of latent position over the course of the SWR as it travels down the center arm toward the right arm. L, R, C correspond to the position of the left, right and center reward wells respectively. The animal's actual position is indicated by the the magenta dashed line. Underneath is the maximum of the 1D decoded position (the most probable position) projected back onto the 2D track for the sorted decoding. Color indicates time. The animal's actual position is denoted by the pink dot. Light grey lines show the animal's 2D position over the entire recording session. **(B)** Decoding using clusterless spikes. The top panel shows multiunit spiking activity from each tetrode. Other panels have the same convention as (A). Underneath is the maximum of the 1D decoded position (the most probable position) projected back into 2D using the clusterless decoding.

Figure 2–Figure supplement 1. Decoding the same SWR in ***Figure 2*** with 2D position using sorted spikes and clusterless decoding.

Figure 2–Figure supplement 2. More examples of SWRs that have continuous trajectories.

Figure 2–Figure supplement 3. Population firing rate on the track is spatially uniform and consistent for each animal.

Figure 2–video 1. Example of a SWR with continuous content. Magenta dot represents the animal's position. Green dot represents the most likely decoded position projected from 1D back to the linearized 2D position. Green line represents the decoded positions in the last 10 ms. Grey lines represent the position of the animal over the course of the recording session.

median difference, 3-6 cm 95% CI, for all time in sessions, 5 fold cross validation). This level of accuracy is just above the size of our 3 cm position bins, and is comparable to that of other studies (***Davidson et al., 2009***; ***Shin et al., 2019***; ***Farooq and Dragoi, 2019***). Importantly, we achieved this using small time bins (2 ms) compared to the commonly used 250 ms time bin that most studies use for decoding during movement, demonstrating the power of our algorithm. For each animal, we also verified that the population spiking rate was relatively consistent over all positions on the track, suggesting that we have comparable sampling of all locations in the environment (***Figure 2–Figure Supplement 3***).

Having established the validity of our decoder in tracking the animal's position, we next assessed the prevalence of spatial content across SWRs. We detected SWRs using a permissive threshold (see Methods) to ensure that we include both the rare large-amplitude events as well as the much more common small-amplitude events. As expected, and providing further support for the validity of the model, we observed replays that are classified as continuous throughout the entirety of the SWR, similar to those seen using standard approaches (***Figure 2–Figure Supplement 2***A-F). However, we also observed many events with spatially coherent content that do not have this structure. For example, there are trajectories that start in one direction and reverse back to the original position (***Figure 3***A, ***Figure 3–Figure Supplement 1***B), representations that remain fixed in one position (***Figure 3***B, ***Figure 3–Figure Supplement 1***G), and trajectories that jump between arms and between dynamics (***Figure 3***C, D, F, ***Figure 3–Figure Supplement 1***F, H, I). We also observed SWRs where the content is spatially incoherent throughout (***Figure 3–Figure Supplement 1***A, D).

Using a 0.80 threshold, 89% (23,071 of 25,844) of SWRs contain at least one of the three dynamics or dynamic mixtures. To ensure that this reflects the spatially tuned firing of the neurons, we trained the encoding model with positions resampled with replacement, a shuffling procedure which disrupts the relationship between spiking and position, for two recording sessions. We then decoded the same SWR events, containing the original spikes. Only 9% of the SWRs are classified in the shuffled datasets, a value that is significantly less than that seen for the real data (p=0.02 for recording session 1, p=0.02 for recording session 2, ***Figure 3–Figure Supplement 2***). This shows that our model does not impose movement dynamics in the absence of spatial information in the data.

Previous work focusing on spatially sequential replay reported that only a minority of events contain sequential spatial content (***Karlsson and Frank, 2009***; ***Foster and Wilson, 2006***; ***Davidson et al., 2009***). We therefore asked what fraction of classified events contain spatially coherent content, which we define as any times with stationary, stationary-continuous mixture, or continuous dynamics (see Methods and ***Figure 1***F). We find that 96% (22,170 of 23,071) of classified SWRs and 86% (22,170 of 25,844) of all SWRs include spatially coherent structure, and that this prevalence of spatially coherent structure is consistent across animals (***Figure 3***G). We then asked what fraction of events contained spatially incoherent content, defined as a SWR containing any times with fragmented or fragmented-continuous mixture dynamics. We find that only 14% (3,295 of 23,071) of classified SWRs include any times with spatially incoherent structure (***Figure 3***G). To further validate this result, we performed a shuffling procedure that, in contrast to our previous shuffle, preserves the local correlations between spikes but reduces global spatial structure. This controls for dynamics that may have been induced by bursts or other local features of the spike train while randomizing the position to spike relationship. To do this, we randomized the order of runs (from one reward well to another) and then circularly permuted the resulting segments of data across all tetrodes uniformly. Using this shuffle, we found that there was still less spatially coherent and more spatially incoherent content compared to the real data (recording session #1, spatially coherent: 99% vs. 71%, p=0.02, real vs. shuffled; spatially incoherent: 22% vs. 61%, p=0.02, real vs. shuffled; recording session #2, spatially coherent: 95% vs. 63%, p=0.02, real vs. shuffled; spatially incoherent: 16% vs. 79%, p=0.02, real vs. shuffled; ***Figure 3–Figure Supplement 3***).

To more directly compare our findings to previous work, we quantified the percentage of classified SWRs that contained continuous content, as would typically be analyzed when using the standard decoder. Here our findings were consistent with previous reports: in our datasets,

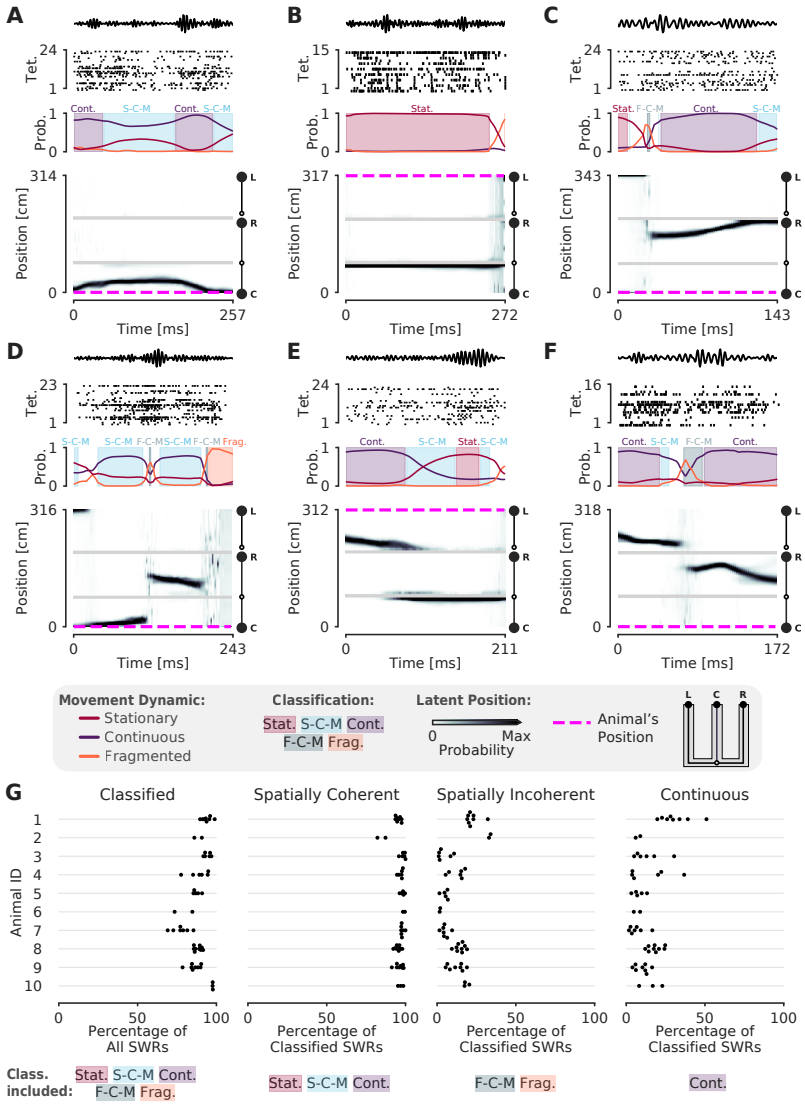

Figure 3. A-F. Examples of SWRs with non-constant speed trajectories. Figure conventions are the same as in *Figure 2*. Filtered SWR ripple (150-250 Hz) trace from the tetrode with the maximum amplitude displayed above each example. **(A)** A SWR where the decoded position starts moving down the center arm away from the animal's position at the center well, slows down, and returns back. **(B)** A SWR where the decoded position persistently stays at the choice point (open circle) while the animal remains at the left well. **(C)** A SWR where the decoded position begins with stationary representation of the left well, then jumps to the middle of the right arm and proceeds up the right arm to the right well. **(D)** A SWR where the decoded position begins with stationary representation of the left well, jumps to the center arm, proceeds away from the center well, jumps to the right arm, proceeds back toward the center well, and then becomes fragmented. **(E)** A SWR where the decoded position begins in the left arm and persists at the end of the center arm. **(F)** A SWR where the decoded position starts in the left arm toward the choice point, jumps to the right arm and proceeds back toward the choice point. **(G)** Classification of SWRs from multiple animals and datasets. Each dot represents the percentage of SWRs for each day for that animal. A SWR is included in the numerator of the percentage if at any time it includes the classifications listed below the column. The denominator is listed in the x-axis label.

Figure 3–Figure supplement 1. More examples of SWRs with non-constant speed trajectories

Figure 3–Figure supplement 2. Shuffling the position data with replacement decreases the percent of SWRs classified

Figure 3–Figure supplement 3. Shuffling the data by swapping the runs and circularly permuting the position increases the percentage of spatially incoherent SWRs and decreases the spatially coherent SWRs

Figure 3–video 1. Example of a SWR with that is not purely continuous. Conventions the same as Figure 2-video1

4,449 of 23,071 or 19% of classified SWRs included time periods where the decoded position was classified as continuous (***Figure 3***G, 17% of all SWRs). Thus, focusing on only high-speed, linear-fit trajectories excludes a large fraction of events where there is evidence for spatially coherent, but not continuously changing, content. We emphasize here that we did not limit our analysis to only those SWRs that had continuous dynamics for the entire SWR, as is assumed by line-fitting approach of the standard decoder. The consideration of this broader class of SWRs allows us to take advantage of the ability of our decoder to capture a range different speeds within each SWR event.

We repeated our classification analysis with a higher classification threshold of 0.95 to ensure that our result was not dependent on the threshold of 0.80. We find that, while this change slightly reduced the total fraction of classified SWRs (19,478 of 26,159 or 74% of all SWRs), an even higher fraction of the classified SWRs (19,317 of 19,478 or 99% classified SWRs) included spatially coherent content. Similarly, SWRs containing spatially incoherent content comprised a small fraction of the classified SWRs (490 of 19,478 or 3% classified SWRs).

Because our model is specified in the context of a latent position associated with different movement dynamics, it allows us to not only classify events in terms of their dynamics, but also to quantify the model's certainty in each position estimate at each moment in time given the model parameters. To do so, we can compute the cumulative spatial size of the 95% highest posterior density (HPD) region of the latent position estimate. The 95% HPD region corresponds to the set of most probable latent positions as estimated by the model. Larger values of the HPD region size indicate the model is less certain about position, because the most probable decoded positions are distributed over more of the track at a given time point. In contrast, smaller values of the HPD region size indicate that the model is more certain about the estimate of position because the extent of the HPD is more concentrated and covers less of the track. Thus, the HPD region size provides a complementary measure of spatial information, and evaluating it allows us to verify that the events we defined as spatially coherent also correspond to events where there is high certainty around the position estimates.

We find that spatially coherent events indeed have smaller HPD regions than events with fragmented dynamics. ***Figure 4***A and ***Figure 4***B show two example SWRs that are classified as having continuous and stationary dynamics, respectively. The most probable positions (as estimated by the model) at each time step in these SWRs is concentrated in a small portion of the track and correspondingly, the HPD region is small throughout the SWRs. In contrast, ***Figure 4***C shows a SWR where the dynamics are fragmented and correspondingly, the most probable positions are more spatially diffuse and the HPD region is much larger. The HPD region size also provides insights into the moment-by-moment structure of each event, which can change over the time course of a SWR. An example of this change is shown in ***Figure 4***D, where the HPD region is small for most of the SWR until the end, at which point the uncertainty of position becomes much higher, reflecting a transition from a spatially coherent to a spatially incoherent representation. Overall, when we examine the average HPD region size for each SWR, grouped by dynamic, we find a clear bimodal split between spatially coherent dynamics and spatially incoherent dynamics (***Figure 4***E). For the spatially coherent dynamics, the average HPD region for each SWR was much smaller than the spatially incoherent dynamics (median 24 cm, spatially coherent vs. median 238 cm, spatially incoherent, p=2.2e-16, one-sided Mann-Whitney-U). The HPD region for the unclassified SWRs was similarly large.

We note here that while the size of the HPD region will be influenced by the dynamic estimated from the data, it remains possible for continuous or stationary dynamics to correspond to a large HPD region (see outliers in ***Figure 4***E), indicating less spatial certainty for those specific events. In general, if the data does not exhibit strong place specific firing, the HPD region will be large, regardless of dynamic classification. To show this, we used the same shuffle as above, resampling position with replacement for two recording sessions and shuffling the relationship between spiking and spatial information when fitting the encoding model. The shuffled decodes have much larger HPD regions than the real data (recording session #1: 55 cm vs. 229 cm, p=0.02, real vs. shuffled;

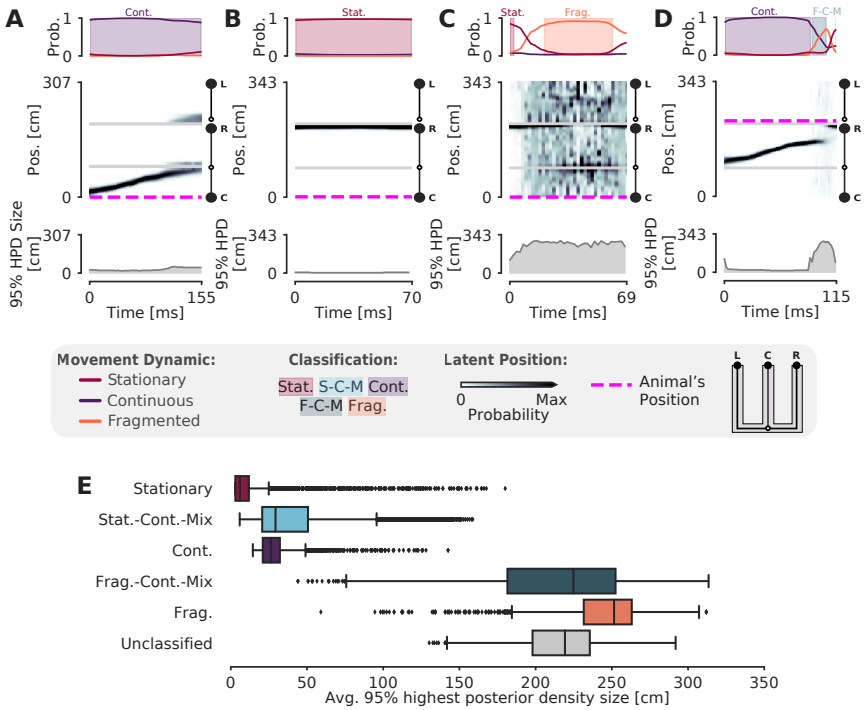

Figure 4. Validation of classification using the 95% Highest Posterior Density. **(A-D)** Examples of the 95 % Highest Posterior Density. In each column: top panel: Probability of dynamic over time. Shading and labels indicate dynamic categories. Middle panel: Posterior distribution of estimated linear position over time. Magenta dashed line indicates animal's position. Bottom panel: HPD region size—the integral of the position bins with the top 95% probable values. **(E)** Average 95% HPD region size for each dynamic category. Median 24 cm for spatially coherent vs. median 238 cm for spatially incoherent, p=2.2e-16, one-sided Mann-Whitney-U

Figure 4–Figure supplement 1. Shuffling the position with replacement increases the average 95% HPD region size

Figure 4–Figure supplement 2. Shuffling the position data by swapping the runs and circularly permuting the data increases the average 95% HPD region size for spatially coherent and incoherent classified times

recording session #2: 72 cm vs. 231 cm, p=0.02, real vs. shuffled, one-sided; ***Figure 4–Figure Supplement 1***). We also compared the HPD regions between the real data and a shuffle that swapped the runs and circularly shuffled position, as above. This also resulted in larger HPD regions than the real data for spatially coherent and incoherent SWRs (recording session #1, spatially coherent: 29 cm vs. 52 cm, p=0.02, real vs. shuffled; spatially incoherent: 181 cm vs. 243 cm, p=0.02, real vs. shuffled; recording session #2, spatially coherent: 26 cm vs. 53 cm, p=0.02, real vs. shuffled; spatially incoherent: 210 cm vs. 263 cm, p=0.02, real vs. shuffled; ***Figure 4–Figure Supplement 2***).

Many replay events are slow or stationary

Surprisingly, our results indicate that most events with coherent spatial content are not dominated by the continuous movement dynamic, but instead correspond to trajectories that are stationary or that move relatively slowly compared to their purely continuous counterparts. We therefore examined these events in more detail. We first note that most of the SWRs (14,989 of 21,433 or 70% of classified SWRs) were classified as having only a single dynamic or dynamic mixture (***Figure 5***A, for specific example SWR see ***Figure 3***B), whereas SWRs with multiple dynamics or dynamic mixtures (such as those in ***Figure 3***A, C, D, E, and F) were less common. Interestingly, the most common category of SWR were those with only stationary-continuous mixtures (8,944 of 21,433 or 42% of classified SWRs, ***Figure 5***A, ***Figure 5–Figure Supplement 1***). These events contain representations that move at slower speeds (***Figure 5***D)—similar to those speeds experienced by the animal in the

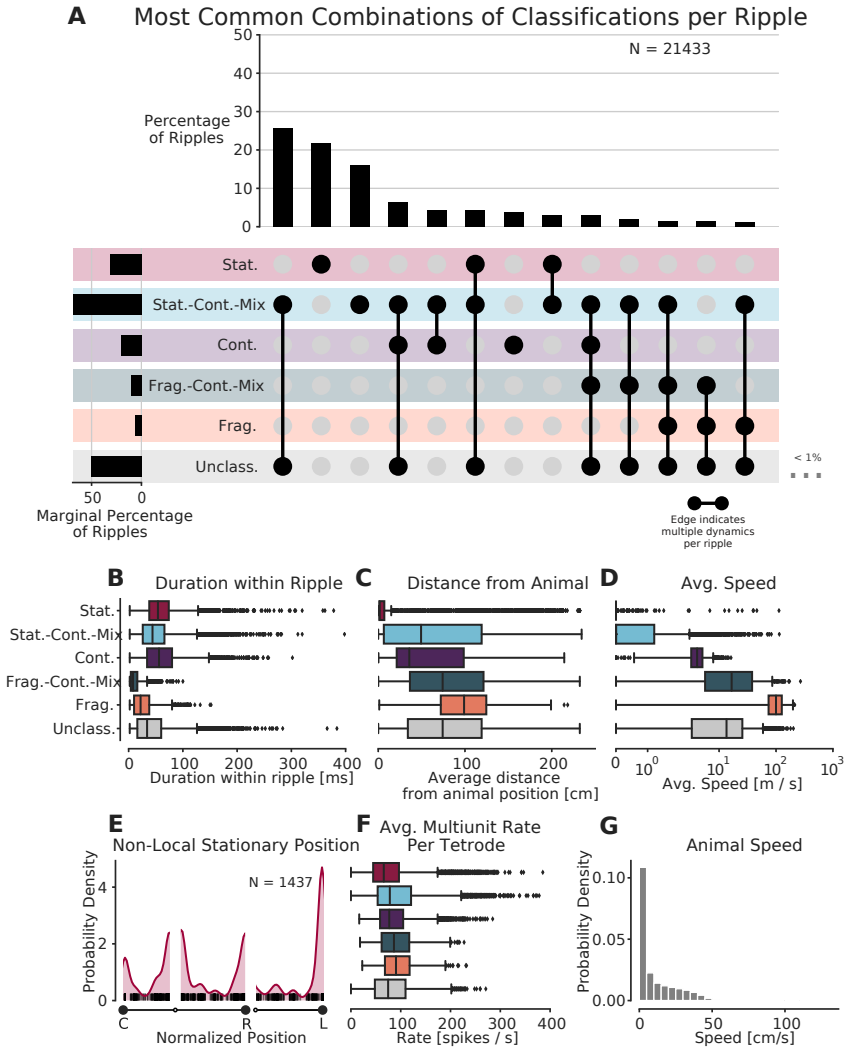

Figure 5. Prevalence of classifications. **(A)** UpSet plot (***Lex et al., 2014***)—which is similar to a Venn diagram with more than 3 sets—of the most common sets of classifications within each SWR. Each row represents a classification and each column represents a set of classifications, where filled-in black dots with an edge between the dots indicates that multiple classifications are present in the SWR (at different times). The sets are ordered by how often they occur as indicated by the bar plot above each category. The total percentage of each classification is indicated by the rotated bar plot on the left. **(B)** Duration of each dynamic category within a SWR. The box shows the interquartile range (25-75%) of the data and the whiskers are 1.5 times the interquartile range. Outlier points outside of this are labeled as diamonds. **(C)** Average distance between the latent position and the animal's actual position for each classification within the SWR. **(D)** Average speed of the classification within the SWR (excluding classifications with durations < 20 ms). Note that these speeds are calculated using the most probable position (MAP estimate) which can be noisy when the probability of position is flat or multimodal.**(E)** Kernel density estimate of the position of stationary trajectories on the W-track at least 30 cm away from the animal's position. The shaded region represents the density estimate while the black ticks represent the observed non-local stationary positions. **(F)** Average tetrode multiunit spike rates for each dynamic category within each SWR (excluding classifications < 20 ms). **(G)** Probability density of animal movement speeds, illustrating prevalence of slower speed real-world movement consistent with stationary and stationary-continuous mixture replay events.

Figure 5–Figure supplement 1. More examples of stationary-continuous-mixtures.

Figure 5–Figure supplement 2. Control analyses for distribution of dynamics

Figure 5–Figure supplement 3. Further quantification of spiking and ripple properties by dynamic

Figure 5–video 1. Example of a stationary-continuous-mixture. Conventions the same as Figure 2-video1

environment (***Figure 5***G)—and are sustained, but slightly shorter, on average, than events with continuous dynamics (median duration: stationary-continuous-mixture 73 ms vs. continuous 94 ms, ***Figure 5***B). Nonetheless, both the slow-moving events and continuous events frequently represent locations that were some distance away from the animal (mean trajectory distance from animal's position: stationary-continuous-mixture 52 cm vs. continuous 43 cm, ***Figure 5***C). This indicates that the content of these SWRs, like those that are classified as continuous, do not represent the animal's actual position. We note that roughly a third of these events persisted for the entire duration of the SWR, but the other two thirds included some time where the model was uncertain about the dynamic. This shows the power of our state space decoder because it allows us identify periods of time when the model has high confidence with each SWR, and use these periods to characterize the SWR overall or to focus on specifically for further analysis.

The second most prevalent classification was exclusively stationary events (4,858 of 21,433 or 23% classified SWRs; 47% of SWRs with stationary dynamics were entirely stationary for the entire duration of the SWR). Unlike the stationary-continuous mixtures, most of these events represented a location close to the animal's actual position (***Figure 5***C). There were, however, a number of stationary events that represented positions far from the animal. We set a threshold of 30 cm distance from the animal to identify non-local stationary events and found such content in ~7% of classified SWRs (1,437 of 21,433 classified SWRs). Of these, 46% (664 of 1,437 stationary SWRs) were stationary throughout the duration of the SWR. These non-local stationary events most commonly represented reward wells or choice points (***Figure 5***E), consistent with (***Yu et al., 2017***), but a small portion of them occurred at other locations of the track. This suggests that these representations can be flexibly deployed to represent discrete locations.

Control analyses

We then carried out a series of control analyses to determine whether our results could be due to differences in spiking statistics across dynamics or interneuron-specific activity patterns. We first calculated the multiunit spiking rates and found that these were similar across the different dynamics. (***Figure 5***F). This indicates that all dynamics, including stationary and stationary-continuous mixtures, were driven by comparable levels of sustained spiking information and could not be explained by the absence of spiking.

We also verified that our classifications were very unlikely to be driven by the firing of interneurons, which are less likely to project to other brain areas and exhibit less spatial specificity than principal neurons in the hippocampus (***Hangya et al., 2010***; ***Wilent and Nitz, 2007***; ***Jinno et al., 2007***)). First, we increased our tetrode participation threshold from two up to five tetrodes active during the SWR, to eliminate the possibility that events with a single tetrode with a high rate interneuron might drive the classification of stationary or stationary-continuous mixture dynamics during SWRs. We found that our results were robust to this increase in threshold (***Figure 5–Figure Supplement 2***A). Indeed, most of our events and most periods of dynamics within each event included spikes from upwards of 10 tetrodes (***Figure 5–Figure Supplement 3***). Next, we removed spikes that had spike widths of less than 0.3 ms to reduce the potential contribution of narrow-waveform inhibitory interneurons (***Fox and Ranck Jr, 1975***). Similarly, this had little effect: we still observed a high proportion of SWRs with stationary and stationary-continuous mixture dynamics (***Figure 5–Figure Supplement 2***B). Using clustered data, we can more robustly categorize units as putative interneurons; thus, we then repeated this analysis on clustered data for 9 of the 10 animals, excluding putative interneurons based on spike width and firing rate and requiring that at least three putative pyramidal cells be active during the SWR. Although there were many fewer events to examine, we still observed many stationary and stationary-continuous mixtures in SWRs (***Figure 5–Figure Supplement 2***D). Finally, if interneurons with less spatial specificity were driving the classification of slower dynamics, we would expect to see the posterior be less spatially concentrated during these events. Instead, we found that stationary and stationary-continuous mixtures tended to have narrow posteriors (as characterized by 95% HPD region size in ***Figure 4***E).

This provides further evidence that our large fraction of slow-speed events were not the solely the result of the firing of interneurons.

We then asked whether the slow spatially coherent events (stationary-continuous mixtures) potentially reflect a slow theta sweep rather than a replay event. We noted that this seems very unlikely given that the stationary-continuous mixtures tended to represent locations far from the animal that would require rapid, long-distance theta sequences (median 50 cm ***Figure 5***C). In addition, the majority of our events occur when the animal is moving at speeds < 2 cm/s, as expected for SWRs (***Figure 5–Figure Supplement 3***B). Further, when we restrict the analyses to SWRs where the animal was moving at speeds < 1 cm/s, we find similar proportion of SWRs with stationary and stationary-continuous mixture dynamics (***Figure 5–Figure Supplement 2***C).

A third possible concern is that burst firing from a single cell is driving the stationary dynamics. There are several lines of evidence against this. First, as mentioned above, most of our SWRs contained spikes from multiple tetrodes, including those with stationary dynamics (***Figure 5–Figure Supplement 3***A), so it is unlikely that a single cell could drive these dynamics. Second, the stationary dynamics have a longer duration than one would expect from a burst alone (***Figure 5***B, 54 ms median duration, 38-74 ms quartiles vs. bursts which have a typical duration of 6-24 ms, (***Ranck Jr, 1973***; ***Harris et al., 2001***; ***Tropp Sneider et al., 2006***)). Third, only ~50% of the spikes had interspike intervals (ISIs) of less than 6 ms during the stationary, stationary-continuous, and continuous dynamics (***Figure 5–Figure Supplement 3***A), indicating that rapid bursts made up at most half of the observed spikes, and likely fewer given that these ISIs are computed for each tetrode, rather than single units.

Finally, we wished to assess whether the spiking information from different hippocampal subfields (CA1, CA2, and CA3), could influence the types of dynamics we observed. This seemed unlikely, given that CA1 and CA3 spiking is tightly coordinated within and across hemispheres during SWR replay (***Carr et al., 2012***) and multiple previous papers have combined recordings across subfields and obtained results similar to those that restricted their analyses to CA1 (***Diba and Buzsáki, 2007***; ***Karlsson and Frank, 2009***). However, we verified that when we used only CA1 tetrodes to decode (***Figure 5–Figure Supplement 2***E), we get similar proportions of dynamics as when using all the tetrodes.

Comparison to standard approaches

Finally, we asked how our method compares to the standard approaches to identifying significant replay events via a non-parametric shuffle test and characterizing their content. This comparison is complicated due to the variety of parameters and inference methods used previously, so here we chose three commonly used approaches and a reasonable parameter set for the comparison. We began with an examination of events based on statistical significance of a line fit to a circularly permuted spatial distribution, which has been the standard in the field. We applied the clusterless algorithm in 20 ms bins to produce probability densities of represented locations, as this binning is comparable to previous methods. We then used the Radon transform method, which finds the line with the highest summed posterior probability of position along the line (***Davidson et al., 2009***; ***Farooq and Dragoi, 2019***), to calculate, for each SWR, a measure of significance based on the assumption that the probability density moves at a constant speed throughout the event. We found that ~12% of the classified events (as defined by our state-space decoder) were identified as significant. Not surprisingly, the fraction of significant events was strongly related to which dynamic they contained. Of the SWRs that contained a period of continuous dynamics, 31% (1,386 of 4,511 SWRs) were significant. Of all SWRs that contained a period of stationary dynamics, 8% (542 of 7,176 SWRs) were significant, and when we restricted the analysis to long stationary events (>100 ms duration), we found that 28% were significant. This increase is not surprising, since the shuffle used for the Radon method penalizes shorter events. An alternative approach, the linear regression method, which samples from the posterior probability and uses linear regression to fit a line to the samples (***Karlsson and Frank, 2009***; ***Carr et al., 2011***; ***Shin et al., 2019***), detects only non-zero

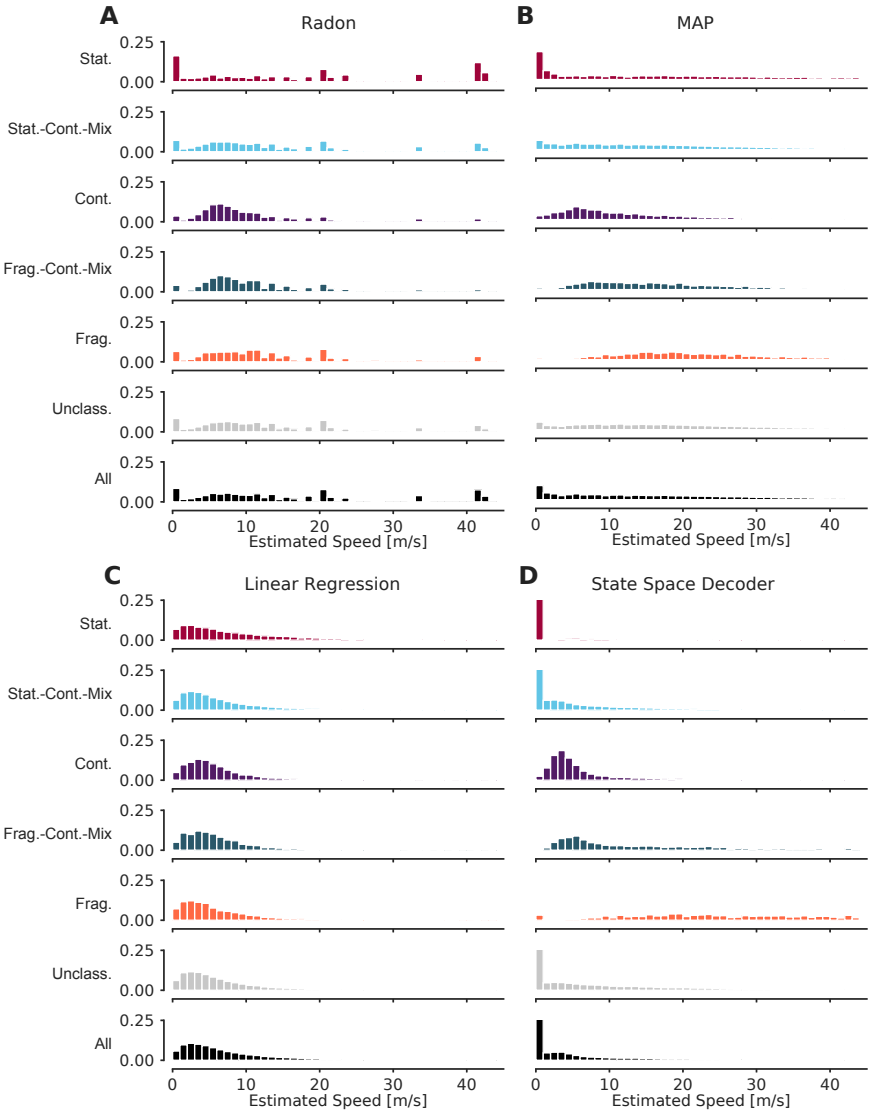

Figure 6. Event speeds calculated using three common "Bayesian" decoder approaches (A-C) compared to using the state space model (D). For each panel, the top five rows show the probability density of the estimated speed for SWRs that contain that dynamic or combination of dynamics. Note that these categories are not mutually exclusive since a SWR can include more than one dynamic or combination of dynamics. The sixth row shows all SWRs containing unclassified dynamics. The final row is the estimated speed for all SWRs. **(A)** Radon transform **(B)** MAP estimate **(C)** Linear regression **(D)** State space decoder presented in this manuscript.

slopes as significant, and thus by definition cannot identify stationary events as significant.

Conclusions about the representational content of events were similarly variable across methods. In addition to the Radon and linear regression methods, we also computed the MAP estimate, the maximum probability position of each 20 ms time bin (***Pfeiffer and Foster, 2015***; ***Stella et al., 2019***; ***Kaefer et al., 2020***). We then compared the speed of the replay for each dynamic and for the entire SWR (***Figure 6***). We found that the methods that rely on fitting a line for the entire SWR (the Radon transform and the linear regression), tend to estimate that the replays proceed at faster speeds (***Figure 6***A,C), whereas the MAP method (***Figure 6***B), which only relies on picking the most probable position for a given time bin and allows for more variable speeds, can capture both fast and slow speeds. In addition, the speeds from the MAP method (***Figure 6***B) are quite similar to the speeds estimated by our decoder (***Figure 6***D). This is to be expected because the MAP method imposes less smoothing (only the 20 ms time bins) whereas the linear methods impose a larger amount of smoothing. Our decoder imposes more smoothing than the MAP estimate but much less than the linear methods and therefore is closer to the MAP estimates. These results indicate that the methods that made more assumptions about the structure of replay (Radon and linear regression approaches) yield different conclusions about which events are potentially meaningful and how those events progress over time from methods that make fewer assumptions (the MAP and state-space approaches). The substantial variability in both the proportions of "significant" events and the speed of the events demonstrates that key conclusions about the events could be very different depending on the specific method and parameters used.

Discussion

We developed a state space model that identifies and quantifies replay in terms of a mixture of three movement dynamics: stationary, continuous, and fragmented. This model is robust: it is relatively insensitive to a range of plausible transition parameters between dynamics and instead strongly reflects the hippocampal place representation structure. We show that this model is interpretable: it has well-defined parameters that capture intuitive movement models and allows us to explain a large fraction of SWRs in terms of these dynamics—far more than have been previously associated with identifiable activity patterns. This model is also flexible: it can be used with sorted or clusterless spikes, it can be used with 1D or 2D positions, and—because it uses a mixture of dynamics—it allows us to confidently discover not only SWRs with spatial representations that move at a constant speed, but also representations that vary in speed, have slower speeds, or that are stationary. In fact, these slower-moving representations constitute the majority of representations seen during SWRs, but would not be found with the typical approaches that assume that events progress at constant speed. The prevalence of these events indicates that hippocampal replay can recapitulate experience at both accelerated and closer to real-world speeds.

Previous work reported that less than half of SWRs or high-population-activity events contained replay events (***Karlsson and Frank, 2009***; ***Foster and Wilson, 2006***; ***Davidson et al., 2009***). This could lead one to the assumption that most events are not spatially coherent or meaningful. Our results indicate that this is not the case: in our dataset, the large majority of SWRs (86%) contained spatially coherent content, as defined by some combination of stationary and continuous dynamics. These events were mostly decoded representations where the most probable positions were highly spatially concentrated (i.e. had small HPD region size), indicating that they were driven by spatially specific firing. This was in contrast to unclassified events or events that included fragmented dynamics, which were associated with highly distributed decoded spatial representations.

Importantly, the spatially coherent events most often expressed a dynamic that was consistent with slower speeds of less than 1 m/s (<100 cm/s), which is consistent with the speed that rats move through their environment. The next most common category were stationary events (***Yu et al., 2017***; ***Farooq and Dragoi, 2019***) that activated a persistent representation of a single location, while the "classical" replay events—continuous, extended trajectories through space—made up only about 19% of classified events. Events with slower dynamics tended to persist for approximately

the same length of time as events with faster continuous dynamics (~50-100 ms), consistent with a brief activation of a representation of a location or a short trajectory through space. Events with slower dynamics most often represented the animal's current location, similar to continuous events identified in previous work (***Davidson et al., 2009***; ***Karlsson and Frank, 2009***); however, slower events were also capable of representing locations far from the animal.

These results challenge the long-held notion that hippocampal replay trajectories necessarily proceed at many times the speed of the animal (***Nádasdy et al., 1999***; ***Lee and Wilson, 2002***; ***Davidson et al., 2009***). Instead, replay events encompass a much richer variety of event speeds that could promote storage and updating of memories on different spatial and temporal scales, including slower moving representations that activate a spatially compact set of nearby locations and stationary representations of single locations. Interestingly, a previous paper (***Farooq and Dragoi, 2019***) reported that these stationary representations were most prevalent in juvenile rats and decreased to "chance" levels when the rats reached adulthood. Our results strongly suggest that they do not disappear, but are a very common occurrence with adult rats. We believe that the difference in these findings is a result of defining "chance" levels of occurrence of these events and using a restrictive decoding model which requires the decoded position to be stationary for the entire event. We also note that we chose to include a stationary dynamic to explicitly evaluate the prevalence of stationary events in adult animals, but in principle, a model with just continuous and fragmented dynamics could capture these events as well, as the continuous dynamic includes movement speeds close to 0 m/s.

The differences between our findings and previous approaches also highlight the importance of identifying the specific question that an analysis seeks to answer. Initial studies of replay (***Lee and Wilson, 2002***; ***Foster and Wilson, 2006***; ***Davidson et al., 2009***; ***Karlsson and Frank, 2009***) focused on the question of whether the observed patterns of spiking seen during replay reflected structured content or "random" spiking as defined by a comparison to one or more shuffles. These and subsequent studies answered the question "are replay events the result of random spiking?" with a clear "no". These studies did not, however, establish that the significance criteria used to show non-random structure has anything to do with the presence of spatial structure or events that are potentially useful for memory processes.

Our method aims to identify structure within events and to capture their complex dynamics to determine what sort of information replay events could provide in service of memory functions. We precisely define patterns of spatial representation that are, and are not, interpretable in terms of the structure of a given environment and the spatial activity seen during movement. In so doing, we find that most, but not all, events contain spiking that corresponds to a location or a trajectory through the animal's environment.

The different questions addressed by our method compared to previous methods can also explain the apparent differences in findings. Our results showed that 19% of SWRs contained continuous trajectories—meaning that portions of the SWR had spatially coherent trajectories that moved at high velocities—but only 4% were continuous throughout the entire SWR. This, *prima facie*, may seem at odds with previous results, which found between 5% and 45% of SWRs were "significant", meaning that they were well fit by a line compared to various shuffled versions of the dataset. Furthermore, these studies found trajectories on average move at much higher speeds (20x normal rat movement speed) than we report here. These previous findings were based largely on a single linear fits to the data for each event, while our model allows the trajectory to change speed during the SWR, which can lead to very different estimates of the speed of the event. Many of the events that had continuous dynamics, and would be classified as having high speed trajectories, also exhibited stationary-continuous mixture dynamics for some portion of the event(see for example ***Figure 2***). Further, because the linear fit includes the whole SWR event, it is highly sensitive to the definition of the start and end of the SWR event, an often arbitrary definition which varies across research groups. In contrast, our model makes an estimate for each time point and only depends on the previous and future time step, which minimizes the dependence of the

event on those boundaries and any possible "noise" activity within the SWR. This, along with our more permissive threshold for SWRs, allows us to characterize more SWRs than previous studies, rather than exclude them from analysis.

Our results may also explain the seeming conflict between ***Stella et al.*** (***2019***), who reported that hippocampal replay trajectories followed Brownian-like diffusion dynamics, and ***Pfeiffer and Foster*** (***2015***), who reported that a subset of replay trajectories contained systematic jumps in position. We found that while replay speeds are, on average, distributed log normally, a small subset of the replays include definitive "jumps" from one position to another. Because our model makes it possible to identify different dynamics within the same SWR, we are able to identify both types of sequences.

We hypothesize that this large set of events supports the ability to store and update representations of individual places and snippets of experience. Previous findings demonstrated that SWRs with movement- or immobility-related activity engage specific patterns of spiking activity outside the hippocampus (***Yu et al., 2017***; ***Jadhav et al., 2016***). Thus, we would expect that the large set of events with slow dynamics would also engage specific patterns of activity outside the hippocampus, perhaps allowing animals to store or update the value of specific locations or short movement trajectories without having to generate full trajectories representing one of multiple ways to arrive at a location (***Yu and Frank, 2015***). Events that contain multiple dynamics (e.g. two continuous representations separated by a brief fragmented representation) might also help bind together experiences that occur in the same environment but at different places, whereas sustained spatially incoherent representations might represent coherent representations in other spatial environments (***Karlsson and Frank, 2009***) or non-spatial experiences. Alternatively, downstream structures may be are attuned to the shifts in dynamics, and that each period of coherent hippocampal dynamics drives its own set of responses elsewhere in the brain. We further note that recent results from a study of activity patterns during movement indicates that individual theta cycles can have different dynamics as well (***Wang et al., 2020***), suggesting that the underlying sequences seen during replay and theta may overlap in their dynamics as well as their content.

Similarly, it is possible that slow and fast dynamics are supported by different subcircuits within the hippocampal formation. ***Yamamoto and Tonegawa*** (***2017***) found that blocking medial entorhinal cortex (MEC) output to CA1 resulted in shorter (and therefore slower) replay trajectories and ***Oliva et al.*** (***2018***) showed that layer 2/3 MEC neurons increased their firing 50 ms before the start of longer duration SWRs (>100 ms). This raises the possibility that these slower replays are driven by CA3-CA1 interaction while the faster trajectories are more externally driven by MEC-CA1, and that these different events could also engage different circuits outside of the hippocampus and entorhinal cortex.

One key aspect to understanding how the hippocampus engages with other brain structures is identifying the times when the neural dynamics have changed. The times of SWRs or increases in multiunit activity has provided an important, but coarse, landmark to this change in neural dynamics. Because our method allows for moment-by-moment characterization of the neural dynamic, we suggest that evaluating the start and end of each dynamic may prove a powerful way identify periods of interest. Examination of such periods may advance our understanding of when and why the hippocampus may be representing current position, past experience, or possible future locations, how inputs to the hippocampus influence these events, and how these events may in turn influence downstream structures.

Our identification of the rich and varied dynamics in SWRs depended on several factors. First, we sought to describe as many SWRs as possible, instead of setting a high threshold for event detection. This more permissive approach was also critical for the identification of events that activate specific, immobility-associated locations in previous work (***Yu et al., 2017***). More broadly, there is no evidence that SWR and high population activity events constitute clearly distinct classes of events that can be identified with a single fixed threshold. Instead, the sizes of these events are well described by a unimodal, long-tailed distribution (***Yu et al., 2017***). High thresholds can

therefore exclude large numbers of real events and potentially lead to erroneous conclusions.

Second, we used clusterless decoding, which enabled us to take advantage of information from neurons farther away from the tetrodes and not just those that could be cleanly and somewhat subjectively separated (***Chen et al., 2012b***; ***Kloosterman et al., 2014***; ***Deng et al., 2016***). We note that more fully automated modern spike sorting algorithms (such as MountainSort (***Chung et al., 2017***)) in combination with better recording technology could reduce the differences in results when using sorted spikes or clusterless decoding, as these sorting methods help reduce experimenter subjectivity involved in spike sorting and identify larger numbers of putative single neurons.

Third, we built an explicit and flexible model that allowed us to identify not just one dynamic, but multiple dynamics consistent with different speeds of motion. We used these dynamics to describe the SWR events, rather than declaring the events as significant or not.

Fourth, our model avoided using large temporal bins and does not make assumptions about the linear progression of events. Instead, the model allows for movement in any direction allowed by geometry of the environment, and further, still accounts for the uncertainty of the data. Because the model gives a moment-by-moment estimate of the position and dynamic, this also minimizes the effect of misidentifying the SWR start and end times. Combined with an acausal estimation of latent position, this approach allows us to identify the specific timing of transitions in dynamics at a fine timescale. In addition to our SWR-based analyses, this approach could also be used to decode activity during movement, where we might expect that place-specific activity would express combinations of stationary and continuous dynamics reflecting different speeds of representational movement during each cycle of the theta rhythm.

Our model can be viewed as an extension to the multiple temporal models approach of ***Johnson et al.*** (***2008***), which used model selection to determine the model with the most appropriate trajectory speed. Our model goes beyond the Johnson model in that it explicitly permits a mixture between the movement speeds, can work for arbitrary track geometries, and uses clusterless decoding. Our model also represents a middle ground between Hidden Markov-style models (***Maboudi et al., 2018***; ***Chen et al., 2016***; ***Linderman et al., 2016***; ***Chen et al., 2012a***), which seek to be environment-agnostic detectors of sequential patterns, and the typical standard decoder, which typically use arbitrarily chosen bin sizes and restrictive assumptions about the nature of the trajectories. In particular, our approach allows for a variety of dynamics while still yielding spatially interpretable results and makes it possible to use bin sizes of any size (here 2 ms).

Code and algorithms for decoding hippocampal replay have not typically been made accessible or easy to use. This is problematic because it can lead to severe variation or errors in code and parameters, limits reproducibility of results, and slows down scientific progress. Accordingly, we have made code for this model publicly available as a python package at the following URL https://github.com/Eden-Kramer-Lab/replay_trajectory_classification. It is easily installable using the pip or conda package installer and contains tutorial Jupyter notebooks in order to facilitate reuse.

State-space models like the one we use here can enable a richer set of analyses of events that take advantage of all of the available data. These approaches can be applied not just to SWRs but to all times, providing a moment-by-moment estimate of the nature of the spatial representation in the hippocampus, including important information about the spatial coherence of that representation. The model can be extended to incorporate other previously experienced environments by training place field models on those environments and including the appropriate movement transition matrices for those environments. It can also be extended to account for task conditions (such as the inbound and outbound conditions in our spatial alternation task) and forward/reverse sequences as in ***Deng et al.*** (***2016***). Finally, models can be built to estimate any covariate, including movement direction (***Kay et al., 2020***). We therefore suggest that this approach has the potential to provide a much richer understanding of neural representations throughout the brain.

Materials and Methods

Simulated Data

Encoding data for ***Figure 1*** and ***Figure 1–Figure Supplement 1*** were generated by simulating 15 traversals of a 180 cm linear track. The spiking of 19 neurons was simulated using an inhomogeneous Poisson process where the instantaneous firing rate of each neuron changes according to a Gaussian place field. The Gaussian place fields had a variance of 36 cm and were spaced at 10 cm intervals over the 180 cm track. The decoding data for ***Figure 1***F was generated by simulating 20,000 linear traversals of the 180 cm track, each with a unique constant speed, starting at 0.5 cm/s and increasing by 0.5 cm/s up to 10,000 cm/s. Each simulated neuron "spiked" when the traversal passed through the peak location of its place field.

Recording Locations and Techniques

Ten male Long Evans rats (500-700 g, 4-9 months old) were trained on a W-track spatial alternation task. 9 rats contributed to previous studies (***Karlsson and Frank, 2009***; ***Kay et al., 2020***, ***2016***; ***Carr et al., 2012***). Neural activity was recorded in CA1, CA2, CA3, MEC, Subiculum, and DG depending on the rat. We only used tetrodes located in the CA1, CA2, and CA3 subregions in this study.

Behavioral Task

All animals performed a W-track spatial alternation task, which is described in detail in ***Karlsson and Frank*** (***2009***). In brief, each day, animals alternate between 20 minute rest sessions in a rest box and 15 minutes run sessions on the W-shaped track equipped with a reward well at each arm end. On the W-track, animals are rewarded at the ends of an arm when that arm is the next correct arm in the sequence. Two rules determine the next correct arm. If the animal is in an outer arm, it must next visit the center arm. If the animal is in the center arm, it has to next visit the less recently visited outer arm. Correct performance of these rules result in the visit sequence: center, left, center, right, center, left, etc. Animals were free to choose any arm at any time. Only run recording sessions with at least 9 putative hippocampal pyramidal cells that fired at least 100 spikes were included in the analysis.

Position of the Animal and Linearization

The animal's 2D position was estimated from digital video (30 Hz) of two infrared diodes placed on the headstage preamplifiers using a semi-automated analysis. In order to decrease the time it takes to run the model, the 2D position of the animal was converted into a 1D position. This is done by first defining a 2D graph representation of the track (herein referred to as the track graph), where edges correspond to segments of the W-track and nodes represent intersection points between those segments. Then, based on the algorithm in ***Newson and Krumm*** (***2009***), we use a Hidden Markov Model (HMM) to assign the position detected in each video frame to the most likely track segment. Using the HMM takes into account the time dependent nature of the data and helps prevents sudden jumps from one track segment to another, which is particularly important near intersections. The observation model of the HMM is Gaussian and it models the likelihood of being on a track segment as the Gaussian distance to that segment with a 5 cm standard deviation. The state transition model is empirically estimated, and changes with each time point to ensure that the Euclidean distance between successive position estimates is similar to the shortest path distance along the graph between successive position estimates. A slight bias of 0.1 is given to the diagonal of the state transition model to encourage staying on the same track segment. The most likely track segment the animal is on is computed using the Viterbi algorithm (***Viterbi, 1967***). After finding the track segment that corresponds to each 2D position, the 2D position is projected onto the nearest point of the track segment. This allows us to define a distance from the center well in terms of shortest path length on the track, where 0 cm represents the center well position. The linear distance can then be converted into a linear

position by assigning each track segment a position in 1D space. 15 cm gaps were placed between the center arm, left arm, and right arms in 1D space to prevent any smoothing done in the model from influencing neighboring segments inappropriately. The code used for linearization can be found at https://github.com/Eden-Kramer-Lab/loren_frank_data_processing.

Sorted Spikes, Multiunit Spikes and Waveform Features

To obtain the neural spiking data used for decoding, electrical potentials from rat hippocampus were recorded at 30 kHz, referenced to a tetrode located in the corpus callosum, and then digitally filtered between 600 Hz and 6 kHz. Spiking events were detected as any potential exceeding a 60 μV threshold on any one of the four tetrode wires of a tetrode. The electrical potential value on each wire of the tetrode at the time of maximum potential of any of the four wires was used as the waveform feature in the clusterless decoding.

For decoding using sorted spikes, the multiunit events were processed further to assign events to putative single cells. Putative single cells were manually identified based on the clustering of three waveform features within a day: peak amplitude, principal components, and spike width. Only putative hippocampal pyramidal cells—identified based on spike width and average firing rate—were included in the analysis.

SWR Detection

Sharp wave ripples were detected using the same method as in ***Kay et al.*** (***2016***). Each CA1 LFP was obtained by downsampling the original 30 kHz electrical potential to 1.5 kHz and bandpass filtering between 0.5 Hz and 400 Hz. This was further bandpass filtered for the ripple band (150-250 Hz), squared, and then summed across tetrodes—forming a single population trace over time. This trace was smoothed with a Gaussian with a 4 ms standard deviation and the square root of this trace was taken to get an estimate of the population ripple band power. Candidate SWR times were found by z-scoring the population power trace of an entire recording session and finding times when the z-score exceeded 2 standard deviations for a minimum of 15 ms and the speed of the animal was less than 4 cm/s. The SWR times were then extended before and after the threshold crossings to include the time until the population trace returned to the mean value. The code used for ripple detection can be found at https://github.com/Eden-Kramer-Lab/ripple_detection. We only analyzed SWRs with spikes from at least 2 tetrodes.

The Model

Let x_k be a continuous latent variable that corresponds to the position represented by the population of cells at time t_k and let I_k be a discrete latent variable that is an indicator for the movement dynamics we wish to characterize: stationary, continuous, and fragmented. The goal of the model is to estimate simultaneously the posterior probability of position and dynamics $p(x_k, I_k \mid O_{1:T})$, where $O_{1:T}$ corresponds to the observed spiking data from time 1 to time T. The observed data can be either spike trains $\Delta N_{1:T}^{(1:C)}$ from C putative cells when decoding with sorted spikes or multiunit spikes $\Delta N_{1:T}^{(1:E)}$ and their associated wave form features $\vec{m}_{k,j}^{i}$ from each tetrode E when decoding with clusterless spikes, where $i \in 1:E$, $k \in 1:T$, $j \in 1:\Delta N_k^i$.

We have previously shown (***Denovellis et al., 2019***) that the posterior probability $p(x_k, I_k \mid O_{1:T})$ can be estimated by applying the following recursive causal filter equation, starting with initial conditions $p(x_0, I_0)$ and iterating to time T:

$$p(x_k, I_k \mid O_{1:k}) \propto p(O_k \mid x_k, I_k) \sum_{I_{k-1}} \int p(x_k \mid x_{k-1}, I_k, I_{k-1}) Pr(I_k \mid I_{k-1}) p(x_{k-1}, I_{k-1} \mid O_{1:k-1}) dx_{k-1}$$

and then applying the acausal smoother equation, starting from the last estimate of the casual filter $p(x_T, I_T \mid O_{1:T})$ and recursively iterating backwards to time 1:

$$p(x_k, I_k \mid O_{1:T}) = p(x_k, I_k \mid O_{1:k}) \sum_{I_{k+1}} \int \frac{p(x_{k+1} \mid x_k, I_{k+1}, I_k) Pr(I_{k+1} \mid I_k)}{p(x_{k+1}, I_{k+1} \mid O_{1:k})} p(x_{k+1}, I_{k+1} \mid O_{1:T}) dx_{k+1}$$

where:

$$p(x_{k+1}, I_{k+1} \mid O_{1:k}) = \sum_{I_k} \int p(x_{k+1} \mid x_k, I_{k+1}, I_k) Pr(I_{k+1} \mid I_k) p(x_k, I_k \mid O_{1:k}) dx_k$$

Therefore, to specify the model, we have to define or estimate the following quantities:

1. $p(x_0, I_0)$ - the initial conditions
2. $Pr(I_k \mid I_{k-1})$ - the dynamics transition matrix
3. $p(x_k \mid x_{k-1}, I_k, I_{k-1})$ - the dynamics movement model
4. $p(O_k \mid x_k, I_k)$ - the likelihood of the observations

For the initial conditions $p(x_0, I_0)$, we set each dynamic I_0 to have uniform $1/3$ probability and each initial latent position to have uniform probability density over all possible positions $\mathcal{U}(\min x, \max x)$, reflecting the fact that we do not have any prior knowledge about which dynamic or position is more likely:

$$p(x_0, I_0) = \frac{1}{3} \mathcal{U}(\min x, \max x)$$

For the dynamics transition matrix $Pr(I_k \mid I_{k-1})$, which defines how likely the dynamic is to change to another dynamic versus persist in the same dynamic, we set it to be:

$$Pr(I_k \mid I_{k-1}) = \begin{array}{c} \\ \\ \\ \end{array} \begin{array}{c} \begin{array}{ccc} I_k = Stationary & I_k = Continuous & I_k = Fragmented \end{array} \\ \begin{bmatrix} 0.98 & 0.01 & 0.01 \\ 0.01 & 0.98 & 0.01 \\ 0.01 & 0.01 & 0.98 \end{bmatrix} \end{array} \begin{array}{l} \\ I_{k-1} = Stationary \\ I_{k-1} = Continuous \\ I_{k-1} = Fragmented \end{array}$$

to encode the prior expectation that each of the dynamics will last the average duration of 100 ms, with a small probability of changing to one of the other dynamics.

For the dynamics movement model $p(x_k \mid x_{k-1}, I_k, I_{k-1})$, which defines how likely the position x_k is to change given the previous position x_{k-1} and current I_k and past dynamics I_{k-1}, we set it to be:

$$p(x_k \mid x_{k-1}, I_k, I_{k-1}) = \begin{array}{c} \begin{array}{ccc} I_k = Stationary & I_k = Continuous & I_k = Fragmented \end{array} \\ \begin{bmatrix} \delta(x_{k-1}) & \mathcal{N}(x_{k-1}, 6.0) & \mathcal{U}(\min x, \max x) \\ \delta(x_{k-1}) & \mathcal{N}(x_{k-1}, 6.0) & \mathcal{U}(\min x, \max x) \\ \mathcal{U}(\min x, \max x) & \mathcal{U}(\min x, \max x) & \mathcal{U}(\min x, \max x) \end{bmatrix} \end{array} \begin{array}{l} \\ I_{k-1} = Stationary \\ I_{k-1} = Continuous \\ I_{k-1} = Fragmented \end{array}$$

where $\delta(x_{k-1})$ is an identity transition matrix where position cannot change from the previous time step, $\mathcal{N}(x_{k-1}, 6.0)$ is a random walk from the previous position with variance 6.0, and $\mathcal{U}(\min x, \max x)$ is a uniform transition that allows transitions to any possible position. As discussed in the Results, this means that when persisting in the same dynamic, the stationary, continuous, and fragmented dynamics are defined by the identity transition, the random walk, and the uniform transition, respectively. When transitioning to or from the fragmented dynamic, we assume we do not have any information about the position, so the transition is uniform. Finally, when the transition is from the stationary to continuous, we assume the position is spatially close where it was previously, so we use a random walk. When the transition is from continuous to stationary, we assume that the position is no longer changing, so we use the identity transition.

Lastly, we evaluate the likelihood of the observations $p(O_k \mid x_k, I_k)$ based on an encoding model fit during the encoding period. We assume the likelihood is the same for each dynamic I_k, so we only need to evaluate $p(O_k \mid x_k)$. It has been shown (***Zhang et al., 1998***; ***Brown et al., 1998***) that the Poisson likelihood with sorted spikes can be computed as:

$$p(O_k \mid x_k) = p(\Delta N_k^{(1:C)} \mid x_k) \propto \prod_{i=1}^{C} [\lambda_i(t_k \mid x_k)\Delta_k]^{N_{t_k}^i} \exp[-\lambda_i(t_k \mid x_k)\Delta_k]$$

where $N^i_{t_k}$ represents a spike at time t_k from cell i, Δ_k is the time bin size, and $\lambda_i(t_k \mid x_k)$ is the instantaneous firing rate of cell i given position x_k. $\lambda_i(t_k \mid x_k)$ is the "place field" of the cell, which can be estimated by fitting a generalized linear model to each cell's spiking during the encoding period.

Likewise, it has been shown (***Chen et al., 2012b***; ***Deng et al., 2016***; ***Kloosterman et al., 2014***) that the clusterless likelihood can be computed as:

$$p(O_k \mid x_k) = p(\Delta N_k^{(1:E)}, \{\vec{m}^i_{k,j}\}^{i=1:E}_{j=1:\Delta N^i_k} \mid x_k) \propto \prod_{i=1}^{E} \prod_{j=1}^{\Delta N^i_k} [\lambda_i(t_k, \vec{m}^i_{k,j} \mid x_k)\Delta_k] \exp[-\Lambda_i(t_k \mid x_k)\Delta_k]$$

where $\lambda_i(t_k, \vec{m}^i_{k,j} \mid x_k)$ is now a generalized firing rate that depends on an associated wave form features $\vec{m}$ and $\Lambda_i(t_k \mid x_k)$ is a marginal firing rate that is equivalent to a place field estimated on multiunit spikes. Both of these rates can be defined as:

$$\lambda_i(t_k, \vec{m}^i_{k,j} \mid x_k) = \mu_i \frac{p_i(x_k, \vec{m}^i_{k,j})}{\pi(x)}$$

and

$$\Lambda_i(t_k \mid x_k) = \mu_i \frac{p_i(x_k)}{\pi(x)}$$

where μ_i is the mean firing rate for tetrode i, $\pi(x)$ is the smoothed spatial occupancy of the animal on the track, $p_i(x_k)$ is the smoothed spatial occupancy only at times when spikes occur in tetrode i, and $p_i(x_k, \vec{m}^i_{k,j})$ is the smoothed occupancy in the space of both space and waveform features. $\pi(x)$, $p_i(x_k, \vec{m}^i_{k,j})$, $p_i(x_k)$ can all be estimated by training a kernel density estimator on each tetrode's spike waveform features and corresponding position during the encoding period.

Encoding - Sorted Spikes

In order to encode how each cell's spiking activity relates to position (the place field), we fit a generalized linear model (GLM) with a Poisson response distribution to each cell's spiking activity during the encoding period, which we define as all movement times (time periods when the running speed is greater than 4 cm/s). We estimate the parameters β, which consist of β_0, the baseline firing rate over time, and β_i, weights for third degree B-spline basis functions f_i over position (or tensor products of the B-splines when position is two dimensional). B-spline basis functions are used because place field firing activity is assumed to vary smoothly over position and this prior knowledge can be exploited to reduce the total number of model parameters needed. Each basis function is spaced every 5 cm over the range of the position and zero constrained so that the change encoded by the parameters is relative to the baseline firing rate. We use a log link function to convert the linear combination of parameters to an instantaneous firing rate over time $\lambda(t_k)$ to ensure the rate is always positive.

$$log(\lambda(t_k)) = \beta_0 + \sum_i f_i(x_k)\beta_i$$

A small L2 penalization term $-\lambda\|\beta_i\|_2^2$ is used to prevent model fitting instability when spiking activity is very low. We set this to 0.5 for all cells. Fitting is done by maximizing the penalized likelihood using a Newton-Raphson algorithm. The code used to fit the GLMs is available at https://github.com/Eden-Kramer-Lab/regularized_glm.

Encoding - Clusterless

In order to relate each tetrode's unsorted spiking activity and waveform features to position, we used kernel density estimation (KDE) to estimate the following distributions: $\pi(x)$, $p_i(x_k)$, and $p_i(x_k, \vec{m}_k)$. We used KDEs of the form:

$$kde(y) = \frac{1}{N h_1 \cdots h_D} \sum_{i=1}^{N} \prod_{d=1}^{D} K_d \left(\frac{y_d - y_{i,d}}{h_d} \right)$$

where y is the data with D dimensions, K is a 1 dimensional Gaussian kernel with bandwidth h_d, $y_{i,d}$ is the data observed during movement, N is the number of observations during the movement period. For $\pi(x)$, $y_{i,d}$ is all positions observed during movement. For $p_i(x_k)$, $y_{i,d}$ is all positions at the time of multiunit spikes during movement. For $p_i(x_k, \vec{m}^i_{k,j})$, $y_{i,d}$ is all positions at the time of multiunit spikes and their associated waveform features during movement. We choose the bandwidth h_d to be 6.0 cm for any position dimension and 24.0 μV for any spike amplitude dimension, because these parameters were found to minimize decoding error in ***Kloosterman et al.*** (***2014***). The mean firing rate μ_k was also estimated during movement.

Decoding

In order to decode $p(x_k, I_k \mid O_{1:T})$, we used a grid based approximation of the latent position x_k that respected the geometry of the track. For 1D linearized positions, we discretized the position space based on the same track graph we used for linearization by finding bins less than or equal to the chosen position bin size for each edge of the graph. For 2D positions, we discretized the position space by binning 2D positions occupied by the animal with equal sized bins of the chosen position bin size, followed by morphological opening to get rid of any holes smaller than the bin size. This was done using Scipy image processing tools (***SciPy 1.0 Contributors et al., 2020***). This grid based approximation allows us to use Riemann sums to approximate the integrals in the causal filter and acausal smoother equations. We chose a position grid with bins of 3 cm in width (and height if the model was computed with 2D positions) in order to get good resolution for the random walk transition matrices (which had 6 cm variance) as well as for the clusterless and sorted spikes decoding (which have 6 cm bandwidth for the KDE position dimensions and 5 cm spline knots for the GLM). For sorted spikes decoding, we evaluated the place field $\lambda_i(t_k \mid x_k)$ on the midpoint of these bins. Likewise, for clusterless decoding, we evaluated the spike amplitudes observed during the decoding period by evaluating the KDE for $p_i(x_k, \vec{m}^i_{k,j})$ for the midpoint of these bins.

We also used these grid bins in combination with the track graph in order to construct the appropriate 1D random walk transition matrices that respected the track geometry. To do this, we inserted the bin centers as nodes in the track graph, and then computed the shortest path distance between all pairs of position bin centers. We then evaluated a zero mean Gaussian with a variance of 6 cm on these distances to get the appropriate transition probability of moving from bin to bin.

Finally, for our SWR analysis of $p(x_k, I_k \mid O_{1:T})$, we decoded each immobility time period (times when the animal's speed was < 4 cm / s) in 2 ms time bins and extracted the SWR times.

Probability of the dynamics

The probability of the dynamics is a quantity that indicates how much that dynamic is contributing to the posterior at a given time. We estimate the probability of the dynamics by integrating out position from the joint posterior:

$$Pr(I_k \mid O_{1:T}) = \int p(x_k, I_k \mid O_{1:T}) dx_k$$

As in other calculations, the integral is approximated with a Riemann sum.

Posterior probability of position

The posterior probability of position is a quantity that indicates the most probable "mental" positions of the animal based on the data. We estimate it by marginalizing the joint probability over the dynamics.

$$p(x_k \mid O_{1:T}) = \sum_{I_k} p(x_k, I_k \mid O_{1:T})$$

Classification of Dynamics

We used a threshold of 0.80 to classify the probability of each state $Pr(I_k \mid O_{1:T})$ into 5 categories. Time periods during sharp wave ripples are labeled as stationary, continuous, or fragmented when

the probability of each individual state is above 0.80. Time periods are labeled as stationary-continuous-mixture or fragmented-continuous-mixture when the sum of stationary and continuous or fragmented and continuous are above 0.80, respectively. Time periods where none of these criterion are met are considered unclassified.

Highest Posterior Density

The 95% highest posterior density (HPD) is a measurement of the spread of the posterior probability and is defined as the region of the posterior that contains the top 95% of values of the posterior probability (***Casella and Berger, 2001***). By using only the top values, this measurement of spread is not influenced by multimodal distributions (whereas an alternative measure like the quantiles of the distribution would be). In this manuscript, we use the HPD region size—the total area of the track covered by the 95% HPD region—to evaluate the uncertainty of the posterior probability of position.

In order to calculate the HPD region size, we first calculate the 95% HPD region by determining the maximum threshold value h that fulfills the follow equality:

$$\int_{\{x:p(x_k|O_{1:T})>h\}} p(x_k \mid O_{1:T})dx = 0.95$$

The 95% HPD region is the set of position bins with posterior values greater than the threshold $\{x : p(x_k \mid O_{1:T}) > h\}$. The HPD region size is calculated by taking the integral of the members of this set:

$$\int_{\{x:p(x_k|O_{1:T})>h\}} \mathbb{1}dx$$

which we approximate with a Riemann sum.

Shuffle Analysis of the Effect of Place Encoding on Classification of Dynamics and HPD region size

In order to confirm that the model classification of dynamics depended on hippocampal place specific encoding, we resampled the position during movement with replacement, but preserved the spike times and spike waveform amplitudes. We fit the clusterless encoding model on the resampled data and then decoded the immobility periods. Like with the non-resampled data, we then extracted the SWR times and determined their classification based on our classification scheme. We repeated this shuffle analysis 50 times and compared this distribution to the real data for two recording sessions on two different animals (animal Bond, day 3, epoch 2 and animal Remy, day 35, epoch 2).

We also performed another shuffle that preserved more of the local correlation structure of spiking by shuffling the order of the well-to-well runs and then circularly shuffling the position within the runs. A well-to-well run starts when the animal moves more than 5 cm from a well and ends when the animal moves back within 5 cm of a well. This largely reflects runs from one well to another (e.g. from the center well to the left well), but for a small number of well-to-well runs, the animal did turn back to the same well (e.g. went back to the center well after leaving the center well). As above, after permuting the order of the well-to-well runs and circularly shuffling the position, we fit the clusterless encoding model and decoded SWRs. We repeated this shuffle analysis 50 times and compared this distribution to the real data for two recording sessions on two different animals (animal Bond, day 3, epoch 2 and animal Remy, day 35, epoch 2).

Distance from animal

The distance of the decoded position from the animal's is defined as the shortest path distance along the track graph between the animal's 2D position projected on to the track graph (see Linearization) and the MAP estimate of the posterior probability of position, $\arg\max_{x_k} p(x_k \mid O_{1:T})$, which is the center of the position bin with the greatest posterior value. The shortest path is found using Dijkstra's algorithm (***Dijkstra, 1959***) as implemented in NetworkX (***Hagberg et al., 2008***).

Estimation of Speed

In order to estimate the speed over time, we first found the MAP estimate of the posterior probability of position, $\arg\max_{x_k} p(x_k \mid O_{1:T})$, which is the center of the position bin with the greatest posterior value. We then computed the first derivative using the Python library Numpy's gradient function, which computes the difference in the forward and backward time direction and then averages them for all points except the boundaries. We then smoothed this quantity with a small Gaussian (2.5 ms standard deviation) and then averaged this for the classified time points. Because our position bin size of 3 cm makes it hard to distinguish between slower speeds in a 2 ms time step, we only analyzed classifications longer than 20 ms.

Non-Local Stationary Position

The non-local stationary position is defined as replay distances at least 30 cm from the animal's position during a time period classified as stationary.

Standard "Bayesian" Decoding and Significance Estimation

For the standard decoder analysis, we first fit a clusterless encoding model in the same manner as described in the subsection Encoding - Clusterless. We then decoded the posterior probability of each position, using 20 ms time bins, using the clusterless likelihood equation as described in the subsection The Model and assuming a uniform prior probability. Finally, to get an estimate of speed of the replay, we used either the Radon transform method, the linear regression method, or the MAP estimate method.

For the Radon transform method, we densely sampled potential line trajectories and picked the line that maximized the probability along that line. The sum of the probability along the line is called the Radon score. We picked the maximum Radon score on lines fit on the combination of center arm and left arm posterior versus center arm and right arm posteriors. P-values were computed by circularly shuffling the posterior probabilities 1000 times and comparing the real data vs. the shuffled distribution.

For the linear regression method, we drew 1000 position samples from each time bin according to the probability of position in that time bin. A linear regression was then fit to the sampled positions using time as a covariate. Fits were made separately for the posterior probabilities corresponding to the center arm with the left arm and the center arm with the right arm. Quality of the fit was assessed using the R-squared and the best fit between the arms was selected based on this metric.

For the MAP estimate, we pick the position bin with the maximum value for each time bin, which corresponds to the most probable position.

Software and Code Availability

Python code used for analysis and generating figures in the paper is available at: https://github.com/Eden-Kramer-Lab/replay_trajectory_paper. Code for the decoder is available in a separate software repository to facilitate code reuse at: https://github.com/Eden-Kramer-Lab/replay_trajectory_classification (doi: 10.5281/zenodo.3713412). Both code bases rely on the following python packages: Numpy (***van der Walt et al., 2011***), Numba (***Lam et al., 2015***), Matplotlib (***Hunter, 2007***), xarray (***Hoyer and Hamman, 2017***), NetworkX (***Hagberg et al., 2008***), Pandas (***McKinney, 2010***) and Seaborn (***Waskom, 2021***). All code is open-source and licensed under the MIT Software License. Decoder code can be easily installed as a python package with all requisite dependencies using pip or conda. See software repositories for specific details.

Acknowledgments

We thank Margaret Carr Larkin and Mattias P. Karlsson for use of their openly available datasets. We also thank Abhilasha Joshi for feedback on the manuscript. This work was supported by the

Simons Foundation for the Global Brain Grants 542971 (U.T.E.) and 542981 (L.M.F.) and the Howard Hughes Medical Institute (L.M.F.).

Interests

The authors declare that no competing interests exist.

Contributions

E.L.D. conceived of the model and did the analysis with help from U.T.E, A.K.G., and M.E.C. E.L.D and L.M.F wrote the manuscript with help from A.K.G. and M.E.C. A.K.G, M.S., J.E.C. contributed data to this manuscript.

References

Bhattarai B, Lee JW, Jung MW. Distinct Effects of Reward and Navigation History on Hippocampal Forward and Reverse Replays. Proceedings of the National Academy of Sciences. 2020 Jan; 117(1):689–697. doi: 10.1073/pnas.1912533117.

Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA. A Statistical Paradigm for Neural Spike Train Decoding Applied to Position Prediction from Ensemble Firing Patterns of Rat Hippocampal Place Cells. The Journal of Neuroscience. 1998 Sep; 18(18):7411–7425. doi: 10.1523/JNEUROSCI.18-18-07411.1998.

Carey AA, Tanaka Y, van der Meer MAA. Reward Revaluation Biases Hippocampal Replay Content Away from the Preferred Outcome. Nature Neuroscience. 2019 Aug; doi: 10.1038/s41593-019-0464-6.

Carr MF, Jadhav SP, Frank LM. Hippocampal Replay in the Awake State: A Potential Substrate for Memory Consolidation and Retrieval. Nature Neuroscience. 2011 Feb; 14(2):147–153. doi: 10.1038/nn.2732.

Carr MF, Karlsson MP, Frank LM. Transient Slow Gamma Synchrony Underlies Hippocampal Memory Replay. Neuron. 2012 Aug; 75(4):700–713. doi: 10.1016/j.neuron.2012.06.014.

Casella G, Berger RL. Statistical Inference, vol. 2 of Duxbury Advanced Series in Statistics and Decision Sciences. Duxbury Pacific Grove, CA; 2001.

Chen Z, Kloosterman F, Brown EN, Wilson MA. Uncovering Spatial Topology Represented by Rat Hippocampal Population Neuronal Codes. Journal of Computational Neuroscience. 2012 Oct; 33(2):227–255. doi: 10.1007/s10827-012-0384-x.

Chen Z, Kloosterman F, Layton S, Wilson MA. Transductive Neural Decoding for Unsorted Neuronal Spikes of Rat Hippocampus. In: *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE* IEEE; 2012. p. 1310–1313. doi: 10.1109/EMBC.2012.6346178.

Chen Z, Linderman SW, Wilson MA. Bayesian Nonparametric Methods for Discovering Latent Structures of Rat Hippocampal Ensemble Spikes. In: *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)* Vietri sul Mare, Salerno, Italy: IEEE; 2016. p. 1–6. doi: 10.1109/MLSP.2016.7738867.

Chung JE, Magland JF, Barnett AH, Tolosa VM, Tooker AC, Lee KY, Shah KG, Felix SH, Frank LM, Greengard LF. A Fully Automated Approach to Spike Sorting. Neuron. 2017 Sep; 95(6):1381–1394.e6. doi: 10.1016/j.neuron.2017.08.030.

Davidson TJ, Kloosterman F, Wilson MA. Hippocampal Replay of Extended Experience. Neuron. 2009 Aug; 63(4):497–507. doi: 10.1016/j.neuron.2009.07.027.

Deng X, Liu DF, Karlsson MP, Frank LM, Eden UT. Rapid Classification of Hippocampal Replay Content for Real-Time Applications. Journal of Neurophysiology. 2016 Nov; 116(5):2221–2235. doi: 10.1152/jn.00151.2016.

Deng X, Liu DF, Kay K, Frank LM, Eden UT. Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter. Neural Computation. 2015 Jul; 27(7):1438–1460. doi: 10.1162/NECO_a_00744.

Denovellis EL, Frank LM, Eden UT. Characterizing Hippocampal Replay Using Hybrid Point Process State Space Models. In: *2019 53rd Asilomar Conference on Signals, Systems, and Computers* Pacific Grove, CA, USA: IEEE; 2019. p. 245–249. doi: 10.1109/IEEECONF44664.2019.9048688.

Diba K, Buzsáki G. Forward and Reverse Hippocampal Place-Cell Sequences during Ripples. Nature Neuroscience. 2007 Oct; 10(10):1241–1242. doi: 10.1038/nn1961.

Dijkstra EW. A Note on Two Problems in Connexion with Graphs. Numerische mathematik. 1959; 1(1):269–271. doi: 10.1007/BF01386390.

Drieu C, Todorova R, Zugaro M. Nested Sequences of Hippocampal Assemblies during Behavior Support Subsequent Sleep Replay. Science. 2018 Nov; 362(6415):675. doi: 10.1126/science.aat2952.

Eichenbaum H, Cohen NJ. From Conditioning to Conscious Recollection: Memory Systems of the Brain. No. 35 in Oxford Psychology Series, New York: Oxford University Press; 2004. doi: 10.1093/acprof:oso/9780195178043.001.0001.

Farooq U, Dragoi G. Emergence of Preconfigured and Plastic Time-Compressed Sequences in Early Postnatal Development. Science. 2019 Jan; 363(6423):168–173. doi: 10.1126/science.aav0502.

Foster DJ, Wilson MA. Reverse Replay of Behavioural Sequences in Hippocampal Place Cells during the Awake State. Nature. 2006 Mar; 440(7084):680–683. doi: 10.1038/nature04587.

Fox SE, Ranck Jr JB. Localization and Anatomical Identification of Theta and Complex Spike Cells in Dorsal Hippocampal Formation of Rats. Experimental Neurology. 1975; 49(1):299–313. doi: 10.1016/0014-4886(75)90213-7.

Frankland PW, Bontempi B. The Organization of Recent and Remote Memories. Nature Reviews Neuroscience. 2005 Feb; 6(2):119–130. doi: 10.1038/nrn1607.

Grosmark AD, Buzsáki G. Diversity in Neural Firing Dynamics Supports Both Rigid and Learned Hippocampal Sequences. Science. 2016 Mar; 351(6280):1440–1443. doi: 10.1126/science.aad1935.

Gupta AS, van der Meer MAA, Touretzky DS, Redish AD. Hippocampal Replay Is Not a Simple Function of Experience. Neuron. 2010 Mar; 65(5):695–705. doi: 10.1016/j.neuron.2010.01.034.

Hagberg AA, Schult DA, Swart PJ. Exploring Network Structure, Dynamics, and Function Using NetworkX. In: *Proceedings of the 7th Python in Science Conference* Pasadena, CA USA; 2008. p. 11–15.

Hangya B, Li Y, Muller RU, Czurkó A. Complementary Spatial Firing in Place Cell-Interneuron Pairs: Place Cell-Interneuron Spatial Correlation. The Journal of Physiology. 2010 Nov; 588(21):4165–4175. doi: 10.1113/jphysiol.2010.194274.

Harris KD, Hirase H, Leinekugel X, Henze DA, Buzsáki G. Temporal Interaction between Single Spikes and Complex Spike Bursts in Hippocampal Pyramidal Cells. Neuron. 2001 Oct; 32(1):141–149. doi: 10.1016/S0896-6273(01)00447-0.

Hoyer S, Hamman JJ. Xarray: N-D Labeled Arrays and Datasets in Python. Journal of Open Research Software. 2017 Apr; 5:10. doi: 10.5334/jors.148.

Hunter JD. Matplotlib: A 2D Graphics Environment. Computing in Science & Engineering. 2007; 9(3):90–95. doi: 10.1109/MCSE.2007.55.

Jadhav SP, Rothschild G, Roumis DK, Frank LM. Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. Neuron. 2016 Apr; 90(1):113–127. doi: 10.1016/j.neuron.2016.02.010.

Jinno S, Klausberger T, Marton LF, Dalezios Y, Roberts JDB, Fuentealba P, Bushong EA, Henze D, Buzsaki G, Somogyi P. Neuronal Diversity in GABAergic Long-Range Projections from the Hippocampus. Journal of Neuroscience. 2007 Aug; 27(33):8790–8804. doi: 10.1523/JNEUROSCI.1847-07.2007.

Johnson A, Jackson JC, Redish AD. Measuring Distributed Properties of Neural Representations beyond the Decoding of Local Variables: Implications for Cognition. In: Holscher C, Munk M, editors. *Information Processing by Neuronal Populations* Cambridge: Cambridge University Press; 2008.p. 95–119. doi: 10.1017/CBO9780511541650.005.

Joo HR, Frank LM. The Hippocampal Sharp Wave–Ripple in Memory Retrieval for Immediate Use and Consolidation. Nature Reviews Neuroscience. 2018 Dec; 19(12):744–757. doi: 10.1038/s41583-018-0077-1.

Kaefer K, Nardin M, Blahna K, Csicsvari J. Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching. Neuron. 2020 Feb; p. S0896627320300416. doi: 10.1016/j.neuron.2020.01.015.

Karlsson MP, Frank LM. Awake Replay of Remote Experiences in the Hippocampus. Nature Neuroscience. 2009 Jul; 12(7):913–918. doi: 10.1038/nn.2344.

Kay K, Chung JE, Sosa M, Schor JS, Karlsson MP, Larkin MC, Liu DF, Frank LM. Constant Sub-Second Cycling between Representations of Possible Futures in the Hippocampus. Cell. 2020 Jan; p. S0092867420300611. doi: 10.1016/j.cell.2020.01.014.

Kay K, Sosa M, Chung JE, Karlsson MP, Larkin MC, Frank LM. A Hippocampal Network for Spatial Coding during Immobility and Sleep. Nature. 2016 Mar; 531(7593):185–190. doi: 10.1038/nature17144.

Kloosterman F, Layton SP, Chen Z, Wilson MA. Bayesian Decoding Using Unsorted Spikes in the Rat Hippocampus. Journal of Neurophysiology. 2014 Jan; 111(1):217–227. doi: 10.1152/jn.01046.2012.

Lam SK, Pitrou A, Seibert S. Numba: A LLVM-Based Python JIT Compiler. In: *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC - LLVM '15* Austin, Texas: ACM Press; 2015. p. 1–6. doi: 10.1145/2833157.2833162.

Lee AK, Wilson MA. Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep. Neuron. 2002 Dec; 36(6):1183–1194. doi: 10.1016/S0896-6273(02)01096-6.

Lex A, Gehlenborg N, Strobelt H, Vuillemot R, Pfister H. UpSet: Visualization of Intersecting Sets. IEEE Transactions on Visualization and Computer Graphics. 2014 Dec; 20(12):1983–1992. doi: 10.1109/TVCG.2014.2346248.

Linderman SW, Johnson MJ, Wilson MA, Chen Z. A Bayesian Nonparametric Approach for Uncovering Rat Hippocampal Population Codes during Spatial Navigation. Journal of Neuroscience Methods. 2016 Apr; 263:36–47. doi: 10.1016/j.jneumeth.2016.01.022.

Liu S, Grosmark AD, Chen Z. Methods for Assessment of Memory Reactivation. Neural Computation. 2018 Aug; 30(8):2175–2209. doi: 10.1162/neco_a_01090.

Maboudi K, Ackermann E, de Jong LW, Pfeiffer BE, Foster D, Diba K, Kemere C. Uncovering Temporal Structure in Hippocampal Output Patterns. eLife. 2018; p. 24. doi: 10.7554/eLife.34467.

McKinney W. Data Structures for Statistical Computing in Python. In: *Python in Science Conference* Austin, Texas; 2010. p. 56–61. doi: 10.25080/Majora-92bf1922-00a.

Michon F, Sun JJ, Kim CY, Ciliberti D, Kloosterman F. Post-Learning Hippocampal Replay Selectively Reinforces Spatial Memory for Highly Rewarded Locations. Current Biology. 2019 Apr; p. S096098221930346X. doi: 10.1016/j.cub.2019.03.048.

Nádasdy Z, Hirase H, Czurkó A, Csicsvari J, Buzsáki G. Replay and Time Compression of Recurring Spike Sequences in the Hippocampus. The Journal of Neuroscience. 1999 Nov; 19(21):9497–9507. doi: 10.1523/JNEUROSCI.19-21-09497.1999.

Newson P, Krumm J. Hidden Markov Map Matching through Noise and Sparseness. In: *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* GIS '09, Association for Computing Machinery; 2009. p. 336–343. doi: 10.1145/1653771.1653818.

Ólafsdóttir HF, Carpenter F, Barry C. Task Demands Predict a Dynamic Switch in the Content of Awake Hippocampal Replay. Neuron. 2017 Nov; 96(4):925–935.e6. doi: 10.1016/j.neuron.2017.09.035.

Oliva A, Fernández-Ruiz A, Fermino de Oliveira E, Buzsáki G. Origin of Gamma Frequency Power during Hippocampal Sharp-Wave Ripples. Cell Reports. 2018 Nov; 25(7):1693–1700.e4. doi: 10.1016/j.celrep.2018.10.066.

Pfeiffer BE, Foster DJ. Hippocampal Place-Cell Sequences Depict Future Paths to Remembered Goals. Nature. 2013 Apr; 497(7447):74–79. doi: 10.1038/nature12112.

Pfeiffer BE, Foster DJ. Autoassociative Dynamics in the Generation of Sequences of Hippocampal Place Cells. Science. 2015; 349(6244):180–183. doi: 10.1126/science.aaa9633.

Ranck Jr JB. Studies on Single Neurons in Dorsal Hippocampal Formation and Septum in Unrestrained Rats: Part I. Behavioral Correlates and Firing Repertoires. Experimental neurology. 1973; 41(2):462–531.

SciPy 1 0 Contributors, Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. Nature Methods. 2020 Mar; 17(3):261–272. doi: 10.1038/s41592-019-0686-2.

Shin JD, Tang W, Jadhav SP. Dynamics of Awake Hippocampal-Prefrontal Replay for Spatial Learning and Memory-Guided Decision Making. Neuron. 2019 Oct; p. S0896627319307858. doi: 10.1016/j.neuron.2019.09.012.

Stella F, Baracskay P, O'Neill J, Csicsvari J. Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. Neuron. 2019 Feb; doi: 10.1016/j.neuron.2019.01.052.

Tang W, Shin JD, Frank LM, Jadhav SP. Hippocampal-Prefrontal Reactivation during Learning Is Stronger in Awake Compared with Sleep States. The Journal of Neuroscience. 2017 Dec; 37(49):11789–11805. doi: 10.1523/JNEUROSCI.2291-17.2017.

Tingley D, Buzsáki G. Routing of Hippocampal Ripples to Subcortical Structures via the Lateral Septum. Neuron. 2020 Jan; 105(1):138–149.e5. doi: 10.1016/j.neuron.2019.10.012.

Tingley D, Peyrache A. On the Methods for Reactivation and Replay Analysis. Philosophical Transactions of the Royal Society B: Biological Sciences. 2020 May; 375(1799):20190231. doi: 10.1098/rstb.2019.0231.

Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, O'Shea DJ, Vyas S, Sun X, Ryu SI, Ganguli S, Shenoy KV. Accurate Estimation of Neural Population Dynamics without Spike Sorting. Neuron. 2019 Jul; 103(2):292–308.e4. doi: 10.1016/j.neuron.2019.05.003.

Tropp Sneider J, Chrobak JJ, Quirk MC, Oler JA, Markus EJ. Differential Behavioral State-Dependence in the Burst Properties of CA3 and CA1 Neurons. Neuroscience. 2006; 141(4):1665–1677. doi: 10.1016/j.neuroscience.2006.05.052.

van der Walt S, Colbert SC, Varoquaux G. The NumPy Array: A Structure for Efficient Numerical Computation. Computing in Science & Engineering. 2011 Mar; 13(2):22–30. doi: 10.1109/MCSE.2011.37.

Viterbi A. Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm. IEEE Transactions on Information Theory. 1967 Apr; 13(2):260–269. doi: 10.1109/TIT.1967.1054010.

Wang M, Foster DJ, Pfeiffer BE. Alternating Sequences of Future and Past Behavior Encoded within Hippocampal Theta Oscillations. Science. 2020 Oct; 370(6513):247. doi: 10.1126/science.abb4151.

Waskom M. Seaborn: Statistical Data Visualization. Journal of Open Source Software. 2021 Apr; 6(60):3021. doi: 10.21105/joss.03021.

Wilent WB, Nitz DA. Discrete Place Fields of Hippocampal Formation Interneurons. Journal of Neurophysiology. 2007 Jun; 97(6):4152–4161. doi: 10.1152/jn.01200.2006.

Wu X, Foster DJ. Hippocampal Replay Captures the Unique Topological Structure of a Novel Environment. Journal of Neuroscience. 2014 May; 34(19):6459–6469. doi: 10.1523/JNEUROSCI.3414-13.2014.

Yamamoto J, Tonegawa S. Direct Medial Entorhinal Cortex Input to Hippocampal CA1 Is Crucial for Extended Quiet Awake Replay. Neuron. 2017 Sep; 96(1):217–227.e4. doi: 10.1016/j.neuron.2017.09.017.

Yu JY, Frank LM. Hippocampal–Cortical Interaction in Decision Making. Neurobiology of Learning and Memory. 2015 Jan; 117:34–41. doi: 10.1016/j.nlm.2014.02.002.

Yu JY, Kay K, Liu DF, Grossrubatscher I, Loback A, Sosa M, Chung JE, Karlsson MP, Larkin MC, Frank LM. Distinct Hippocampal-Cortical Memory Representations for Experiences Associated with Movement versus Immobility. eLife. 2017; 6. doi: 10.7554/eLife.27621.

Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ. Interpreting Neuronal Population Activity by Reconstruction: Unified Framework With Application to Hippocampal Place Cells. Journal of Neurophysiology. 1998 Feb; 79(2):1017–1044. doi: 10.1152/jn.1998.79.2.1017.

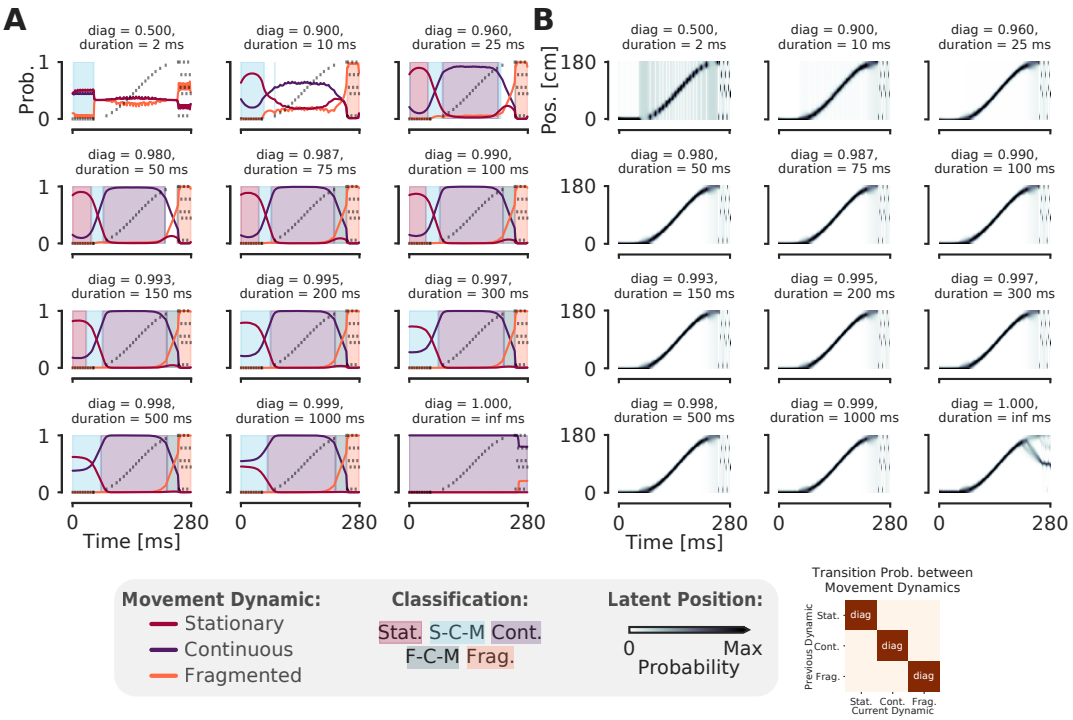

Figure 1–Figure supplement 1. The model is robust to change of the probability of persisting in the same dynamic for a wide range of plausible expected durations (25 ms to 150 ms). **(A)** Each panel shows the probability of each dynamic on simulated data example from Figure 1 with a different diagonal value—which governs the probability of remaining in that dynamic. The corresponding expected duration of staying in the dynamic is listed as duration. The off-diagonal values—the probability of switching to one of the other dynamics—are set to be equally likely with the remainder of the probability, as in Figure 1D. The diagonal increases from left to right, top to bottom, until the case where the diagonal is one and the off-diagonal is zero—i.e. the case where there is no probability of switching to another dynamic. Shaded regions correspond to the classification as in Figure 1F. **(B)** The probability of position over time for each diagonal value. Conventions the same as in (A).

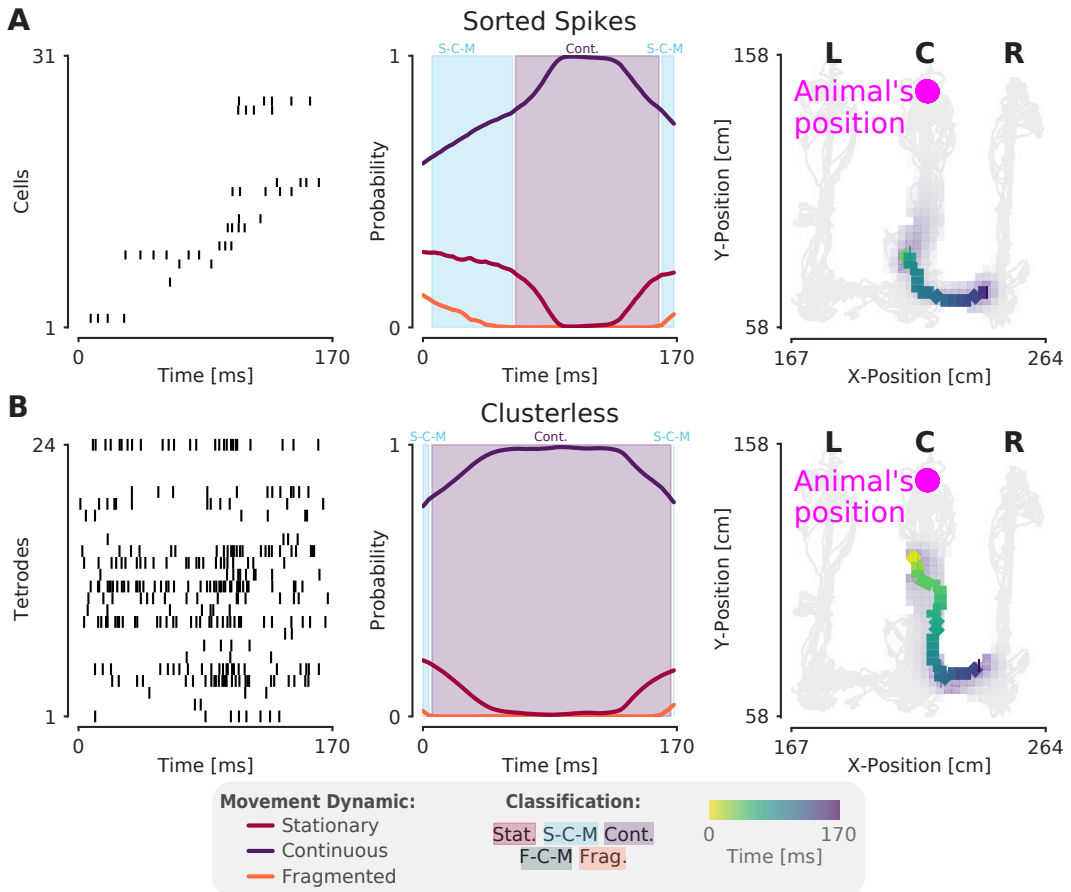

Figure 2–Figure supplement 1. Decoding the same SWR in **Figure 2** with 2D position using sorted spikes and clusterless decoding. **(A)** The left panel shows the spikes from cells arranged by the linear position of the peak of place field as in **Figure 2**. The middle panel shows the probability of each dynamic over time from the 2D decode. Shaded regions indicate classification category as in **Figure 1**F and **Figure 2**. The rightmost panel shows the most probable estimate of the latent position (MAP estimate) with color indicating time. The latent position posterior summed over time is shown in the purple shading. The light grey lines represent the position of the animal over the entire recording session and the magenta dot represents the animal's position. **(B)** Same as in (A), but with clusterless decoding.

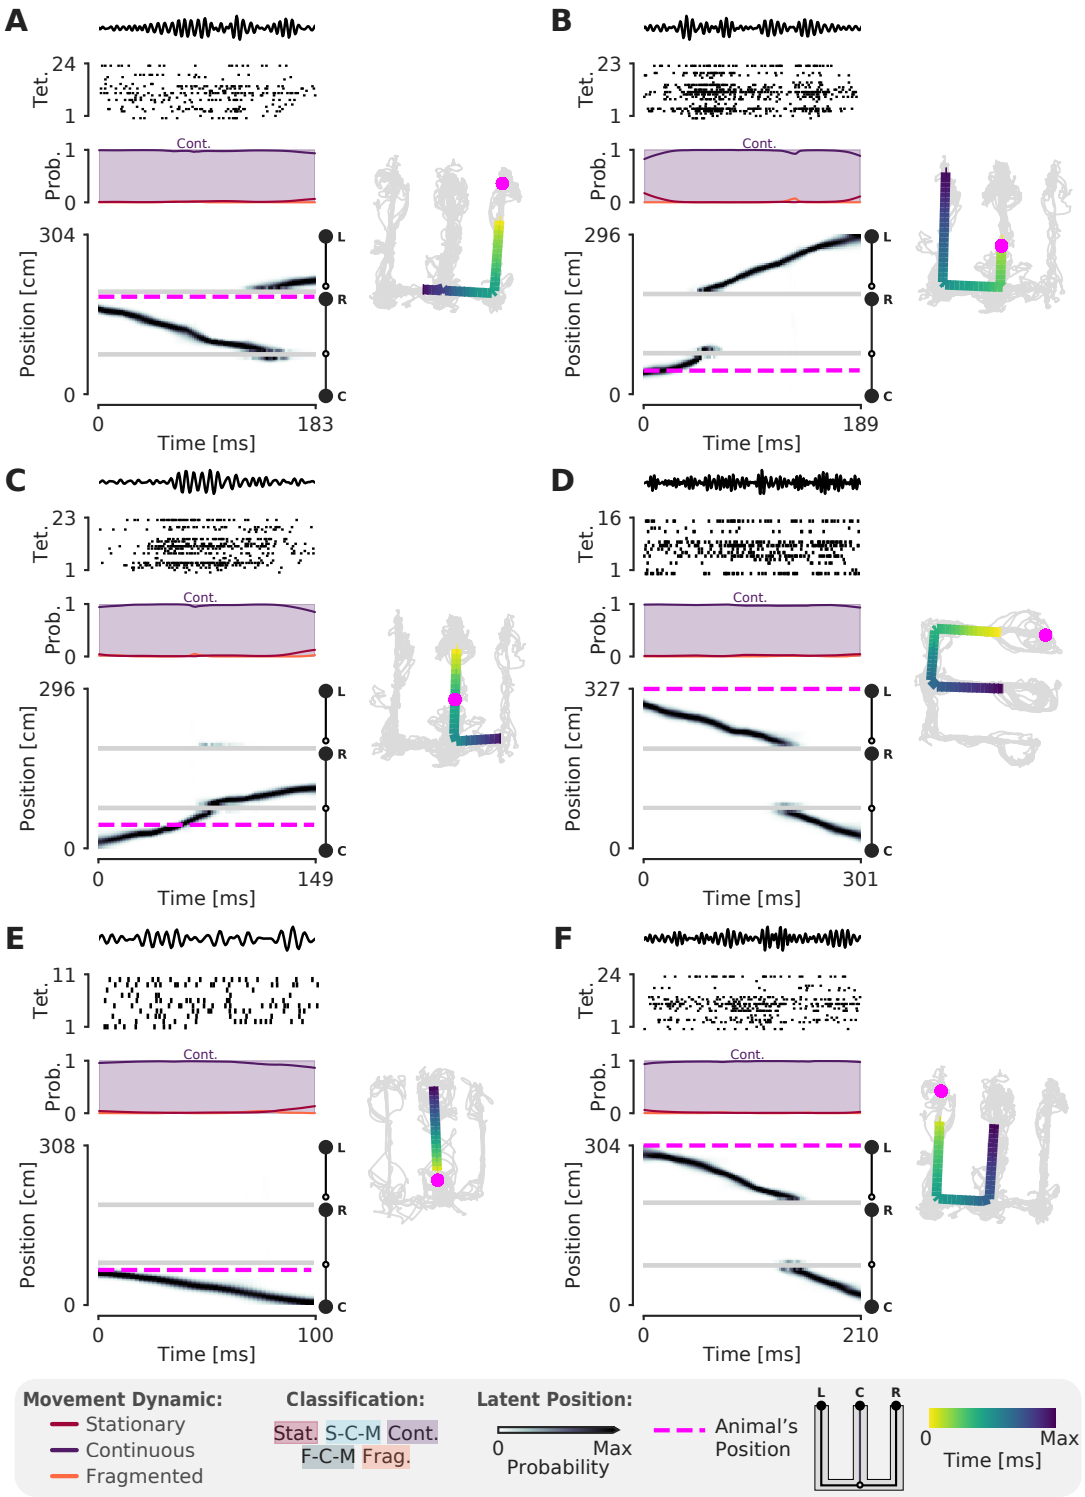

Figure 2–Figure supplement 2. A-F. More examples of SWRs that have continuous trajectories. Left panel uses the same conventions as **Figure 2**A and **Figure 2**B. Right panel shows the 1D MAP estimate projected back to 2D as in **Figure 2**C. Color indicates time. Light grey lines indicate the animal's position over the entire recording session. Magenta dashed line represents the animal's position during the SWR.

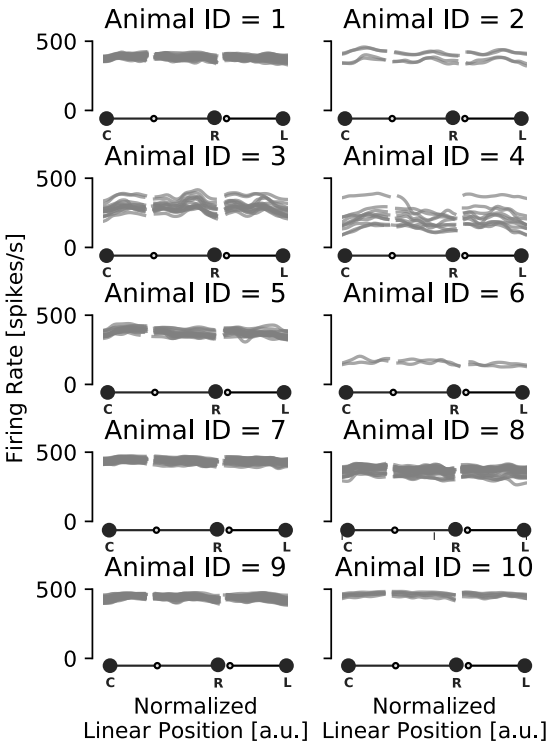

Figure 2–Figure supplement 3. Multiunit rate of encoding spikes over track positions for each animal. Each grey line represents a recording session for that animal.

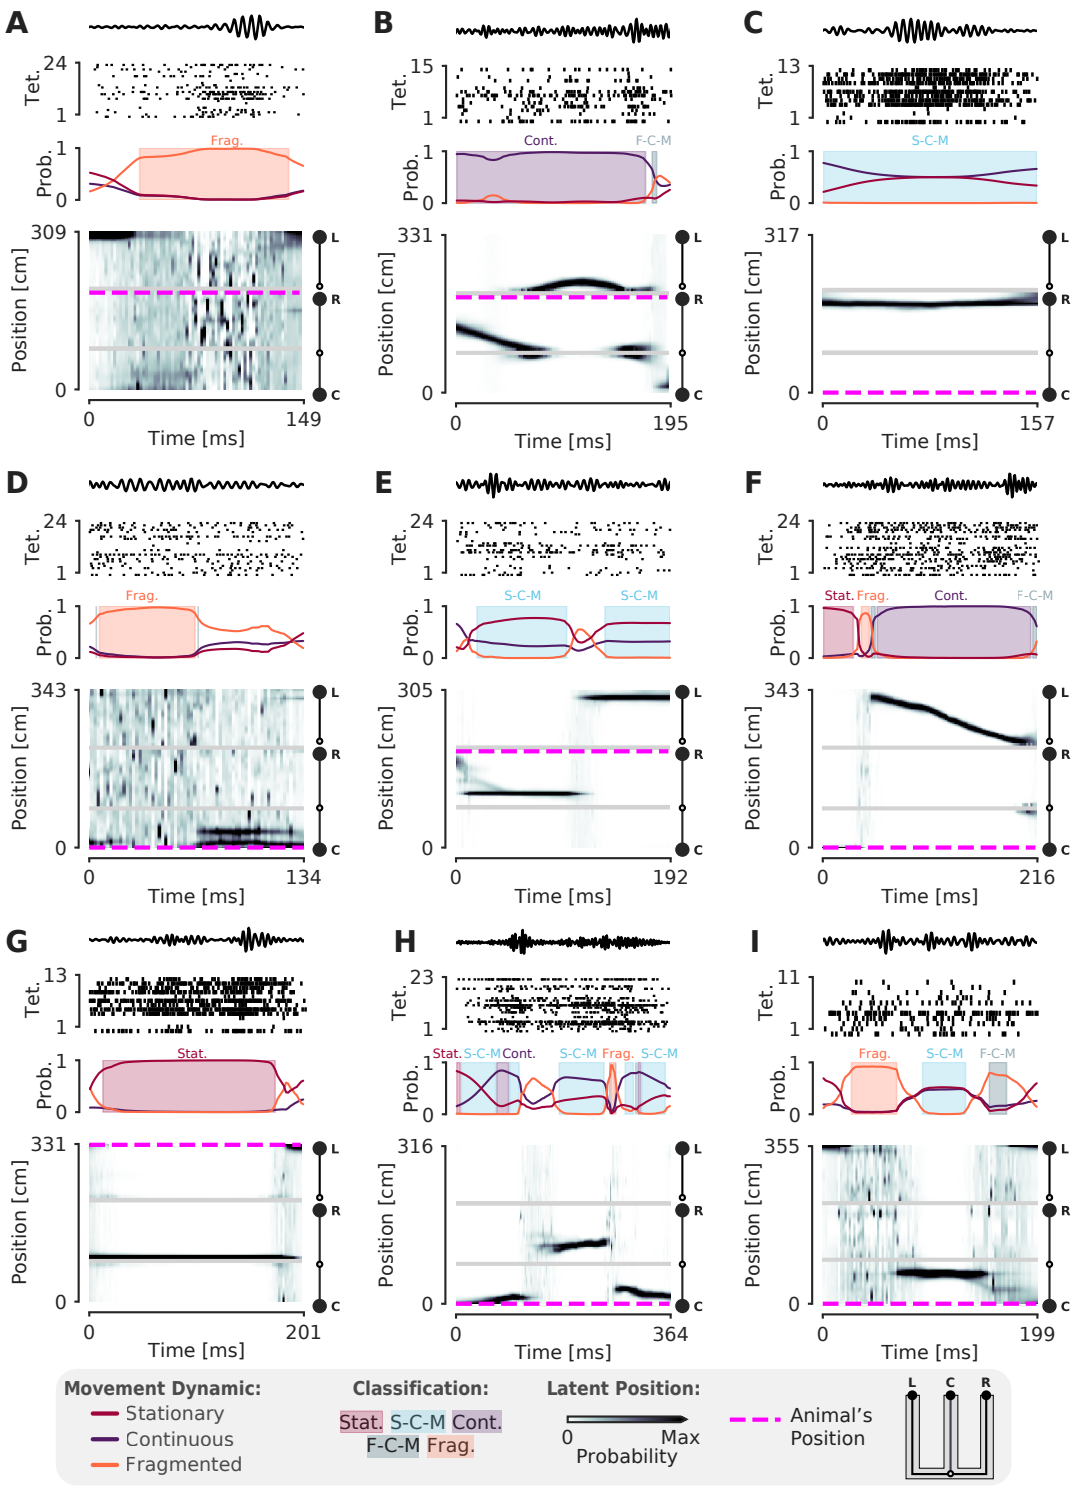

Figure 3–Figure supplement 1. A-I. More examples of SWRs with non-constant speed trajectories. Conventions are the same as in Figure 3.

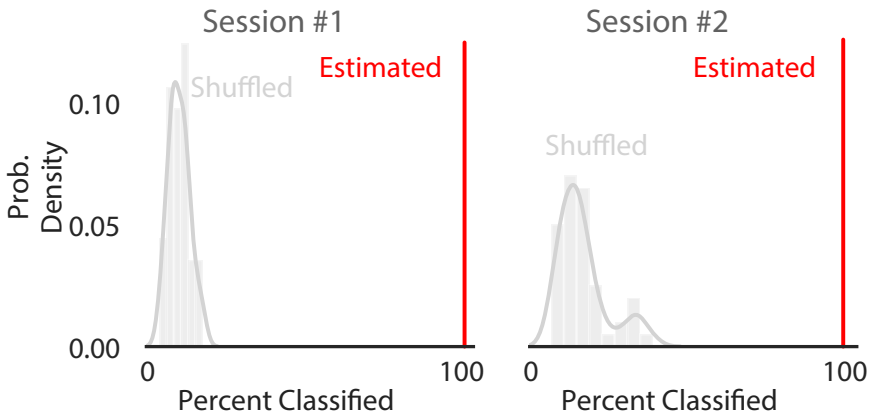

Figure 3–Figure supplement 2. Comparison of percentage of SWRs classified (that is, a SWR containing at least one of the five classifications) on real vs. position shuffled data for two recording sessions from different animals. Red line represents the percent of SWR events classified in that recording session for real data. The histogram represents the distribution after 50 shuffles of the position data. Position data was shuffled by resampling with replacement from the set of all observed positions in that recording session, destroying position information but preserving spiking timing and position occupancy.

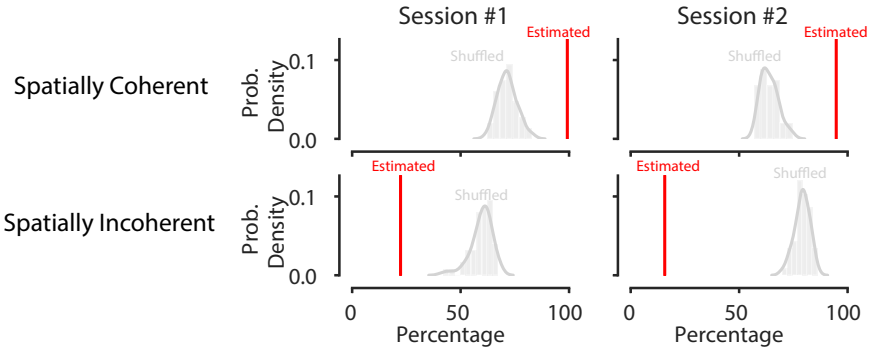

Figure 3–Figure supplement 3. Shuffling the data by swapping the runs and circularly permuting the position increases the percentage of spatially incoherent SWRs and decreases the spatially coherent SWRs. The red line represents the percent of spatially coherent or incoherent SWRs in that recording session for actual data. The histogram represents the distribution after 50 shuffles of the run from well to well as well as the circularly shuffled position (for each tetrode). This preserves local spatial correlations in the data but breaks the global spatial relationships between the spikes and the data.

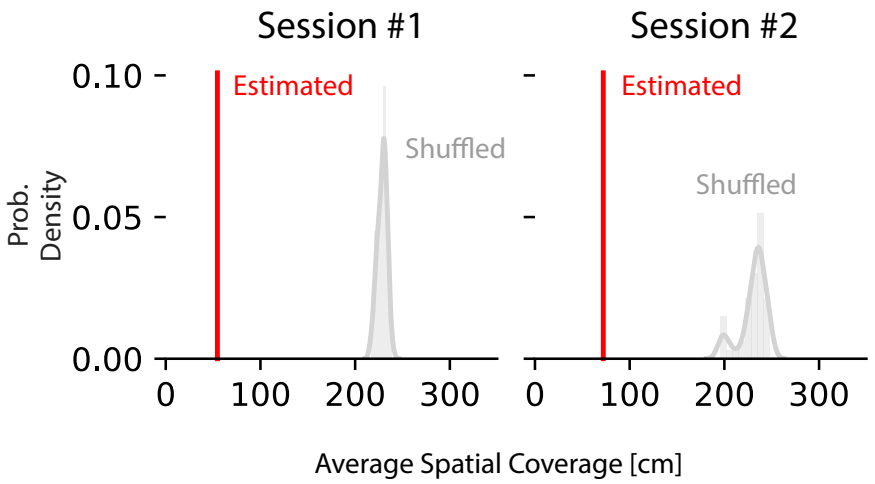

Figure 4–Figure supplement 1. Average 95% HPD region size increases on position shuffled data for two recording sessions from different animals. Red line represents the average 95% HPD region size for all ripples. The histogram represents the distribution after 50 shuffles of the position data. Position data was shuffled by resampling with replacement from the set of all observed positions in that recording session, destroying position information but preserving spiking timing and position occupancy.

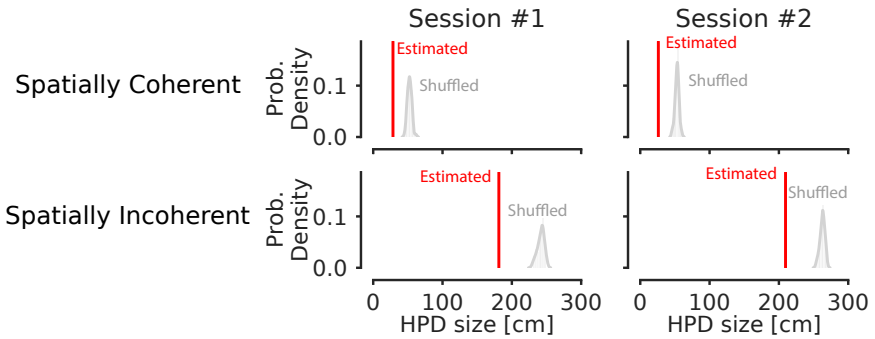

Figure 4–Figure supplement 2. Shuffling the position data by swapping the runs and circularly permuting the data increases the average 95% HPD region size for spatially coherent and incoherent classified times. Data from two recording sessions from different animals. Red line represents the average spatial position spanned by the highest posterior density on real data for all ripples. The histogram represents the distribution after 50 shuffles of the position data. The histogram represents the distribution after 50 shuffles of the run from well to well as well as the circularly shuffled position (for each tetrode). This preserves local correlations between spikes but reduces global spatial structure.

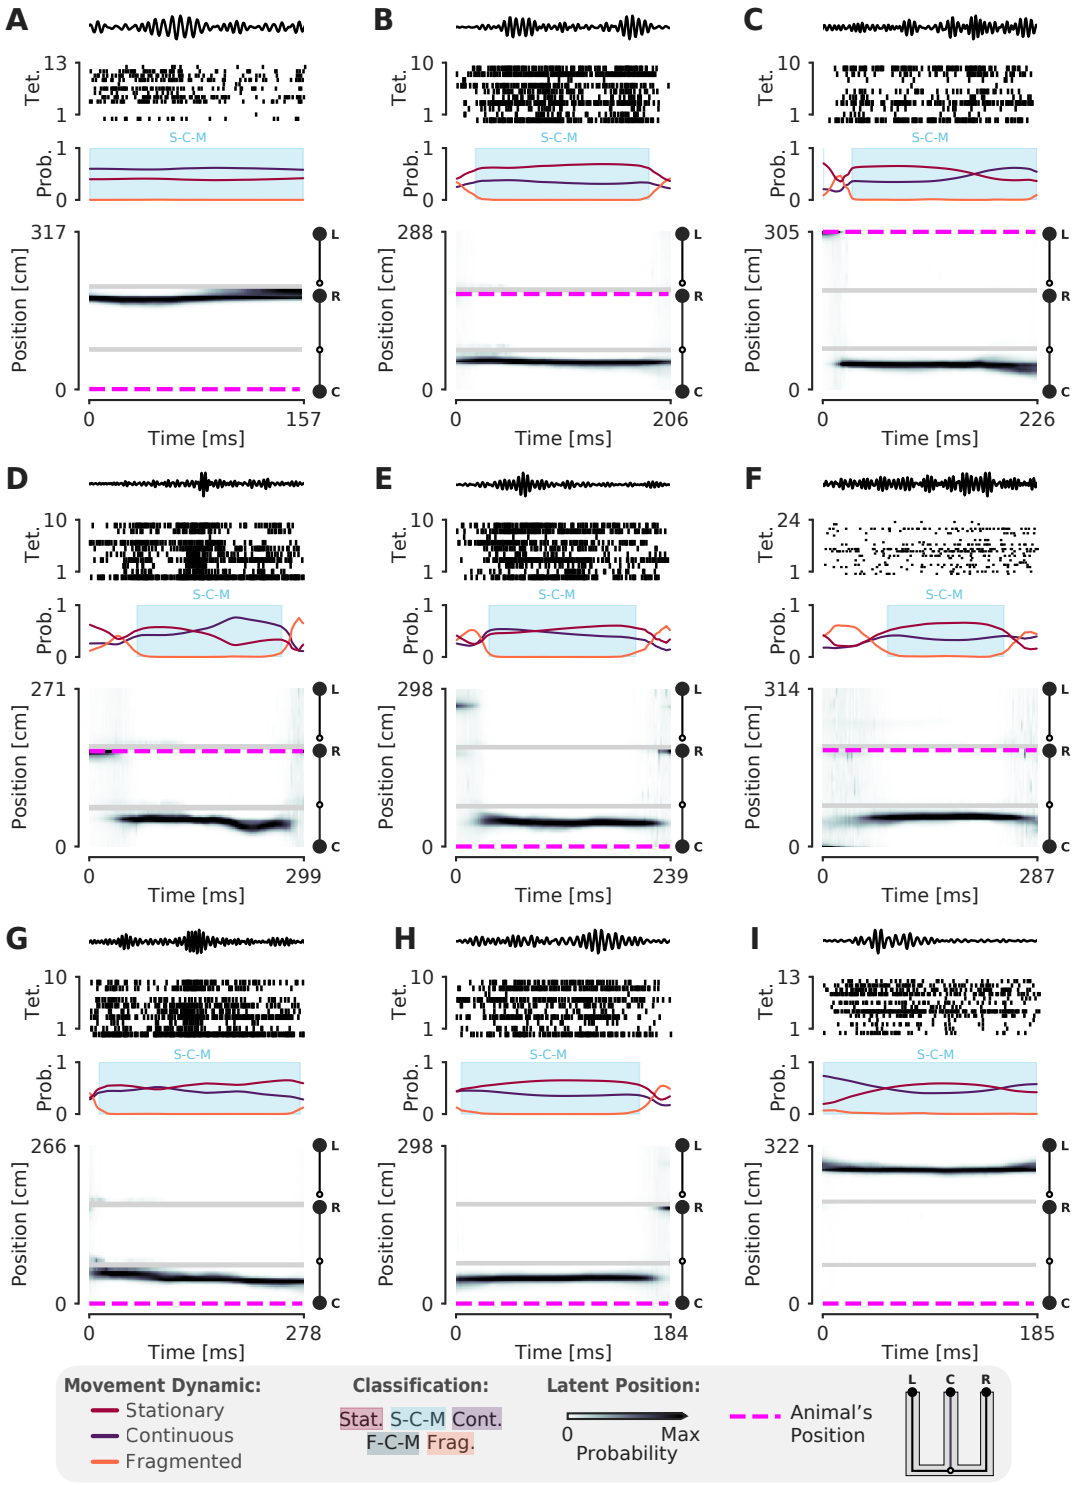

Figure 5–Figure supplement 1. More examples of stationary-continuous-mixtures. Conventions are the same as in **Figure 3**.

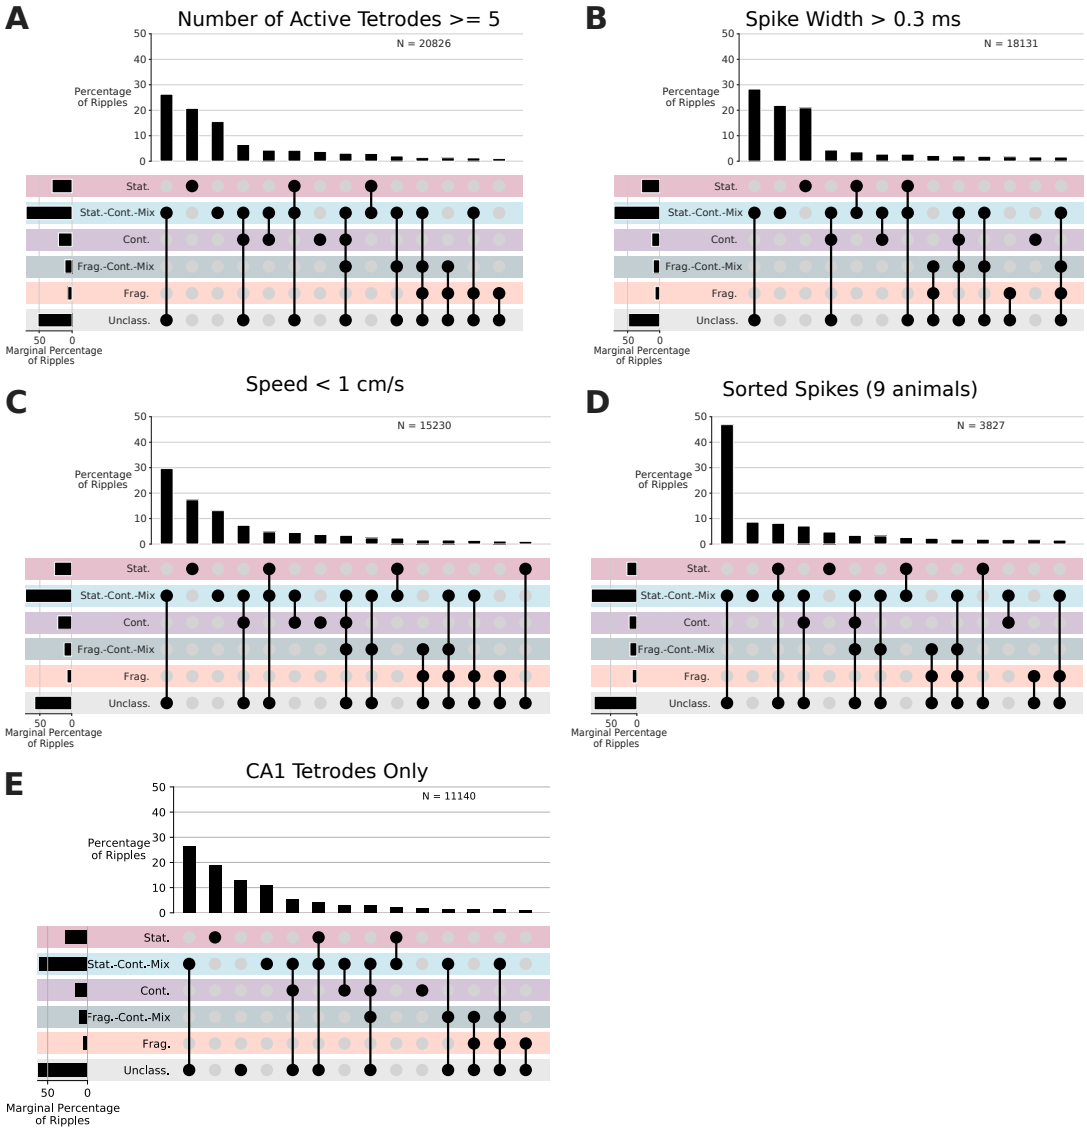

Figure 5–Figure supplement 2. Control analyses for distribution of dynamics. Conventions are the same as in **Figure 5**.

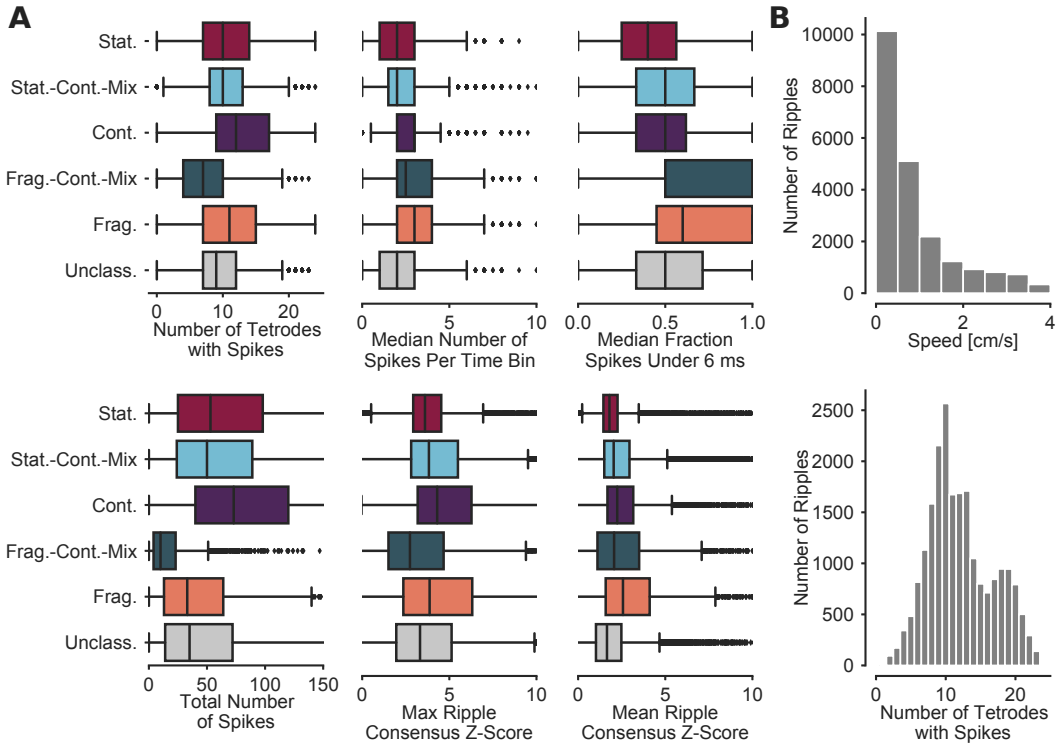

Figure 5–Figure supplement 3. **(A)** Further quantification of the dynamics with respect to the spiking and ripple properties. **(B)** Speed and number of tetrodes with spikes for entire SWR