Hippocampal replay of experience at real-world speeds

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

[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

Representations of past and possible future experiences play a critical role in memory and decision-making processes. The 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 real-world speeds. 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, most events progress at real-world, rather than accelerated, movement speeds, consistent with actual experiences.

Hippocampus | Replay | State Space | Dynamics

Correspondence: ***loren@phy.ucsf.edu***

Introduction

The brain has the remarkable ability to retrieve representations of past events and generate representations of possible future events. These "generative" (26) mental events may not occur in the same time span as actual events: memories can be retrieved and used in less time than was required for the original experience, and mental simulations can span a series of actions more quickly than would be required to perform those actions. Thus, a potential neural substrate for retrieval and mental simulation would be expected to support rapid reinstatement of patterns of brain activity associated with past or potential future experience.

Hippocampal "replay" events are well suited to that role (2). As animals move through space, neurons in the hippocampus preferentially fire at specific locations in an environment, and thus sets 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 are thus 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 (10, 25, 30, 36).

While the existence of these sequential events is well established, the current consensus is that only a minority (~10-45%) of hippocampal SWRs contain statistically identifiable, sequential replay (10, 24, 35, 40). One might therefore conclude that the retrieval of stored representation during SWRs is limited to sequences of locations, and that all such retrieval events are sequential in nature. Our subjective experience is that we can retrieve memories of locations without having to mentally traverse a long distance. If the rodent hippocampus is capable of that sort of retrieval, then one might expect to see SWRs where the neural activity corresponds not to a very rapid trajectory through space but instead to either a stable pattern associated with a single location or perhaps a pattern more similar to one seen during a real experience. Importantly, this sort of activity could still be useful, as it would take place during the short duration of an SWR (~100 ms) and could thus support rapid retrieval and use of representations.

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 (22), but another report suggested that events that replay a single location are only seen in young animals (41). Thus, while it would seem useful to replay memories associated with single locations and perhaps even experiences occurring at real-world speeds, 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 multiple 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 (13, 16, 25, 41). 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 approach for detecting sequential replay involves fitting a line to the resulting set of probability distributions, which relies on the assumption that the representations progress at a constant, non-zero speed (13, 16, 25). A statistical test is then used to determine whether the fit is better than the fit to shuffled versions of the data.

While the standard decoder approach identifies constant speed events, it does not consider events that are rejected by the statistical test, and the use of a fixed size temporal bins acts as a boxcar smoother that limits the potential movement speeds of the representation. The linear fit is also problematic: it has the potential to reject or mischaracterize real events that do not move at constant speeds. Approaches that focus on the order of cell activity within each event (17, 30) 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 exclude events that have stationary representations of a single location (15, 22).

Recognizing the problems with the linear fit, several studies have moved away from the constant velocity assumption, using an estimated represented location at each time bin and connecting each with a line. For example, using this approach, Pfeiffer and Foster (38) found that awake replays can alternate between representing a single location and sequential spatial trajectories. On the other hand, Stella et al. (41) reported that replays during sleep are spatially continuous and follow Brownian diffusion dynamics. Both methods still used large time bins and neither took into account the uncertainty of the decoded estimates, however, making it hard to identify the source of the different conclusions.

An ideal approach to identifying and quantifying the dynamics of generative representations would circumvent these problems. It would use all of the available spiking data to yield the most accurate decoded positions. It would be applicable to either thresholded events or to all of the data to permit an unbiased assessment of the prevalence and nature of generative activity. It would use very small temporal bins (1 or 2 ms) to allow for very rapid representational movement and 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 can be explored systematically to understand their influence on the results.

We therefore developed a state space model of generative activity 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 (11, 33) by characterizing represented trajectories 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 much slower speeds, more consistent with real-world experiences, suggesting that high speed replay is not the most common manifestation of generative activity. 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, which is spaced every 10 cm along the virtual track. We then construct a 280 ms spiking sequence (Figure 1A) 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 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")

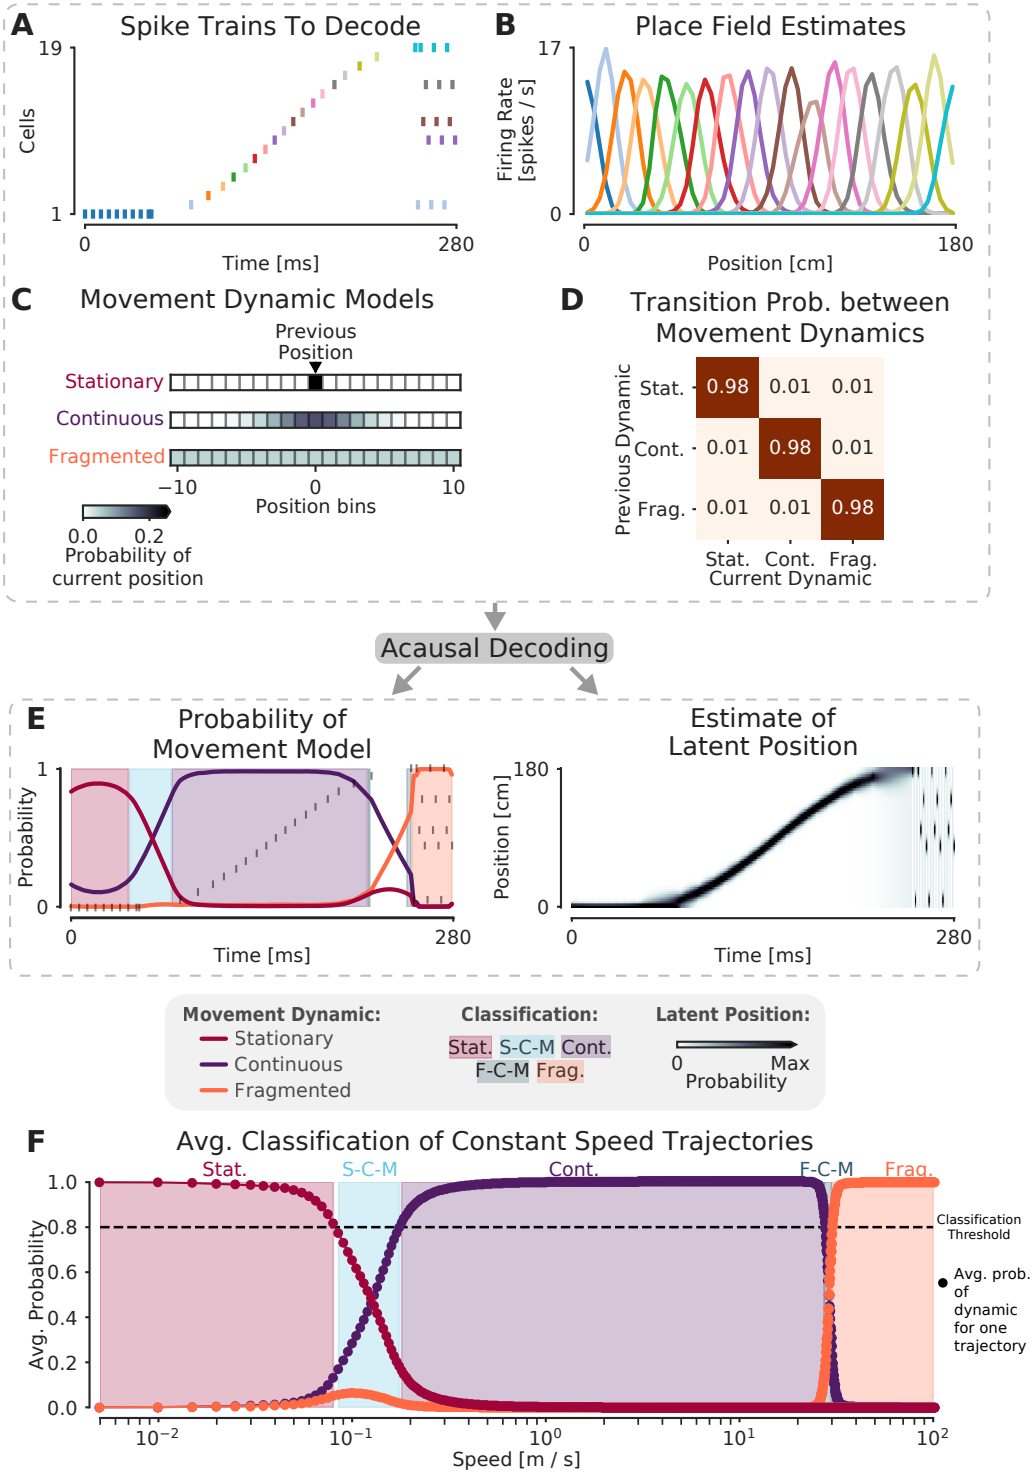

Figure 1. The model can capture different sequence dynamics on simulated data. **(A)** We construct a firing sequence of 19 simulated place cells that exhibits three different movement dynamics. For the first 60 ms, one cell fires repeatedly, representing one stationary location. For next 190 ms, the cells fire in sequence, representing a rapid continuous trajectory along the virtual track. Finally, 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, which is determined from the probability. Marginalizing the posterior across dynamics also provides an estimate of latent position over time (right panel). **(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 0.8 threshold (dotted line) to classify each sequence based on the dynamic or dynamics which contribute maximally to the posterior (shaded colors).

decoder (10, 38, 41). We compute an estimate of how each cell's firing rate varies over position during movement (i.e. the place field, Figure 1B). 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.

Next, we specify movement dynamic models for how latent position—the "mental" position of the animal represented by the cells—evolves over time for a given dynamic (Figure 1C). 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 (15, 22), continuous (10), 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 (38). We therefore define state transition models to capture each of these movement dynamics.

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 the next position will be the same as the last position (subject to 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 (10, 38), 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). Finally, 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.

Finally, we specify how likely each movement dynamic is to persist in time versus change to another dynamic (Figure 1D) via a state transition matrix between dynamics. In order to be conservative, we assume that each movement dynamic is likely to dominate for ~100 ms, 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 a particular movement dynamic is high, data providing clear evidence for a change in dynamic can, within the time frame of ~10 ms, drive a change to a different estimated dynamic, as in our simulated example. In line with this, note that below we show 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) and the probability of latent position over time, irrespective of movement dynamic (by summing over the movement dynamics; Figure 1E, left and right plot respectively).

An examination of the two summaries shown in Figure 1E reveals that that the model successfully captures the dynamics of the population spiking activity in Figure 1A. 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.

Importantly, this approach allows the model to to capture a wide range of movement speeds for the latent position. 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 movement 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 1F) from 1 cm/s to 10,000 m/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 0.8 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 0.8. 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 1F) 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 sum-

mary, we can characterize the speed or set of speeds in 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 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 0.8 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 1A with different probabilities of staying in the same dynamic versus switching to another dynamic (Figure S1). We found that for a large range of plausible probabilities of staying in one of the dynamics (between 0.96 and 0.993 which corresponds to an expected duration between 25 and 150 ms), the model still correctly identified the sequence dynamics with high probability (Figure S1A), demonstrating that data itself is the most influential element in the model. Furthermore, the probability of latent position remains relatively consistent across the range of these probabilities as well (Figure S1B).

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 2A, top panel) from 31 cells recorded in hippocampus while 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 being in the continuous dynamic is high throughout this event, but the probability of being in a stationary dynamic was also noticeable at the beginning and end of the SWR (Figure 2A, middle panel). Using our speed classification scheme, this means that the speed of the replay starts slower—as 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. This shows that the replay most likely travels down the center arm and up the right arm (Figure 2A, 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 2C). Importantly, when we apply the same model using the 2D position of the animal, we get a similar result as when we use the 1D linearized position (Figure S2A).

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 (42), is time-consuming, and often requires inclusion decisions that can vary widely across experimenters. We therefore adopted a "clusterless" approach which directly relates spikes and their multiunit 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 (6, 11, 26, 28). Applying a clusterless decoder to the same SWR event, we get similar classification of the sequence (Figure 2B, D), both with 1D linearized position and 2D position (Figure S2B). 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 2D 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 (3, 25, 26), position linearized to 1D). One major goal of our work is to assess the overall prevalence of spatial content across SWRs, so we detect 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 observe many replays that are classified as continuous throughout the entirety of the SWR, similar to those seen using the standard decoder (Figure S3A-F). However, we also observe many events with spatially coherent content that do not have this structure. For example, we observe trajectories that start in one direction and reverse back to the original position (Figure 3A, Figure S4B), trajectories that remain fixed in one position (Figure 3B, Figure S4G), and trajectories that jump between arms and between dynamics (Figure 3C, D, F, Figure S4F, H, I). We also observe SWRs that are spatially incoherent throughout the SWR (Figure S4A, D).

Using a 0.8 threshold, we are able to classify 89% (23,382 of 26,160) of SWRs as containing at least one of the three dynamics or common dynamic mixtures. To ensure that this depends on the sequential firing of spatially tuned cells in the

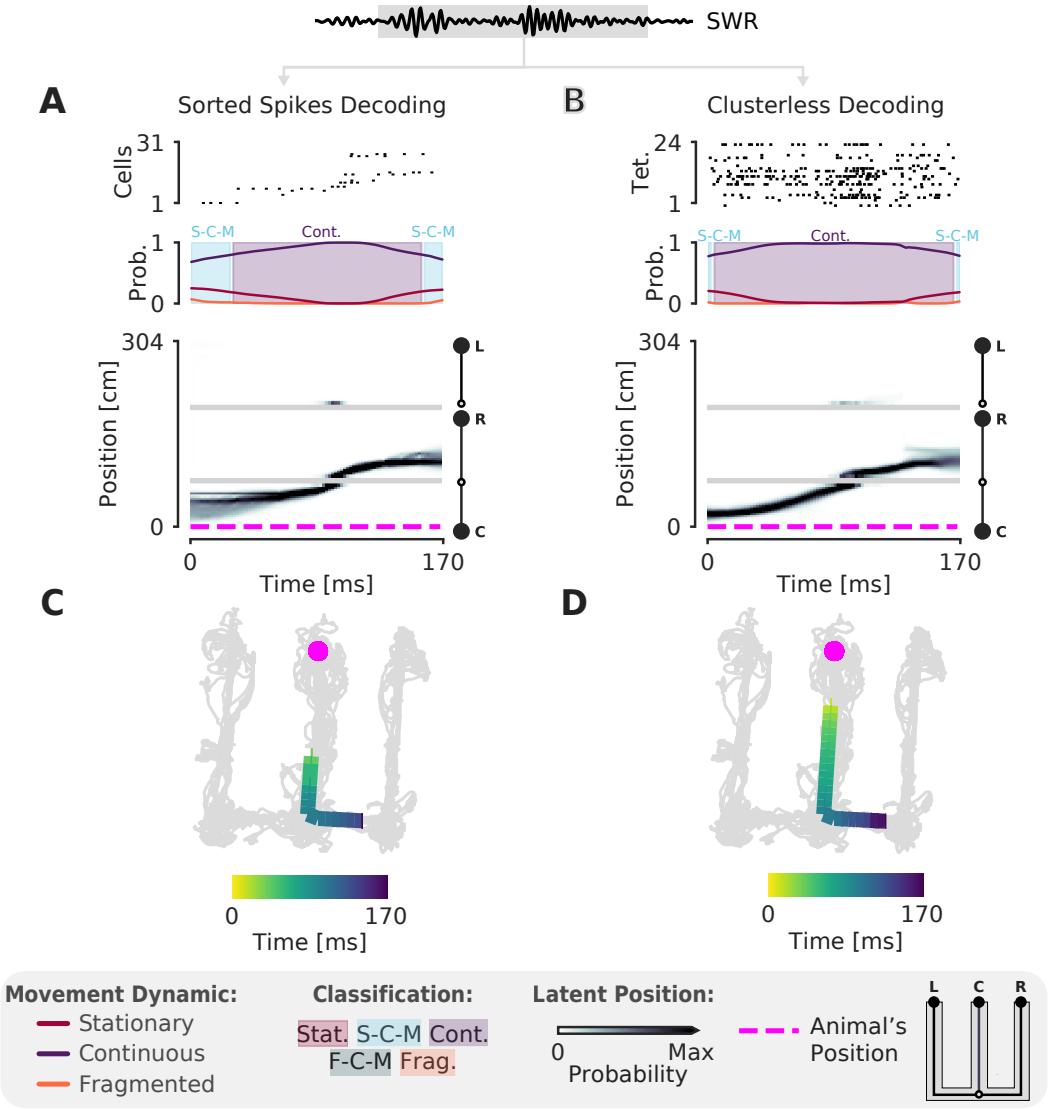

Figure 2. The model can decode hippocampal replay trajectories using either sorted and clusterless spikes from the same SWR event. **(A)** Decoding using sorted spikes. The top panel shows 31 cells on a W-track ordered according to linearized position by their place field peaks. The middle panel shows the probability of each dynamic over time as in Figure 1F. Shaded regions correspond to the speed classifications as in Figure 1G. 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.

hippocampus, we resample position with replacement for two recording sessions, shuffling the relationship between spiking and spatial information and then fitting the encoding model. We then decode the same SWR events containing the same 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 S5).

Previous work focusing on spatially sequential replay reported that only a minority of events contain sequential spatial content (10, 16, 25). We therefore asked what fraction of classified events contain spatially coherent content, which we define as SWRs containing any times with stationary, stationary-continuous mixture, or continuous dynamics (see Methods and Figure 1F). We find that 96% (22,478 of 23,382) of classified SWRs and 85% of all SWRs include spatially coherent structure and that this prevalence of spatially coherent structure is consistent across animals (Figure 3G). 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% (3,353 of 23,382) of classified SWRs have time points with spatially incoherent structure (Figure 3G).

To better 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

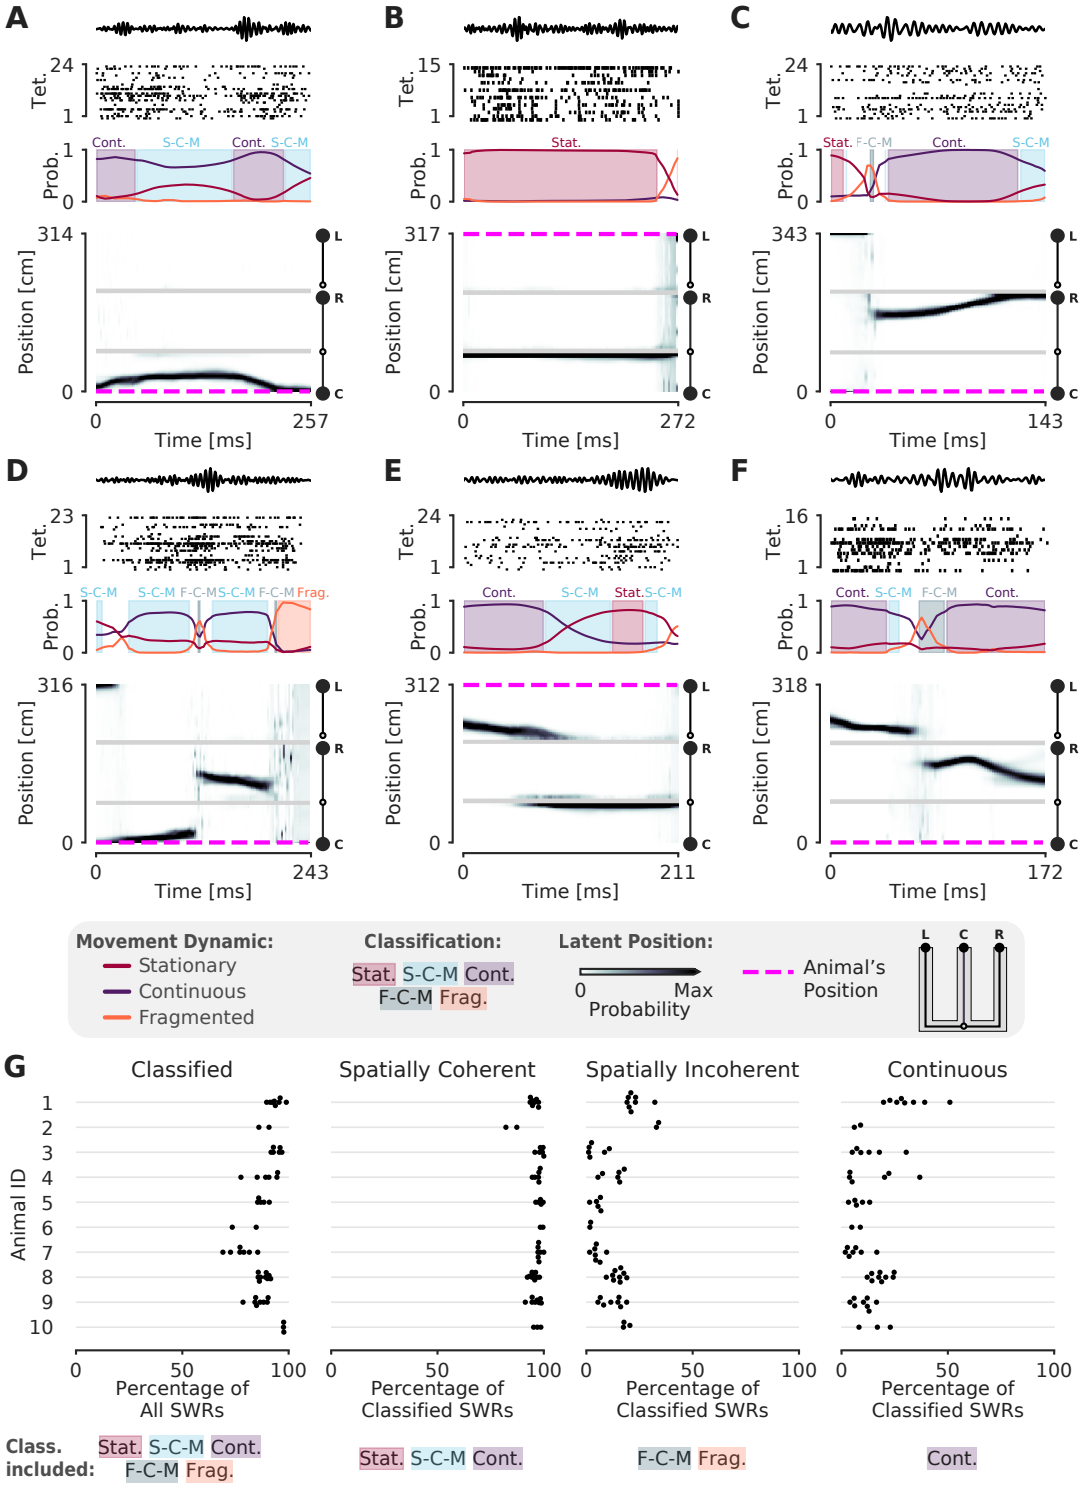

Figure 3. A-F. Six examples of SWRs with non-constant speed trajectories. Figure conventions are the same as in Figure 2. Filtered SWR ripple (150-250 Hz) trace from the tetrode with the maximum amplitude displayed above each example. **(A)** A SWR where the decoded position starts moving down the center 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.

previous reports: in our datasets, 4,511 of 23,382 or 19% of classified SWRs had time periods where the decoded position was classified as continuous (Figure 3G, 17% of all SWRs). Thus, focusing on only high speed trajectories excludes a large fraction of events where there is evidence for spatially coherent, but not continuously changing, content.

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 0.8. We find that, while this change slightly reduced the total fraction of classified SWRs (19,180 of 25,845 or 74% of SWRs), an even higher fraction of the classified SWRs (19,020 of 19,180 or 99% classified SWRs) had spatially coherent content. Similarly, SWRs containing spatially incoherent content were a small fraction of the classified SWRs (484 of 19,180 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. 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 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 provides a complementary measure of spatial coherence, 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 4A and Figure 4D show two example SWRs that are classified as having stationary and continuous 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 4C 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 4B, 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 4E). 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). Additionally, 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 4E), 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 resampled position with replacement for two recording sessions, shuffling the relationship between spiking and spatial information and then fitting the encoding model. The shuffled decodes have much larger HPD regions than the real data (epoch #1: 55 cm vs. 229 cm, p=0.02, real vs. shuffled; epoch #2: 72 cm vs. 231 cm, p=0.02, real vs. shuffled; one-sided Mann-Whitney-U, S6).

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 (16454 of 23382 or 70% of classified SWRs) are well described by a single dynamic or mixtures of dynamics (Figure 5A). Surprisingly, the most common category is stationary-continuous mixtures (9860 of 23382 or 42% of classified SWRs, Figure 5A, FigureS7). These events contain representations that move at slower speeds (Figure 5D), and are slightly shorter, on average, than events with continuous dynamics (median duration: stationary-continuous-mixture 72 ms vs. continuous 95 ms, Figure 5B). Nonetheless, both these 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 51 cm vs. continuous 42 cm, Figure 5C). This indicates that the content of these SWRs, like those that are classified as continuous, do not represent the animal's actual position.

The second most prevalent classification was exclusively stationary events (5149 of 23382 or 22% classified SWRs). Unlike the stationary-continuous mixtures, most of these events represented a location close to the animal's actual position (Figure 5C). 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 1551 of 23382 (~7%) of classified SWRs. About half of these (731 classified SWRs) were stationary throughout the duration of the SWR. These non-local stationary events most commonly represented reward wells or choice points (Figure 5E), consistent with (22), 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.

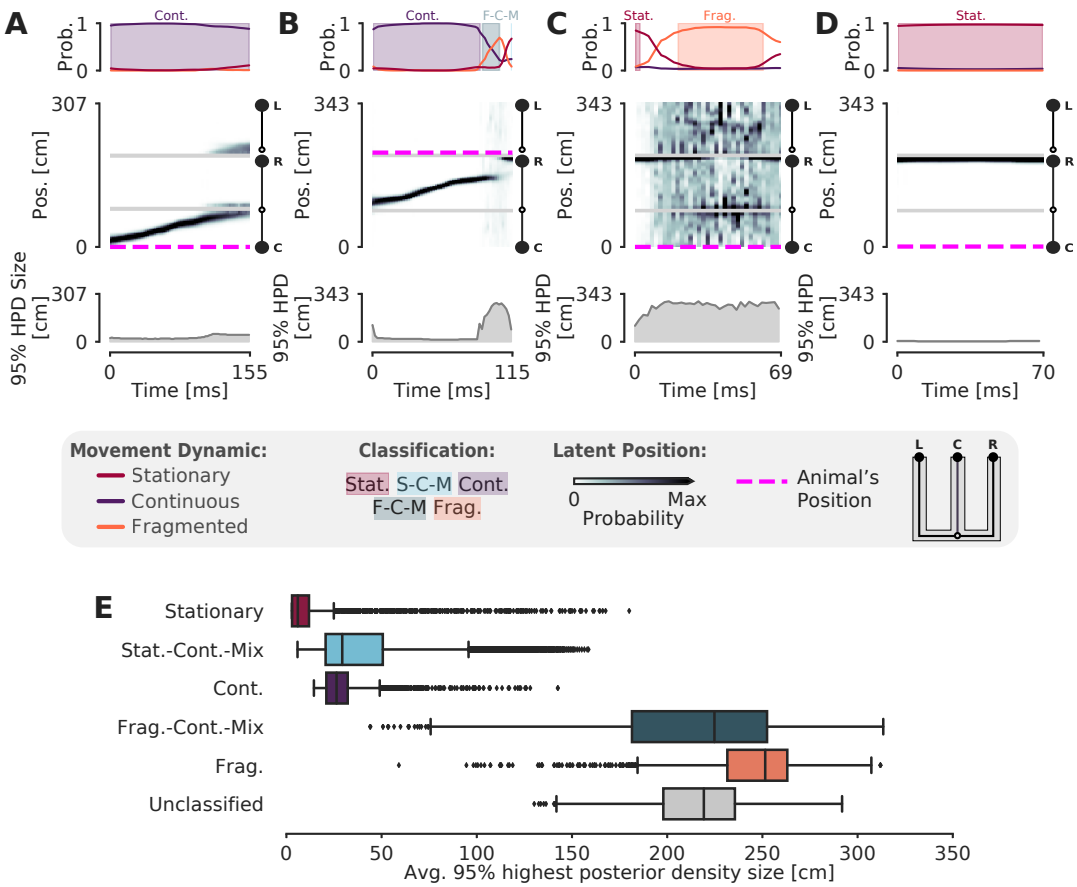

Figure 4. Validation of classification using the 95% Highest Posterior Density. **(A-D)** Examples of the 95 % Highest Posterior Density. In each column: top panel: Probability of dynamic over time. Shading and labels indicate dynamic categories. Middle panel: Posterior distribution of estimated linear 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

Finally, we examined the overall levels of neuronal spiking across dynamics. We found that all the dynamics had similar average multiunit spiking rates (Figure 5F). This indicates that all dynamics, including stationary and stationary-continuous mixtures, were driven by sustained spiking information and could not be explained by the absence of spiking.

Discussion

We developed a state space model that identifies and quantifies generative activity 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 greater than has been previously captured. 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 more or less linearly at a constant speed, but also representations that vary in speed, have slower speeds, or that are stationary. These slower-moving representations constitute the majority of representations seen during SWRs, but would not be found with the typical standard decoder. 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 (10, 16, 25). 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 (85%) contained spatially coherent content, as defined by some combination of stationary or continuous dynamics. These events included decoded representations where the most probable positions were highly spatially concentrated (i.e. had small HPD size). This was in stark 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 <1 m/s. The next most common category were stationary events (15, 22) that activated an extended representation

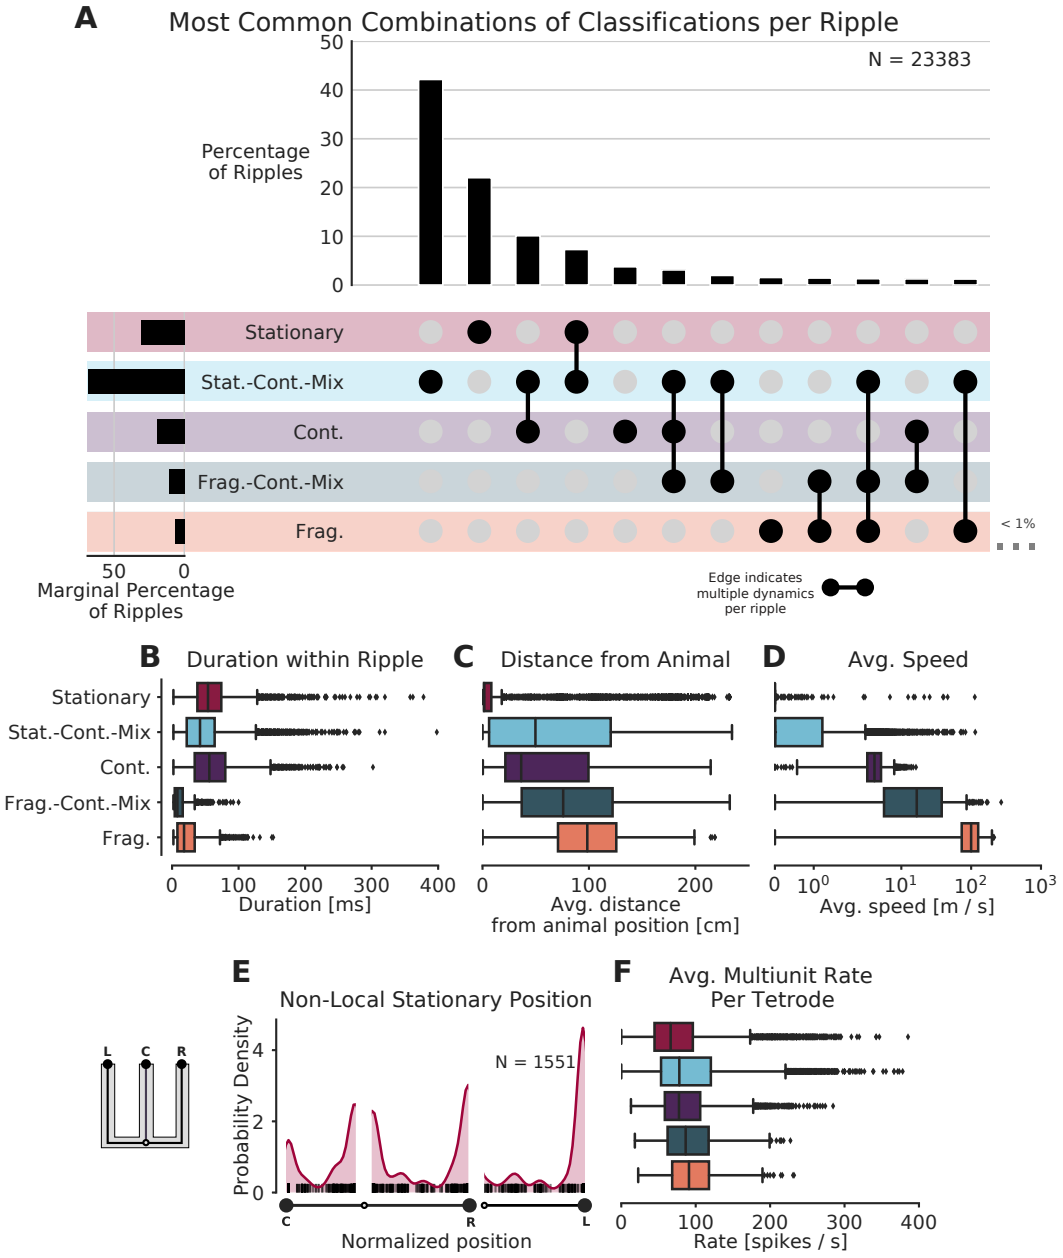

Figure 5. Prevalence of classifications. **(A)** UpSet plot (31)—which is similar to a Venn diagram with more than 3 sets—of the most common sets of classifications within each SWR. Each row represents a classification and each column represents a set of classifications, where filled-in black dots with an edge between the dots indicates that multiple classifications are present in the SWR (at temporally distinct 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).

of a single location, while the "classical" replay events—continuous, extended trajectories through space—made up only about 19% of classified events. Importantly, events with slower dynamics tended to persist for approximately the same length of time as events (~50-100 ms) with faster continuous dynamics, consistent with a brief activation of a representation of a location or a short trajectory through space. Events with slower dynamics could also represent either the animal's current location or locations far from the animal, and similarly, previous work has shown that events with continuous dynamics most often begin at the animal's current location (10, 25).

These results challenge the long-held notion that hippocampal replay trajectories necessarily proceed at many times the speed of the animal (10, 30, 36). Instead, there is a much richer variety of speeds the animal can use to retrieve information about prior experience or generate possible future events, including slower moving representations that activate a spatially compact set of nearby locations and stationary representations of single locations. Interestingly, a previous paper (15) reported that these stationary representations were most prevalent in juvenile rats and decreased to "chance" levels when the rats reached adulthood. Here, we show 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, using large time bins, and using a restrictive decoding model. We also note that we chose to include a stationary distribution 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.

We hypothesize that this large set of events supports the ability to rapidly retrieve representations of individual places and snippets of real-world experience. Previous findings demonstrated that SWRs with movement- or immobility-related activity engage specific patterns of spiking activity outside the hippocampus (21, 22). 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 evaluate the value of specific locations or short movement trajectories without having to generate full trajectories to or through them (45).

Our results may also explain the seeming conflict between Stella and colleagues (41), who reported that hippocampal replay trajectories followed Brownian-like diffusion dynamics, and Pfeiffer and Foster (38), who reported that replay trajectories contained jumps in position. We found that while replay speeds are on average distributed log normally, there are still definitive "jumps" from one position to another in a small subset of the replays. Because our model makes it possible to identify different dynamics within the same SWR, we are able to identify both types of sequences.

Additionally, as the start and end time of each SWR event are also detection method- and parameter-dependent, we suggest that instead, evaluating the start and end of each dynamic as an alternative landmark system may prove a powerful way identify periods of interest. Examination of such periods may shed light on when and why the hippocampus may be representing current position, past experience, or possible future locations, to understand 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 (22). 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 (8). 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 (7, 11, 28). We note that more fully automated modern spike sorting algorithms (such as MountainSort (9)) in combination with better recording technology may 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 but 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 avoids large temporal bins and does not make assumptions about the linear progression of events. Instead, the model can use any size of bin and allows for movement in any direction. Finally, we used an acausal estimator of latent position, which allowed us to identify the specific timing of transitions in dynamics at a fine timescale.

Our model can be viewed as an extension to the multiple temporal models approach of Johnson and colleagues 2008 (23), 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 (5, 7, 32, 33), 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. We believe this is not ideal 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 generative 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 trajectory task context and forward/reverse sequences as in Deng et al. (11). Finally, models can be built to estimate any covariate, including movement direction (26). We therefore suggest that this approach has the potential to provide a much richer understanding of neural representations throughout the brain.

ACKNOWLEDGEMENTS

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.).

COMPETING FINANCIAL INTERESTS

The authors declare that no competing interests exist.

AUTHOR 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, K.K., M.S., J.E.C. contributed data to this manuscript.

References

1. **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.
2. **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.
3. **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.
4. **Casella G**, Berger RL. Statistical Inference, vol. 2. Duxbury Pacific Grove, CA; 2001.
5. **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.
6. **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.
7. **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.
8. **Cheng S**, Frank LM. New Experiences Enhance Coordinated Neural Activity in the Hippocampus. Neuron. 2008 Jan; 57(2):303–313. doi: 10.1016/j.neuron.2007.11.035.
9. **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.
10. **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.
11. **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.
12. **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.
13. **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.
14. **Dijkstra EW**. A Note on Two Problems in Connexion with Graphs. Numerische mathematik. 1959; 1(1):269–271.
15. **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.
16. **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.
17. **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.
18. **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.
19. **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.
20. **Hunter JD**. Matplotlib: A 2D Graphics Environment. Computing in Science & Engineering. 2007; 9(3):90–95. doi: 10.1109/MCSE.2007.55.
21. **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.
22. **Jai YY**, 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.
23. **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.
24. **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.
25. **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.
26. **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.
27. **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.
28. **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.
29. **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.
30. **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.
31. **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.
32. **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.
33. **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.
34. **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.
35. **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.
36. **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.
37. **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* ACM; 2009. p. 336–343.

38. **Pfeiffer BE**, Foster DJ. Autoassociative Dynamics in the Generation of Sequences of Hippocampal Place Cells. Science. 2015; 349(6244):180–183.
39. **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.
40. **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.
41. **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.
42. **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.
43. **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.
44. **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.
45. **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.
46. **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.

Materials and methods

Simulated Data. Encoding data for Figure 1 and Figure S1 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 1F 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 (3, 25–27). 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 (25). 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 epochs 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 (37), 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 (44). 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 (27). Each CA1 LFP was obtained by down-sampling 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.

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 (12) 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:

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

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 (1, 46) 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 $N_{t_k}^i$ 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 (6, 11, 28) that the clusterless likelihood can be computed as:

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

where $\lambda_i(t_k, \vec{m}_{k,j}^i \mid x_k)$ 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:

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

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 $p_i(x_k, \vec{m}_{k,j}^i)$ is the smoothed occupancy in the space of both space and waveform features. $\pi(x)$, $p_i(x_k, \vec{m}_{k,j}^i)$, $p_i(x_k)$ 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:

$$kde(y) = \frac{1}{Nh_1 \cdots h_D} \sum_{i=1}^{N} \prod_{d=1}^{D} K_d \left(\frac{y_d - y_{i,d}}{h_d} \right)$$

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. (28). 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 (39). 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 0.8 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 0.8. 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 0.8, respectively. Time periods where none of these criterion are met are considered unclassified.

Shuffle Analysis of the Effect of Place Encoding on Classification of Dynamics. 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 epochs on two different animals (animal Bond, day 3, epoch 2 and animal Remy, day 35, epoch 2).

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 (4). In this manuscript, we use the HPD size to evaluate the uncertainty of the posterior probability of position. We 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 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.

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 defined as the center of the position bin with the greatest posterior value. The shortest path is found using Dijkstra's algorithm (14) as implemented in NetworkX (18).

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 defined as 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.

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 classifier 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 (43), Numba (29), Matplotlib (20), xarray (19), NetworkX (18), Pandas (34) and Seaborn (doi: 10.5281/zenodo.883859). All code is open-source and licensed under the MIT Software License. Classifier code can be easily installed as a python package with all requisite dependencies using pip or conda. See software repositories for specific details.

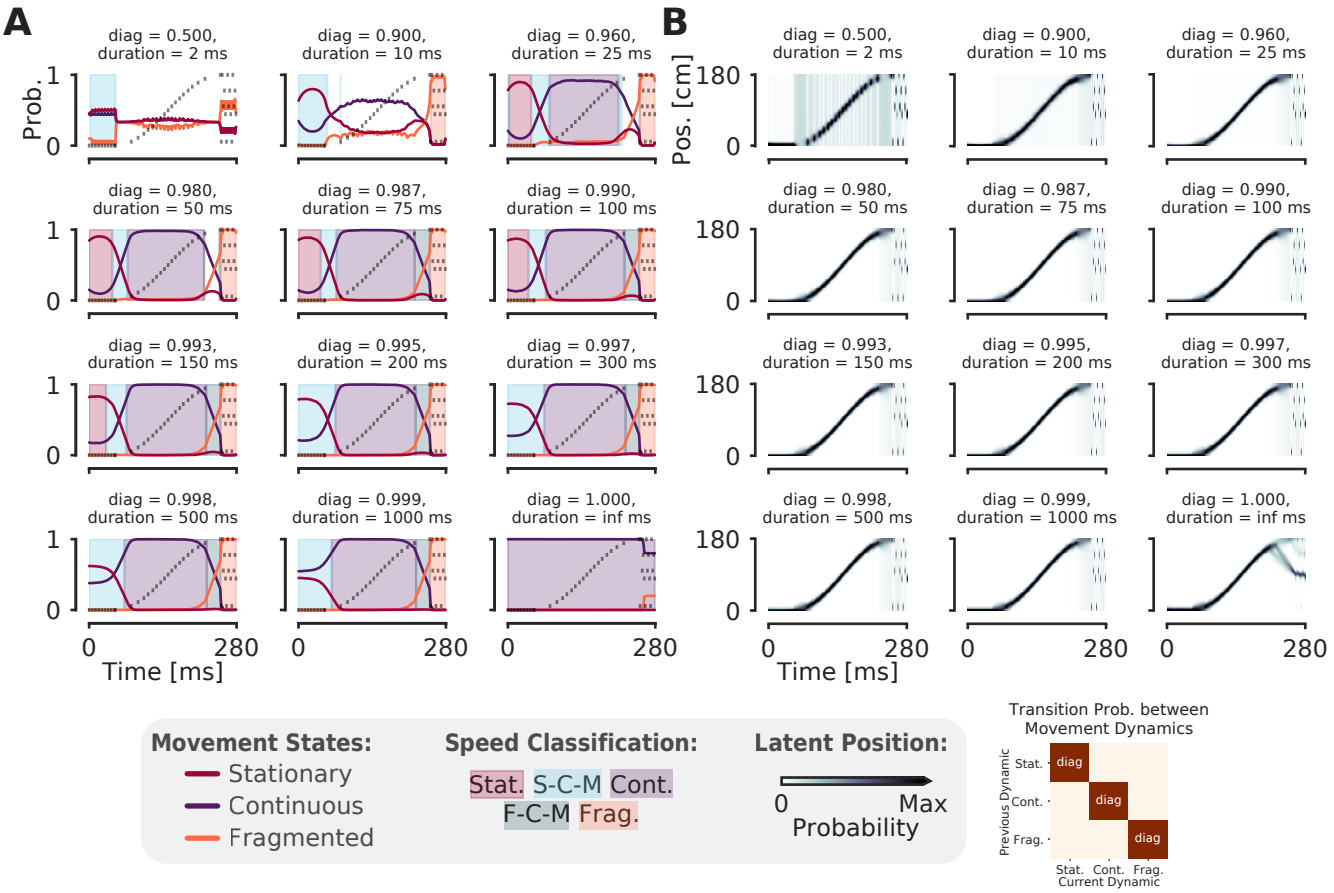

Figure S1. The model is robust to change of the probability of persisting in the same dynamic for a wide range of plausible expected durations (25 ms to 150 ms). **(A)** Each plot shows the probability of each dynamic on simulated data example from Figure 1 with a different diagonal value—which governs 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 1. **(B)** The probability of position over time for each diagonal value. Conventions the same as in (A).

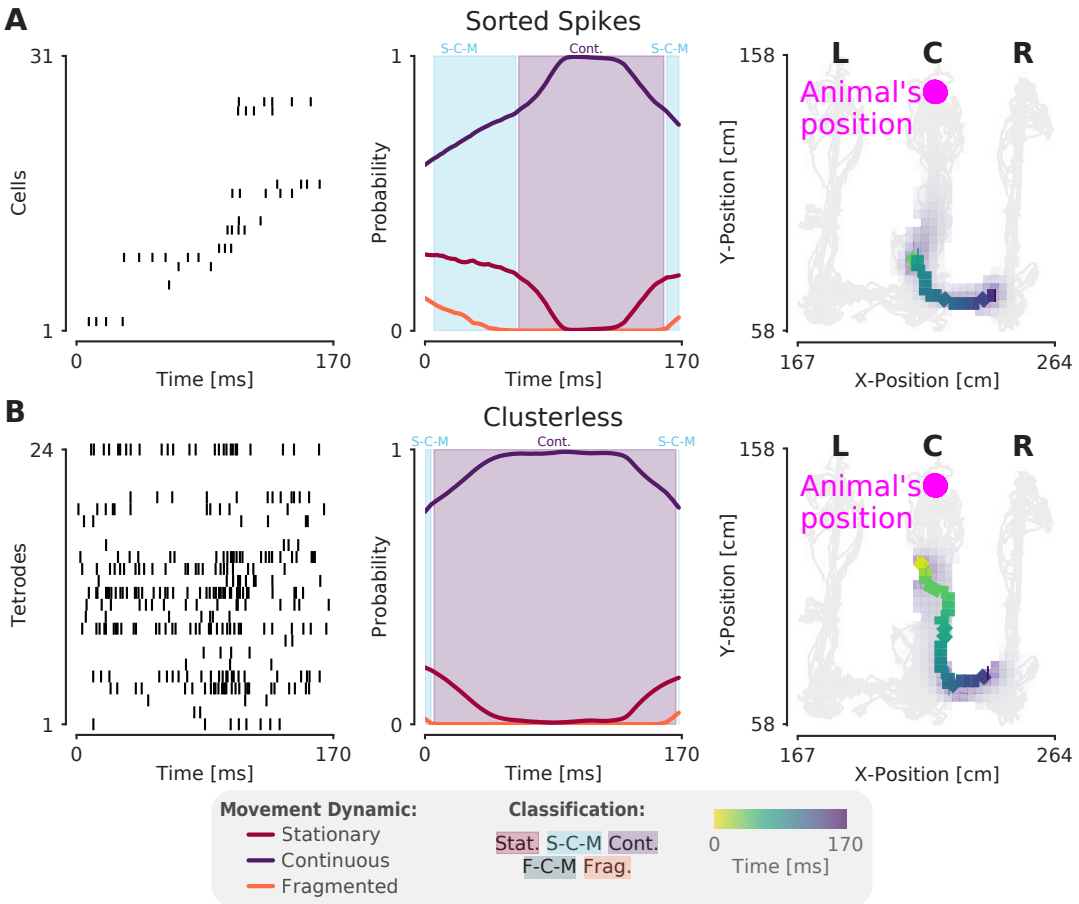

Figure S2. Decoding the same SWR in Figure 2 with 2D position using sorted spikes and clusterless decoding, respectively. **(A)** The left plot shows the spikes from cells arranged by the linear position of the peak of place field as in Figure 2. The middle plot shows the probability of each dynamic over time from the 2D decode. Shaded regions indicate classification category as in Figure 1G and 2. The rightmost plot 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 epoch and the magenta dot represents the animal's position. **(B)** Same as in (A), but with clusterless decoding.

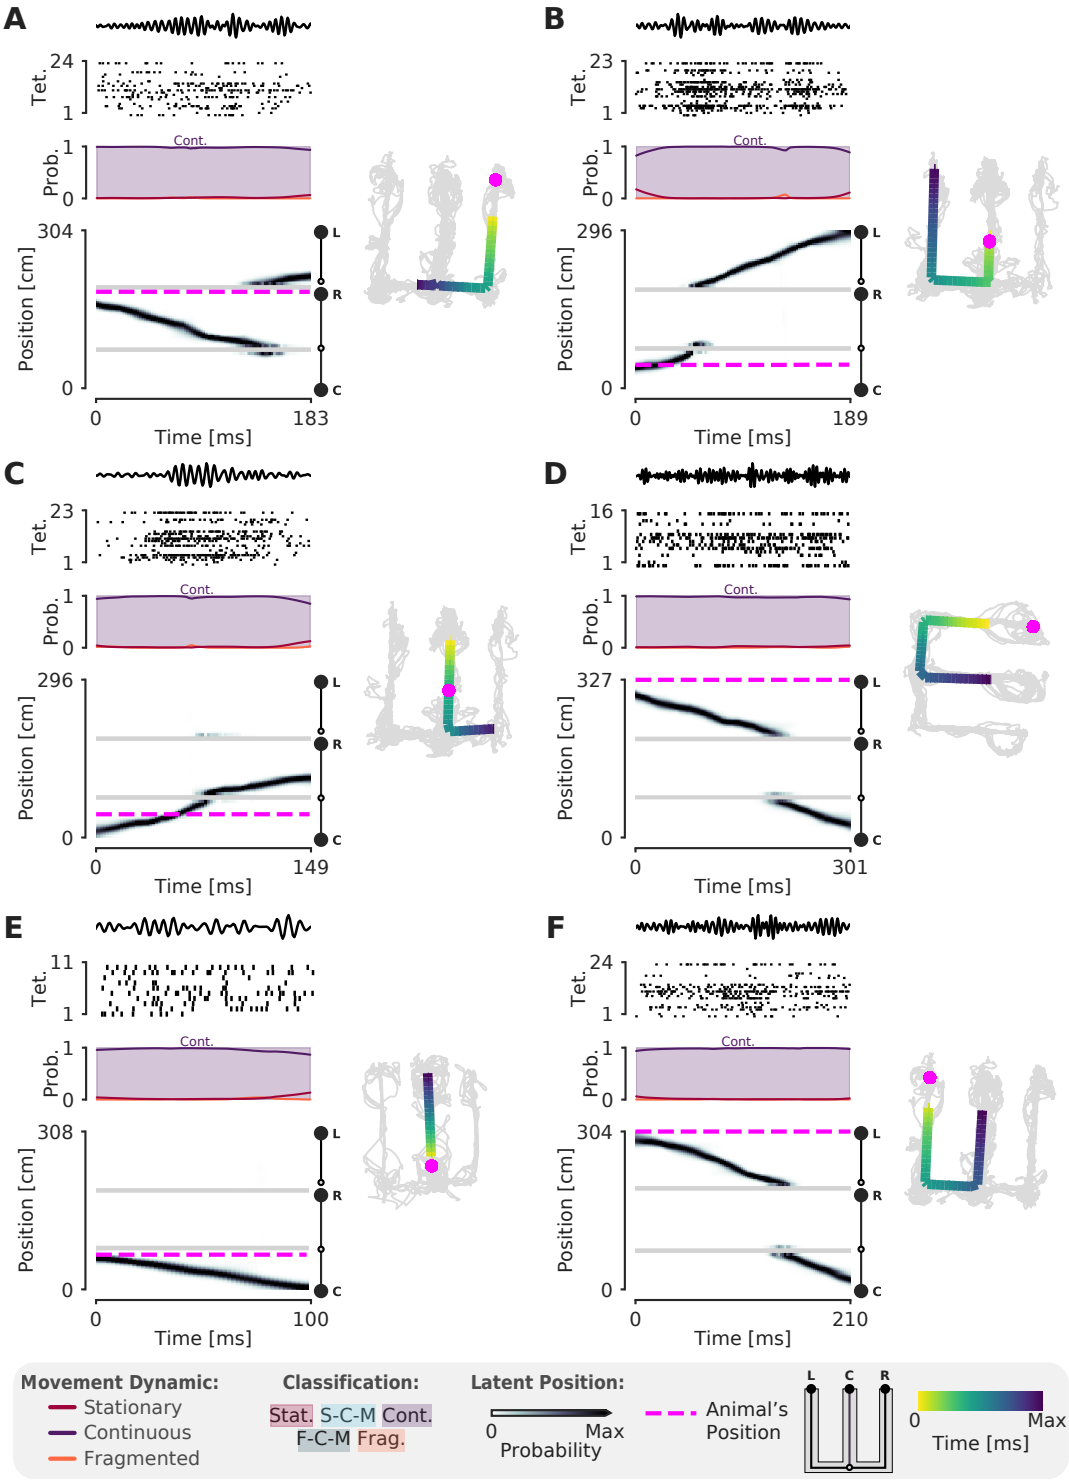

Figure S3. A-F. More examples of SWRs that have continuous trajectories. Left panel uses the same conventions as Figure 2A and 2B. Right panel shows the 1D MAP estimate projected back to 2D as in Figure 2C. Color indicates time. Light grey lines indicate the animal's position over the entire epoch. Magenta dashed line represents the animal's position during the SWR.

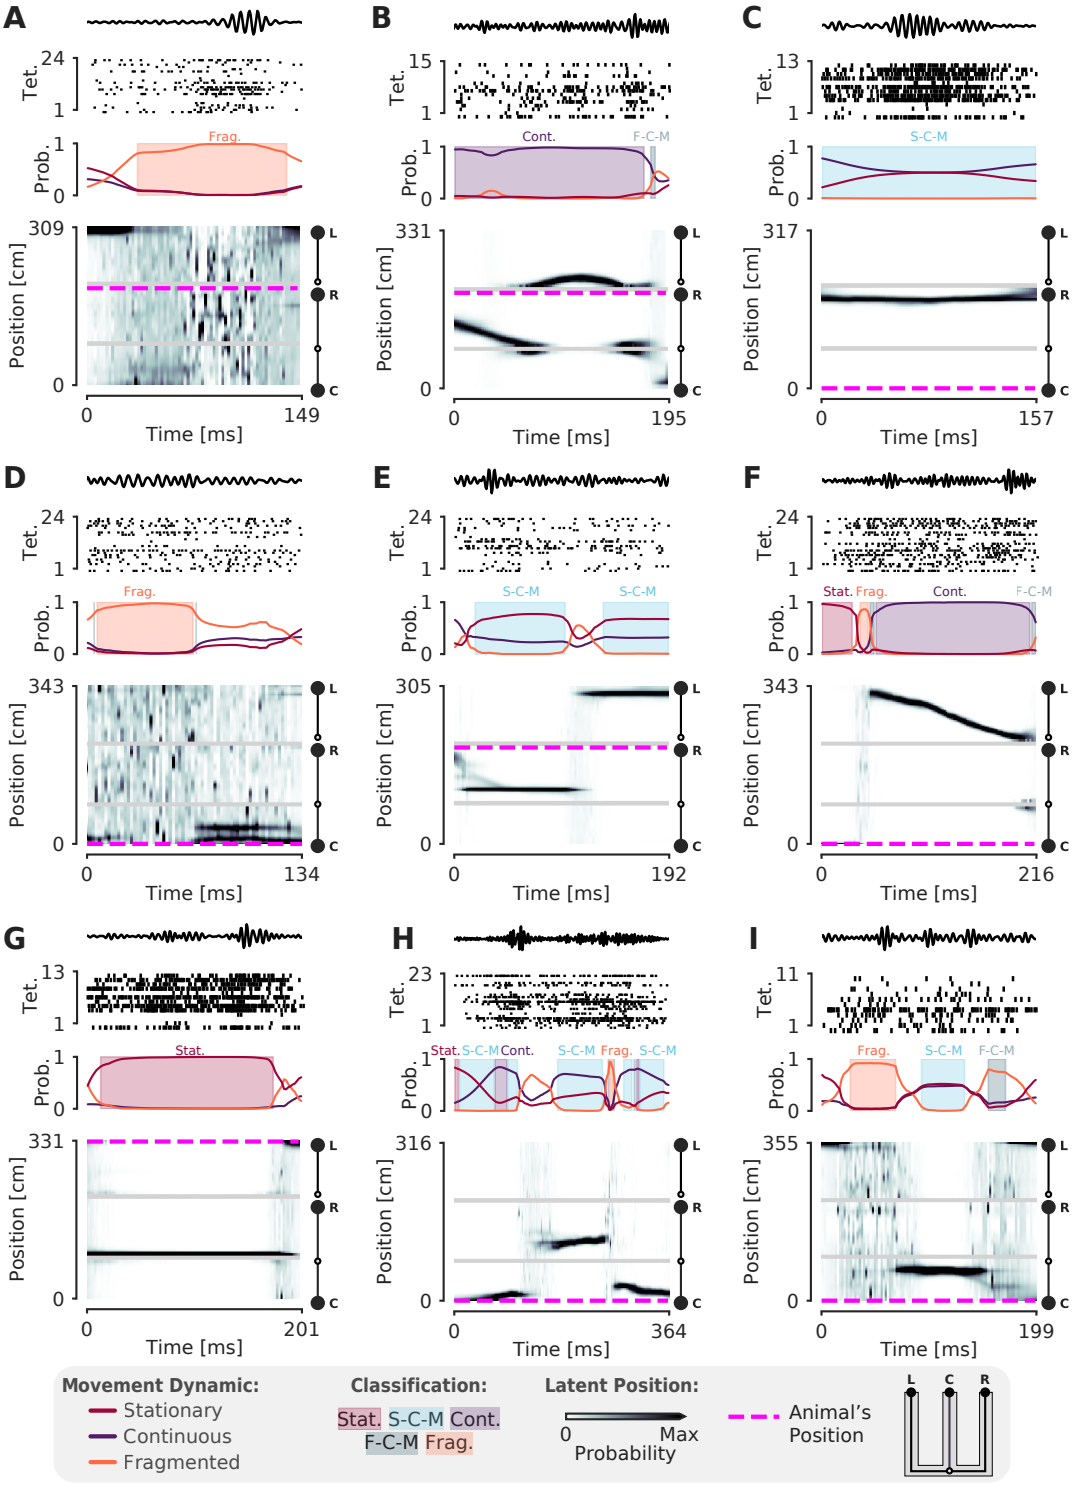

Figure S4. A-I. More examples of SWRs that would not be well-characterized by using the standard decoder. Conventions are the same as in Figure 3.

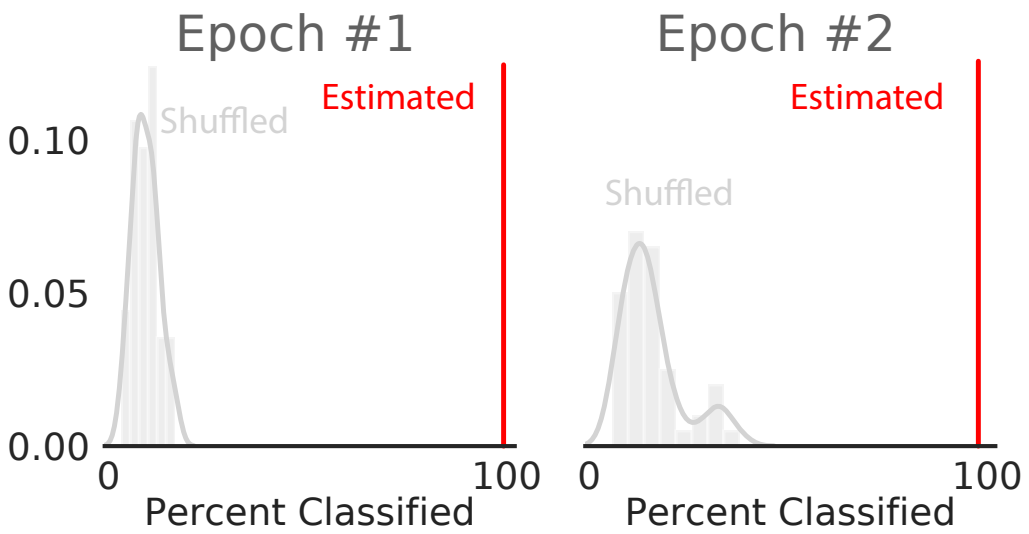

Figure S5. Comparison of percentage of SWRs classified (that is, a SWR containing at least one of the five classifications) on real vs. position shuffled data for two epochs from different animals. Red line represents the percent of SWR events classified in that epoch for real data. The histogram represents 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 epoch, destroying position information but preserving spiking timing and position occupancy.

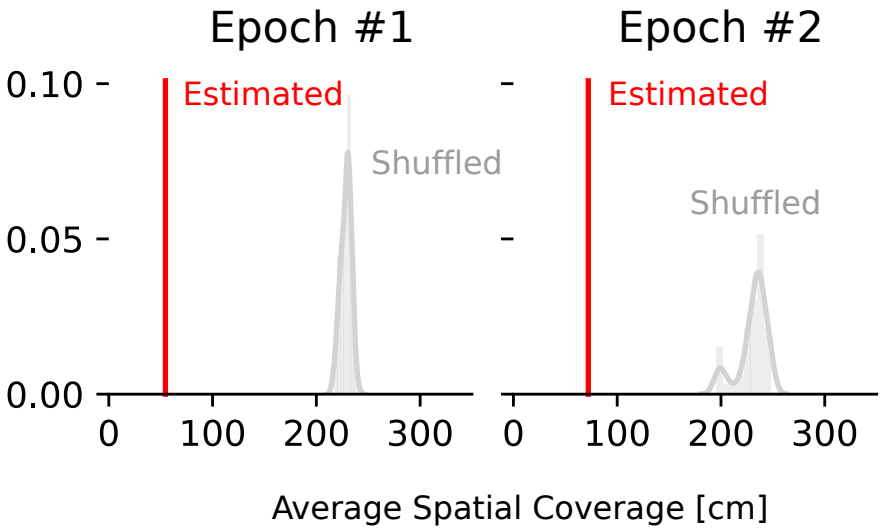

Figure S6. Comparison of average spatial coverage of the 95% highest posterior density on real vs. position shuffled data for two epochs from different animals. Red line represents the average spatial position spanned by the highest posterior density on real data for all ripples. The histogram represents the distribution after 50 shuffles of the position data. Position data was shuffled by resampling with replacement from the set of all observed positions in that epoch, destroying position information but preserving spiking timing and position occupancy.

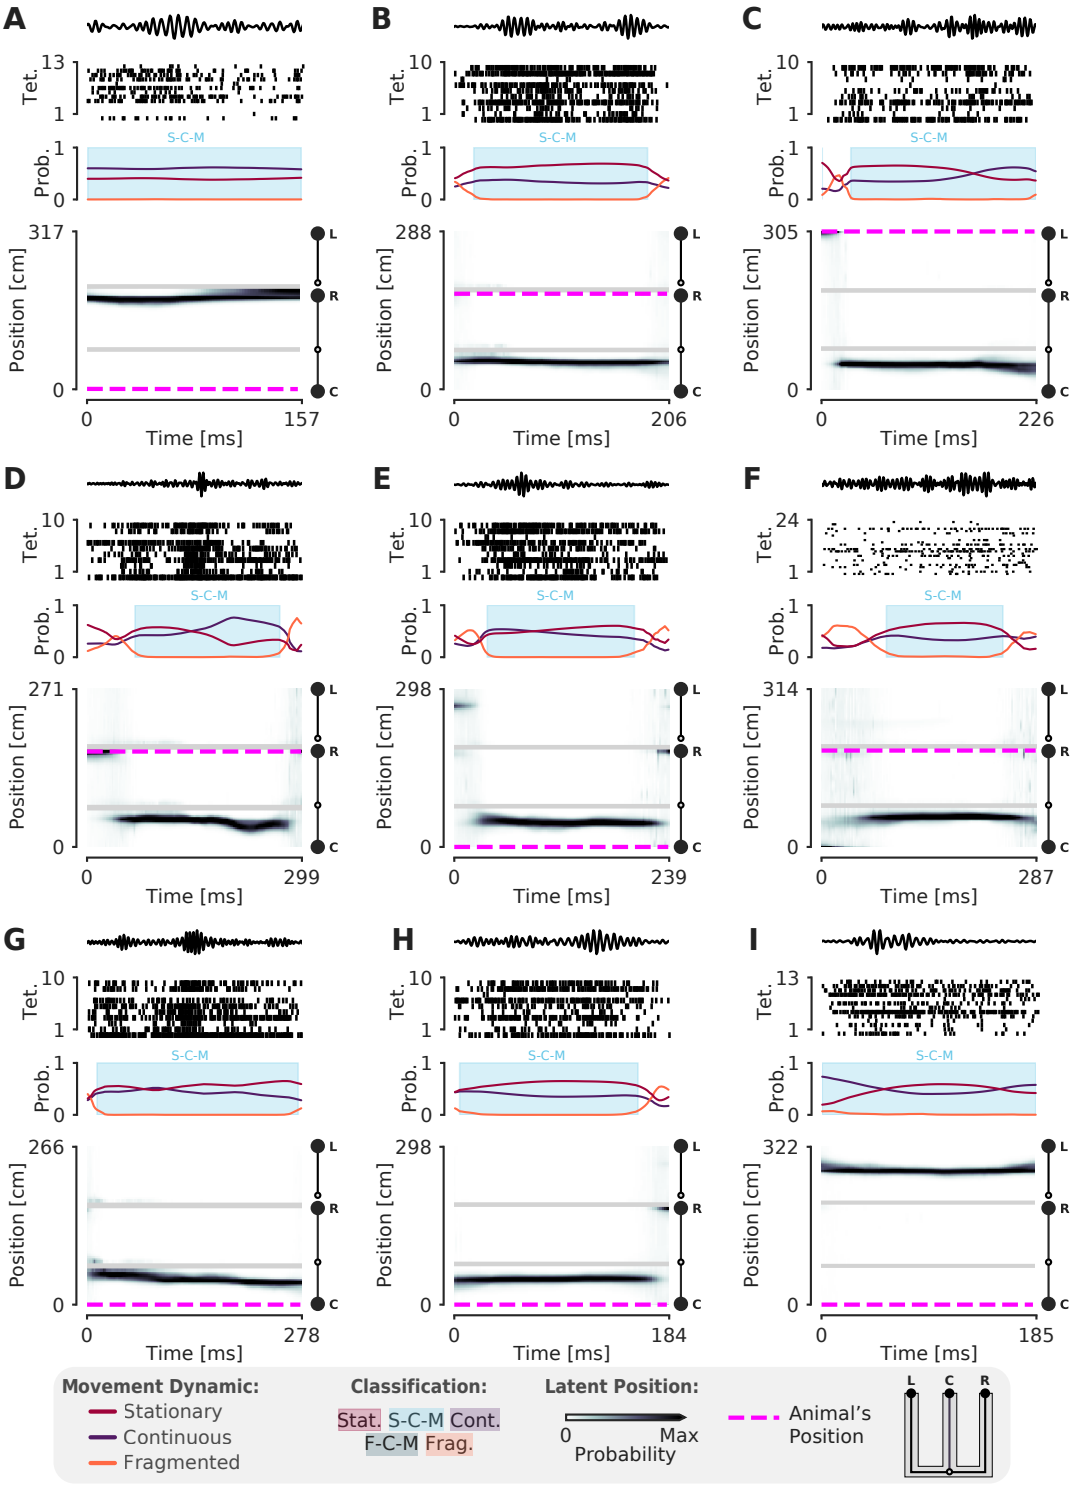

Figure S7. More examples of stationary-continuous-mixtures.