Drift-diffusion dynamics of hippocampal replay

Zhongxuan Wu[1,3,4,5] and Xue-Xin Wei[1,2,3,4,5,*]

1) Department of Neuroscience, The University of Texas at Austin
2) Department of Psychology, The University of Texas at Austin
3) Center for Perceptual Systems, The University of Texas at Austin
4) Center for Learning and Memory, The University of Texas at Austin
5) Center for Theoretical and Computational Neuroscience, The University of Texas at Austin

* Correspondence: weixx@utexas.edu

Abstract

Replay in the hippocampus and other brain regions during sharp-wave ripples (SWRs) is thought to play major roles in learning, memory, and planning. Toward clarifying the basis and functions of replay, it is critical to characterize its dynamical structure. Standard methods rely on restrictive assumptions for detecting replay events and may lead to inconsistent and inaccurate metrics for characterizing replay dynamics. We develop a computational framework based on state-space modeling in conjunction with drift-diffusion dynamics, an approach that provides an appropriate combination of flexibility and interpretability. In this approach, two parameters corresponding to drift and diffusion can be directly mapped to the speed and variability, respectively, of a single replay event. Further, to capture the potentially rich population-level dynamics expressed during SWRs, our model allows switching between multiple well-motivated types of dynamics (i.e., drift-diffusion, stationary and uniform). Applications of our method to all detected SWRs from rat hippocampal recordings lead to insights into a number of important open questions. First, our results show that, while a substantial fraction of SWRs are stationary, those with drift-diffusion dynamics generally have a much higher speed than the animal's movement. Second, we find that replay dynamics in 1-D environments are heterogeneous and are inconsistent with a random walk model proposed previously. The mean-squared displacement of replay trajectories scales quadratically with time, supporting the presence of substantial drifts. Third, we find that only a tiny fraction of SWR events (less than 1%) exhibit sequential structure at the timescale of 100 ms before the animal had spatial experience in an environment. This suggests that, while neural activity in the hippocampus may be coordinated before the animal's spatial experience (i.e., preplay), the level of coordination is much weaker than that after spatial experience. Overall, the proposed analysis approach enables precise and unambiguous interpretations of population dynamics during SWRs, which more broadly can provide a better understanding of the functions and underlying mechanisms of replay.

1 Introduction

Neural activities in the mammalian hippocampus during sharp-wave ripples (SWRs) have been found to "replay" previous experience in forward [1, 2] or reverse [3, 4] order, both during sleep [5–8] and awake immobility [3, 9]. Replays in the hippocampus are also found to be coordinated with neural activities in other brain regions, e.g., prefrontal cortex [10–12], visual cortex [13], and brain-wide spontaneous dynamics [14]. While initially reported in rodents, experimental evidence for replay in the human brain has also been reported in recent studies [15–18]. As a form of selective and experience-dependent neural activity, replays have been proposed to play an important role in learning, memory consolidation, and planning [19–22], although their precise functions remain to be determined.

Surprisingly, even given decades of research on hippocampal replay, the methodology of detecting replays and, at the same time, identifying their properties on an event-by-event basis remain an ongoing challenge, both statistically and also with respect to defining the phenomenon [19, 23]. Most studies rely on template matching, where putative replay activity is compared to a template generated from the population activity during "online" periods, typically taken to be running [5, 6]. More specifically, this template can be defined using either a sequential order of neurons or the sequential structure in a decoded variable, conventionally spatial location. However, many SWR events remain ambiguously classified according to this approach, rendering a challenge for subsequent analyses and, further, suggesting that many replays may have structure that is not well-captured by existing approaches. As an example, in prior methods, the speed and quality (coherence) of replays are entangled. Replay events with high speed tend to be significant despite having lower coherence, even though these variables are ideally independently estimated. Indeed, recent studies suggest that replays may exhibit richer dynamics than traditionally considered or assumed [24, 25], and the properties of replay dynamics may have major implications for learning [26]. In this way, the lack of a principled and appropriately general computational framework for characterizing the dynamics of replay constitutes a major obstacle to analyzing and interpreting ongoing experimental data and findings, respectively.

Here we look to address this problem by developing a principled probabilistic method for unbiased characterization of the dynamics of replay. This method builds upon prior work [24, 25, 27–29] and overcomes multiple limitations in accurate identification of replay dynamics. We systematically validate our method using simulated datasets, and demonstrate that our method can accurately recover the dynamics and drift-diffusion parameters under neural data scenarios. We will show that applications of our modeling framework to neural population recordings from the rodent hippocampus lead to new insights into the precise neural dynamics of the hippocampal replay. In particular, our results have direct implications for three open questions regarding hippocampal replay: (i) whether the speed of most replay events is comparable to real-world running; (ii) whether replays follow a random walk; (iii) whether "preplay" events exist. By enabling more flexible and interpretable characterization of replays, this new analysis framework may provide a path toward better understanding of the functions and mechanisms of replay.

2 Results

A replay event is generally thought to be a sequence of neural activity that encodes a traversal of a particular state space (e.g., positions on the track). Most previous studies use a two-step procedure to analyze replays: (i) decode spatial location from neural population activity, typically using a Bayesian decoder that yields decoded posterior probabilities; (ii) apply a "manually" specified form of template matching to the decoded posterior probabilities to detect putative sequential structure that would qualify as a (traditional) replay. Despite the success of this approach in identifying classical replays, the major limitation is the rigidity of the pre-specified template, which effectively disregards the potential richness of replay dynamics. A second major limitation is that the estimation in the second step fails to incorporate uncertainty, which limits quantitative evaluation of robustness. To address these issues, we develop a novel set of state-space models based on drift-diffusion dynamics to characterize the structure of replay events. By basing dynamics on drift diffusion, a general yet interpretable and parsimonious process, our approach rigorously formalizes past intuition regarding replays, and moreover allows for direct probabilistic inference of replay dynamics from neural data.

2.1 Drift-diffusion models of replay dynamics

Consider an experiment where the animal is running on a one-dimensional track (Fig. 1A,B), a common behavioral paradigm in studying replays. Although we focus on modeling replays in 1-D environment, our model can be generalized to more complex environments. We assume that hippocampal activity during a replay event encodes a sequence of positions on the track (Fig. 1C). However, in contrast to previous work, we propose to model changes in encoded positions during replays as evolving according to the dynamics specified by drift-diffusion. Formally, these dynamics of replayed positions z during putative replays (population activity during SWRs) evolves according to (Fig. 1D):

$$\mathrm{d}z = \lambda \mathrm{d}t + \sigma \mathrm{d}W \tag{1}$$

where λ and σ are the drift and diffusion parameters, respectively. $\mathrm{d}W$ denotes the standard Brownian increment. This model can be implemented in discrete time steps, with the dynamics described by the state-transition matrix. We further assume that, given a latent state, the activity of individual neurons follows an independent Poisson process with firing rate dependent on location-specific fields (place fields) –an assumption shared with the conventional Bayesian decoding method for analyzing replays (see Methods).

In this framework, the expected latent state at the next time step is determined by the current state and the drift rate, with the variability of the states determined by the diffusion parameter. In this way, both parameters (drift and diffusion) are highly interpretable. More specifically, the drift parameter is analogous to the speed (or "virtual speed") of the replay event. The diffusion parameter corresponds to how noisy the replay trajectory is, thus providing a quantitative measure of the variability of a replay event. Different values of two parameters are able to capture a variety of dynamical patterns while retaining

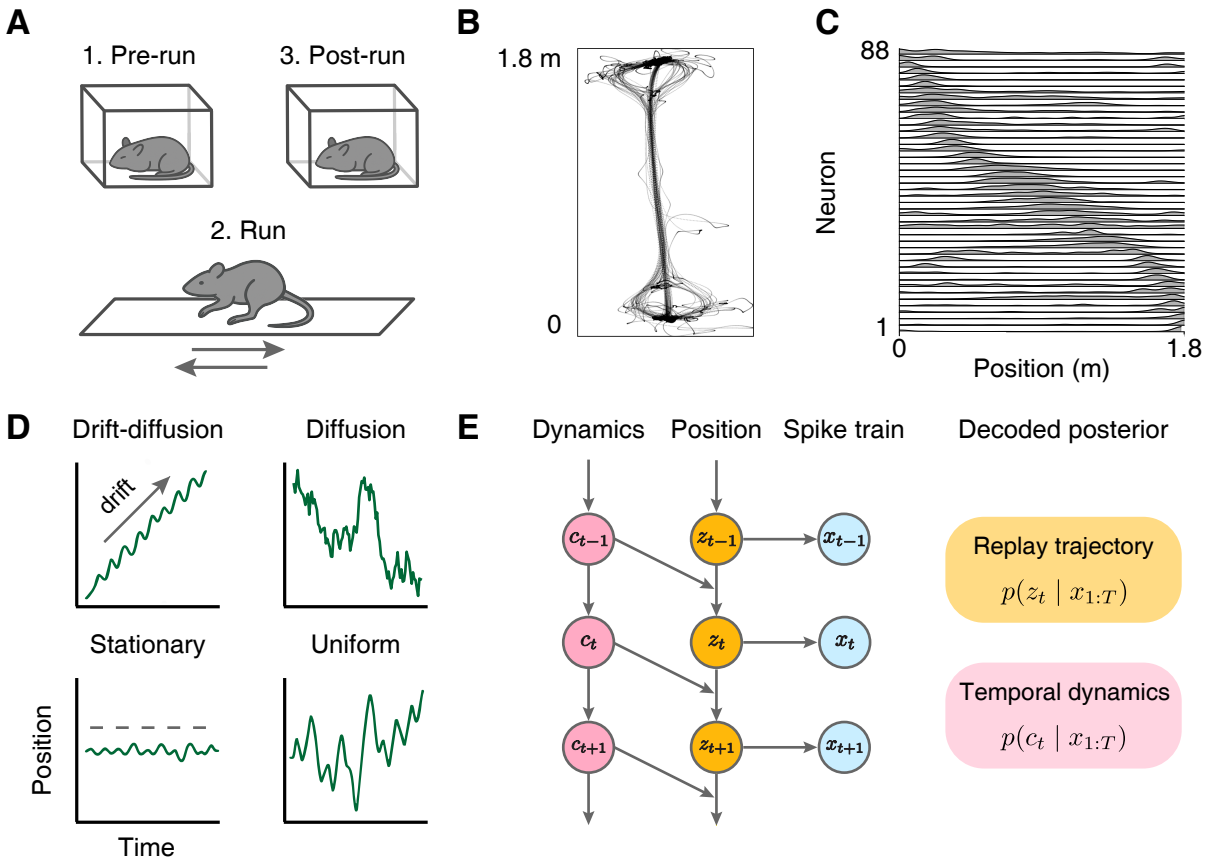

Figure 1: Animal behavior and the overall computational modeling framework. **A**. Schematic diagram of a commonly used behavioral paradigm for studying hippocampal replay. Each recording session consists of three behavioral epochs: a pre-run sleep epoch in an isolated box, an active run epoch on a 1-D linear track, and a post-run sleep epoch in the same box used for the pre-run. **B**. Head-location tracking data during the run epoch from an example recording session. The animal performed repeated back-and-forth traversals along the 1.8 m track. The track is 6.35 cm wide, with two ends expanding to 12.7 cm for reward wells. Data replotted from [30]. **C**. Place fields of the recorded CA1 neurons, estimated from spikes during locomotion (> 5 cm/s) using 2 cm spatial bins and Gaussian smoothing (σ = 5 cm). Only half of the recorded population is plotted for clarity. Data from [30]. **D**. Potential spatiotemporal dynamics underlying the latent trajectory during sharp-wave ripples (SWRs). *Drift-diffusion*: a mixture of deterministic drift and stochastic diffusion, $dz = \lambda\, dt + \sigma\, dW$, where the drift parameter λ sets the "speed" and diffusion parameter σ sets the amplitude of Brownian motion fluctuations. *Diffusion*: pure Brownian motion, $dz = \sigma\, dW$, representing an unbiased random walk. *Stationary*: dz = 0, indicating no movement over time. *Uniform*: drift-diffusion process with near-zero λ and large σ, implemented as independent resampling of position from a uniform distribution at each time step. Note that the uniform trajectory has no temporal continuity. **E**. Left, probabilistic graphical model of a switching HMM, where a discrete variable of dynamics (regime) c_t determines a regime-specific transition probability for latent position z_t. z_t then generates the observed spike counts x_t with Poisson emissions parameterized by estimated place fields. Right, applying probabilistic inference to the switching HMM produces a marginal posterior of trajectory $p(z_t \mid x_{1:T})$ in conjunction with a marginal posterior of regime variable $p(c_t \mid x_{1:T})$. For derivations of marginal posteriors, see Methods and Supplementary Information.

interpretability (Fig. 1D). In the case of zero drift and nonzero diffusion, the latent trajectory is equivalent to one-dimensional Brownian motion, a pattern recently proposed to capture replay dynamics [31]. In the case where both drift and diffusion parameters are near zero, the trajectory is largely stationary, a dynamical pattern emphasized in [24]. Finally, in the case where the drift rate nearly zero, and diffusion is large, the discretized trajectory would behave similarly to a collection of independently sampled states

having little coherence. We refer to this type of event as a uniform event. Conceptually, this is the opposite regime of replay events discussed in previous studies.

Extended modeling framework to incorporate switching dynamics. A fundamental matter that remains unresolved regarding replays is whether and to what extent individual SWR events express multiple dynamical patterns serially [24]. To capture the potential richness of the dynamics of an SWR event, we incorporate the switch between different types of dynamics within an individual event (Fig. 1E). Importantly, both the type of dynamics and the states (i.e., positions) can be inferred from the neural activity based on our model. The resulting one-step framework represents the most powerful way to date to clarify whether and to what extent there are mixtures of dynamics. Specifically, we propose that a replay event may be serially comprised of three types of dynamics (regimes): drift-diffusion, stationary, and uniform sampling (Fig. 2A). For a given event, in addition to the latent position evolution, there is also a finite probability of switching to a different type of dynamics. We assume that the transition probability matrix between regimes has equal self-transition probability q and identical switching probabilities $\frac{1-q}{2}$ of switching to the other two regimes. The parameter q is optimized for individual replay events and thus can capture variability of dynamical switching which may be expressed across different events, as suggested in previous work [33]. The transition matrix between different regimes can be easily incorporated into the model by conditioning the state transition matrix for each type of dynamics. We develop inference procedures for estimating the parameters from the neural population activity, leveraging established techniques in HMMs and state-space models [34–41]. For each SWR event, we optimize the model parameters $\{\lambda, \sigma, q\}$ by maximizing the likelihood of observing the neural population activity. With optimized parameters, we can calculate the posterior distribution of dynamics and positions (see Methods and Supplementary Information A for details). Fig. 2B-H show example events from analyzing a hippocampal population recording dataset [30].

2.2 Validation of the modeling framework

Our full model of replay dynamics can be interpreted as a specific type of switching HMM. Although identifiability has been established for related models [42, 43], two potential issues are relevant for our formulation: (i) under maximum-likelihood estimation (MLE), are the drift and diffusion parameters governing the latent-position transition probability recoverable, and (ii) can dynamic regimes be correctly classified? A practical constraint is that replay events are brief with a typical length of 100-250 ms and recordings comprise only dozens to a few hundred place cells. This relatively limited quantity of data may make estimation challenging. Thus, it is materially important to evaluate the performance of our inference procedures under conditions that are relevant to this experimentally relevant constraint.

We first validated parameter identifiability for the basic drift-diffusion model without switching dynamics. We generated latent position trajectories from a two-dimensional grid of drift and diffusion values, then simulated spike trains from synthetic place fields conditioned on each trajectory. By fitting our model to these synthetic data, we can assess how well the model parameters can be recovered (Fig. 3A). In general, both drift and diffusion parameters were satisfactorily recovered (Fig. 3B-C). We found that the

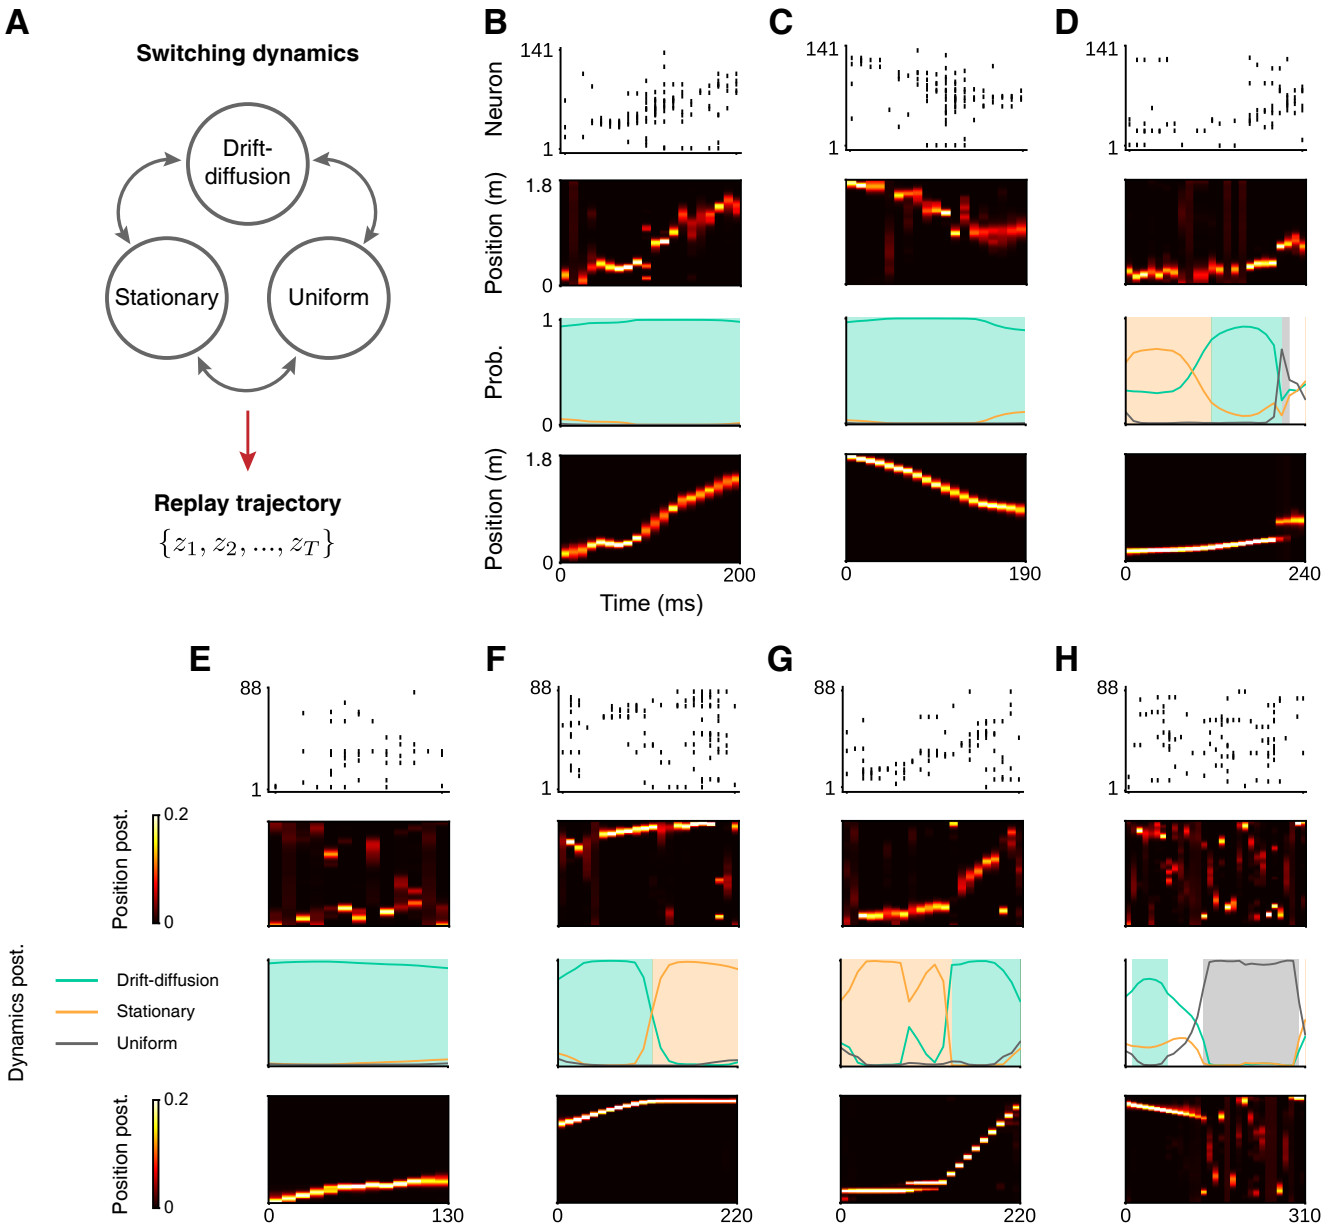

Figure 2: Decoding the heterogeneous dynamics in rat hippocampal SWRs using switching HMMs. A. Schematic diagram of the switching dynamics in our framework. At each time step, the model can switch between three basic dynamics—drift-diffusion, stationary and uniform (see definitions in Fig. 1D). The probabilities of these basic dynamics together determine the temporal transition of latent positions $p(z_2 \mid z_1), p(z_3 \mid z_2), \ldots, p(z_T \mid z_{T-1})$. **B-H**. Seven representative SWRs from a rat hippocampal recording dataset [30] are each shown as four subpanels. Top to bottom: (1) Spike raster of the SWR. Spikes are binned using 10 ms non-overlapping time bins. For visualization, place cells are sorted by the peak location of their place fields. (2) Posterior probability of the latent positions decoded from a memoryless Bayesian decoder [32]. (3) Posterior probability of the three basic dynamics decoded from our model, $p(c_t \mid x_{1:T})$. Background shading indicates the assigned type of dynamics based on the posterior. If $p(c_t = \text{any type of dynamics} \mid x_{1:T}) > 0.7$, we assign that type at time step t. Otherwise, if $p(c_t = \text{drift-diffusion}) + p(c_t = \text{stationary}) > 0.7$, we label such time steps stationary, as drift-diffusion reduces to stationary in this case. Time steps that do not meet any of these criteria are labeled unclassified. Color code: drift-diffusion, green; stationary, orange; uniform, gray; unclassified, none. (4) Posterior probability of the latent positions decoded from our model, $p(z_t \mid x_{1:T})$.

diffusion parameter was considerably more variable than the drift, and variance of the diffusion parameter tended to increase with the true diffusion magnitude. We also found that the diffusion estimates tend to be underestimated because each replay event only contains a limited number of time steps (see SI for a further discussion on the origin of this estimation bias).

We next evaluated identifiability in the switching HMM, the model that enables inference of switching between dynamical regimes. As before, we simulated latent trajectories with regime switches and generated spike trains from synthetic place fields, then optimized the full model. We first asked whether dynamics classes could be correctly decoded. We found that fitting yielded robust regime classification overall, as shown by confusion matrices comparing inferred and true regimes for simulated SWRs (Fig. 3D-E). As expected, when both drift and diffusion are near zero and observed replay sequence is short, the separability between drift-diffusion and stationary regimes is reduced. With a substantial drift magnitude, separability improves and classification accuracy between drift-diffusion and stationary increases markedly (Fig. 3E). Consistent with the basic drift-diffusion analysis, the full model also yields accurate parameter recovery. MLE estimates of both drift and diffusion recover the ground truth across simulated conditions reasonably well (Fig. 3F-G). However, there did remain modest underestimation of the diffusion parameter, attributable to short events. Together, these results show that our full model can reliably identify the dynamical regimes and the drift-diffusion parameters under experimentally realistic data sizes.

Comparison to prior methods. We compared our modeling framework with the traditional 2-step analysis procedure, which first uses Bayesian decoding to map the neural population activity onto the posterior distribution over positions [32], then estimates descriptive statistical parameters (such as absolute correlation, slope of replay trajectory, maximum jump distance) [44, 45] from the resulting posteriors (see Fig. S1A). Using simulated data, we compared the drift/diffusion parameters to the key metrics used in prior studies. Based on extensive simulations, we found that, while the three key metrics may capture some structure in replays as conceived when they were developed, each measure is relatively confounded with the estimated speed and quality of the candidate replay event (see SI text and Fig. S1 for detailed explanation). Thus, in practice the traditional 2-step approach may not allow for an unambiguous interpretation of the replay dynamics. In contrast, our method adopts a more principled approach, namely explicit modeling of drift-diffusion, as a means to systematically disambiguate the speed (drift) and quality (diffusion) of a replay event. Importantly, this approach enables parameter estimation in a single step, allowing a more principled modeling of uncertainty.

2.3 Applications to experimental data to investigate open questions about replay

We next apply our modeling framework to investigate the structure and dynamics of hippocampal replay. We will investigate three questions that remain unresolved in research on replay. First, do hippocampal replays depict trajectories that evolve at real-world speeds, or are even stationary [24]? Second, does hippocampal replay evolve similarly to Brownian motion (random walk) [31]? Third, how does the structure of the hippocampal "preplay" events (during pre-run sleep) compare with replay events [44, 47]?

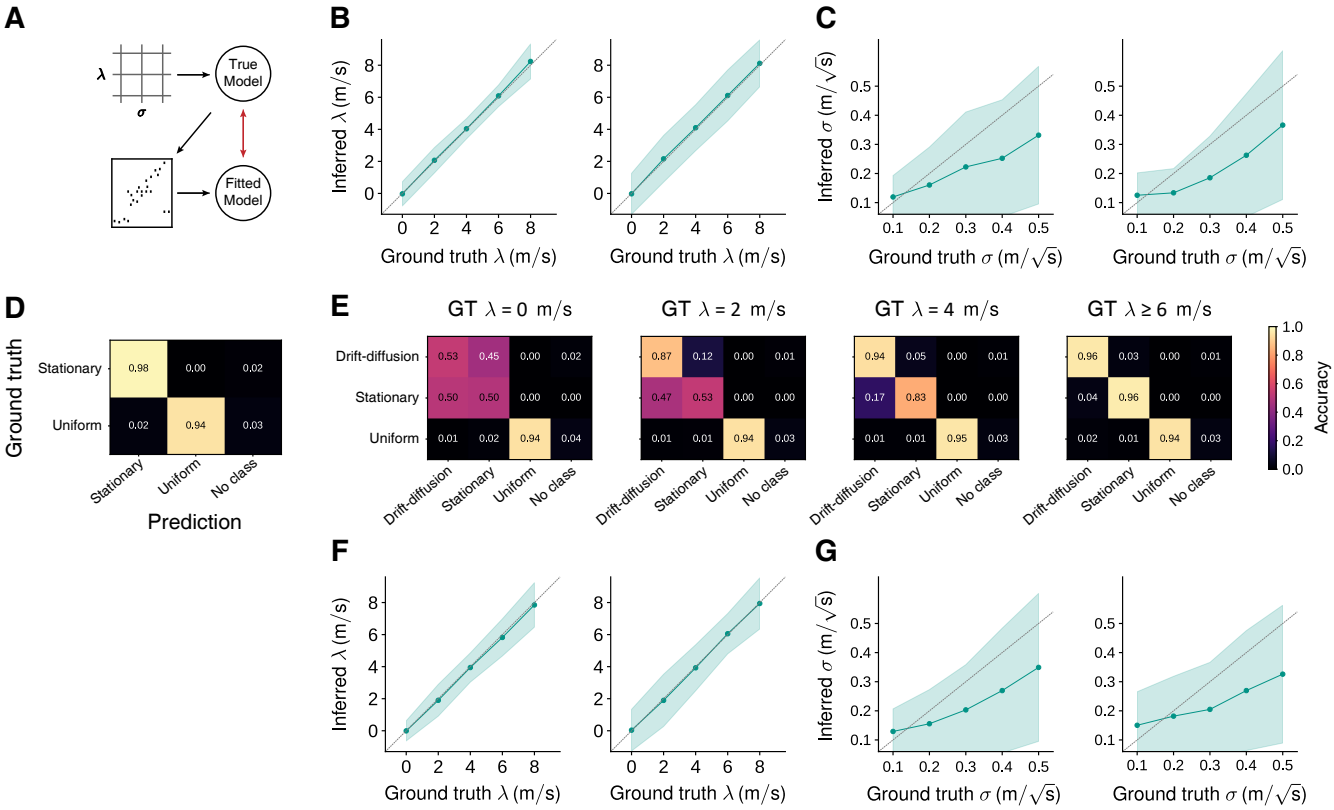

Figure 3: Identifiability of the model under realistic experimental conditions. **A**. Schematic diagram of the simulation pipeline. We considered two simulation settings: pure drift-diffusion (**B-C**), in which only drift-diffusion regime was included, and regime switching (**D-G**), in which the ground-truth dynamics could switch between drift-diffusion, stationary, and uniform regimes. In the pure drift-diffusion setting, ground-truth drift-diffusion parameter pairs (λ, σ) were first sampled on a grid. The sampled parameters determine the drift-diffusion dynamics which was used to generate multiple trajectories on a 2 m linear track. Each trajectory was then used to generate one 150 ms SWR event with population activity of N = 120 simulated place cells whose place fields uniformly tile the track. In the regime-switching setting, each SWR consisted of two concatenated 150 ms parts drawn from different regimes. Within each part, trajectories were generated according to the sampled regime. If the regime is drift-diffusion, the data generation followed the same procedure as in the pure drift-diffusion setting. If the regime is stationary or uniform, trajectories were generated using their corresponding dynamics and population activities were generated using the same place cell population. Each simulated SWR event was analyzed with our inference pipeline to obtain marginal posterior of dynamics and, if the ground-truth regime is drift-diffusion, parameter estimates $\hat{\lambda}$ and $\hat{\sigma}$. The marginal posterior of dynamics was classified into a discrete regime sequence $\{\hat{c}_t\}$ following the convention in Fig. 2B-H. The inferred $\hat{\lambda}$, $\hat{\sigma}$ and $\{\hat{c}_t\}$ were compared with the ground truth to quantify recovery performance. **B**. The parameter recovery of λ when σ is held at 0.2 (left) or 0.4 (right) $\mathrm{m}/\sqrt{\mathrm{s}}$. **C**. The parameter recovery of σ when λ is held at 0 (left) or 2 (right) m/s. Solid lines show the mean over 100 repeats and shaded bands span ± 1 standard deviation. The diagonal line denotes perfect recovery. **D**. Confusion matrix of regime classification accuracy for simulated SWRs when drift-diffusion regime is not involved. **E**. Confusion matrices of regime classification accuracy for simulated SWRs when drift-diffusion regimes with different ground-truth (GT) drift values are involved. From left to right, the ground truth drift is 0, 2, 4, ≥ 6 (m/s), respectively. Color intensity reflects the classification accuracy. **F**. The parameter recovery of λ when σ is held at 0.2 (left) or 0.4 (right) $\mathrm{m}/\sqrt{\mathrm{s}}$. **G**. The parameter recovery of σ when λ is held at 0 (left) or 2 (right) m/s. Solid lines show the mean over 100 repeats and shaded bands span ± 1 standard deviation. The diagonal line denotes perfect recovery.

We analyze a previously published dataset [30] that contains hundreds of simultaneously recorded neurons from the hippocampus when the rats navigated on a one-dimensional linear track (see Fig. 1A,B

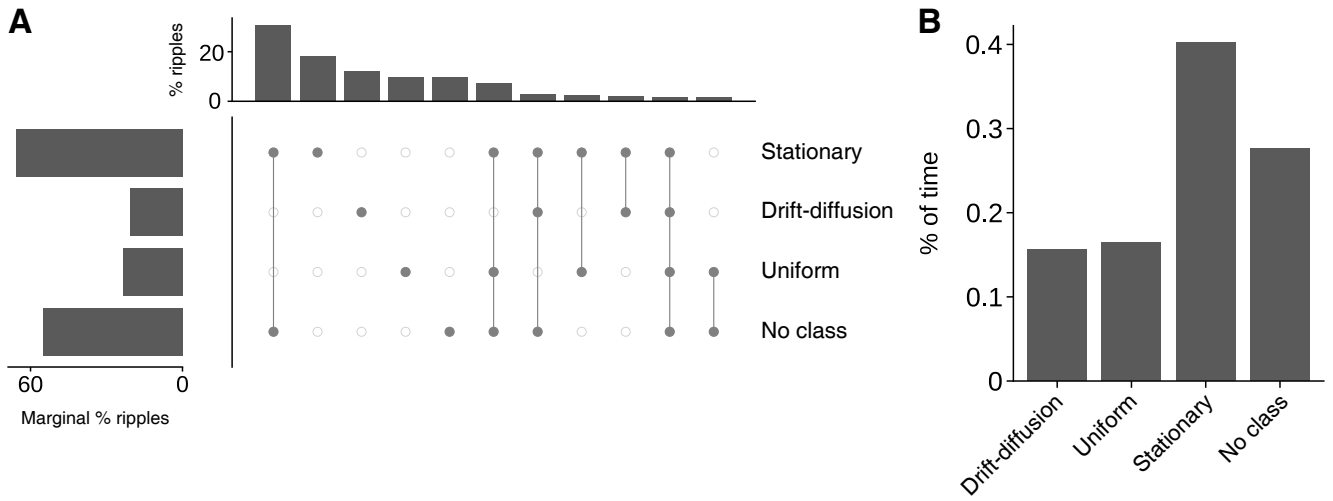

Figure 4: Composition of spatiotemporal dynamics of SWRs. For each SWR, the marginal posterior of dynamics was classified into a time series of regime class $\{\hat{c}_t\}$ following the convention in Fig. 2B-H. The time series was then partitioned into several segments, each defined as a period in which a single regime class persisted for $\geq$ 20 ms. **A.** The UpSet plot [46] summarizes the co-occurrence of the segments in all SWRs. Bottom: Dot–line matrix marks which segments are present in each combination. Filled circles denote the presence of a given segment in that combination, connected by lines. Top: Bars show the most frequent segment combinations in SWRs, ranked by the prevalence percentage. Left: Bars show the marginal percentage of SWRs that contain at least one segment of each regime class. **B.** Fraction of total SWR time assigned to each regime class.

for an overview of the task paradigm).

Application 1: Do most replay events exhibit real-world moving speed or even remain stationary?

Recent work suggests that a large portion of SWR events were either stationary or had a replay speed that is comparable to the moving speed of the animal ($<$ 1 m/s) [24]. This result seems to contrast with many previous reports that suggest that replay is compressed temporally, and its speed is often ten times faster than the animal's typical movement speed [3, 7, 48]. As our method leads to an unbiased characterization of the structure of the replay, it may shed light on this apparent discrepancy.

We fit our full model to the neural population responses during SWRs collected in [30] (a total of N=16,523 events from 4 rats). A summary of the results is shown in Fig. 4. We find that 12.2% of the events in this dataset can be classified as drift-diffusion dynamics only. In addition, about 6.7% of events are best described as a mixture of drift-diffusion and other types of dynamics. Thus, 18.9% of the events contain at least a segment of drift-diffusion period (i.e., at least 20 ms). This fraction is generally consistent with the previously reported fraction of replay events, analyzed using standard Bayesian decoding methods [4, 48]. Unlike prior methods, our method enables a more principled characterization of the potentially rich dynamics within individual events (see examples in Fig. 2B-H and Fig. 5A).

In light of the findings reported in [24], we next focus on the stationary dynamics in SWRs. We find that more than 60% of the SWR events contain at least a period of stationary dynamics. About one-third of

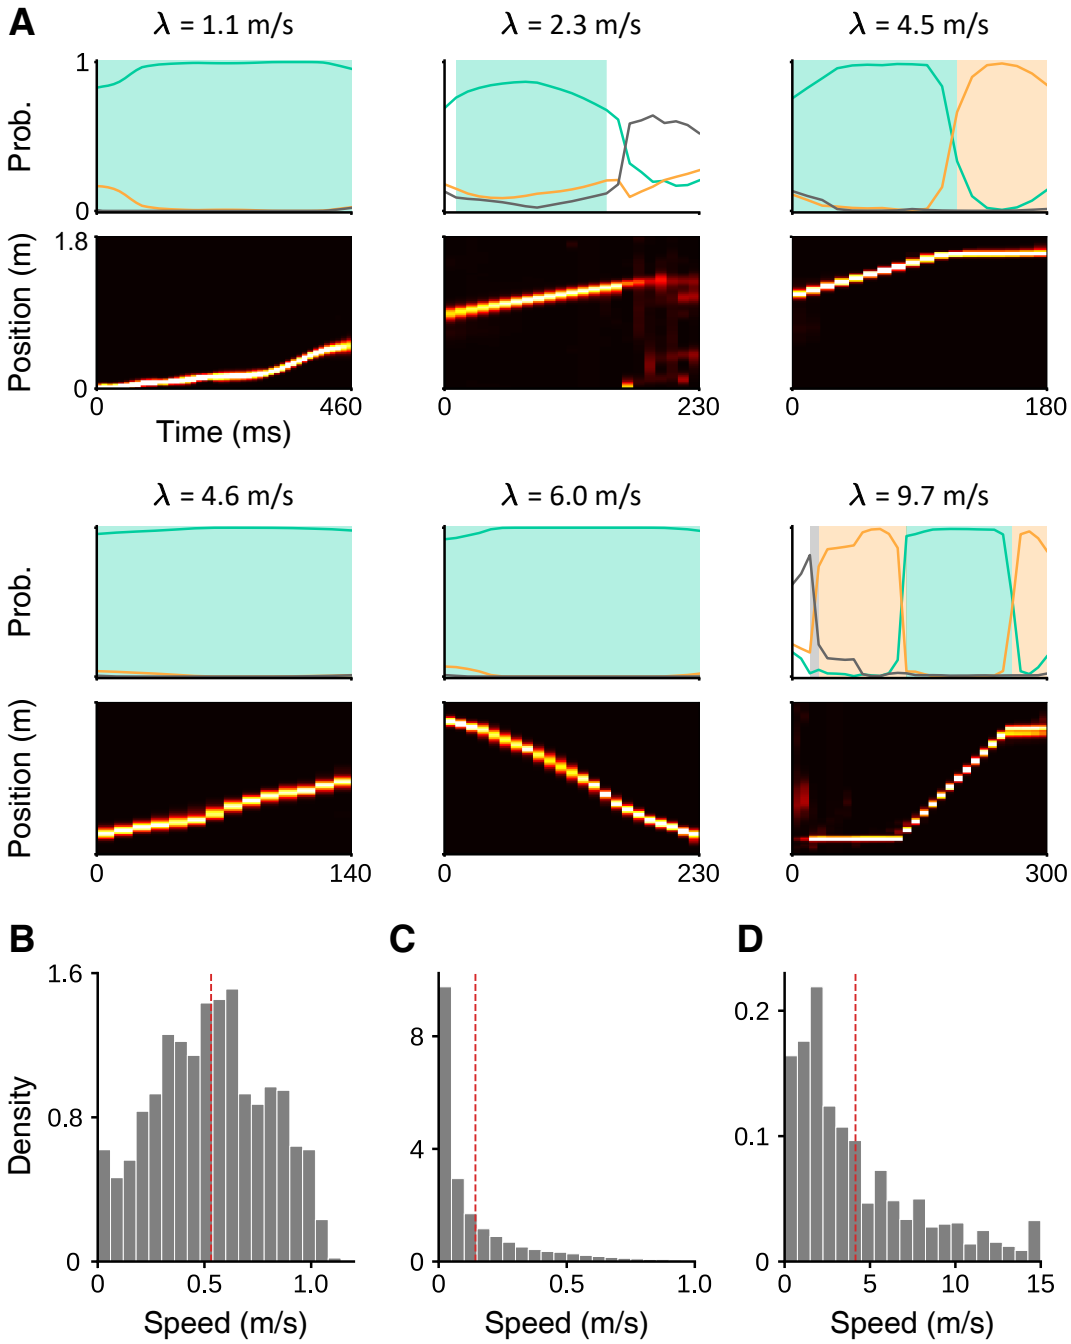

Figure 5: Speed distributions for running, behavior, and replay events. **A**. Example SWR events exhibiting different speeds. Plotting conventions follow Fig. 2B-H. **B**. Histogram of per-lap running speed, estimated by fitting a linear regression to the running trajectory of each lap on the linear track. **C**. Histogram of instantaneous behavioral speed, estimated using a finite-difference derivative of positions during running. **D**. Histogram of replay speed, estimated using the inferred drift rate λ. Only drift-diffusion events, namely SWRs whose marginal posterior of drift-diffusion regime exceeded 0.7 for at least 50 ms, are considered in **D**. The red dashed lines in **B-D** indicate the mean of each speed distribution: run, 0.53 m/s; behavior, 0.14 m/s; replay, 4.1 m/s.

these events are stationary only. The remaining two-thirds are mixed with other types of dynamics. We also computed the total fraction of time steps that are classified as stationary dynamics, and find that the stationary period accounts for 40.2% of all SWR events. One potential explanation of the stationary

dynamics during awake immobility is that the neural activity encodes the animal's current location despite the presence of high ripple power. Consistent with this hypothesis, we find that the replay events are concentrated around the animal's current location when SWRs occur. For 79.6% of the time, the decoded position is within 20 cm of the animal's current location (Fig. S4).

We further examine a subset of SWR events with marginal posterior of drift-diffusion dynamics that exceed 0.7 for at least 50 ms. These events are particularly interesting because they exhibit strong sequential structures and thus are consistent with the classic notion of replay events. These selected events account for 17.3% of the total events (N=2861 events). Fig. 5A shows example SWR events with varying replay speed. The distribution of the replay speed inferred for these events is shown in Fig. 5D, with a median of 4 m/s. The speed for these replay events is substantially larger than that of running (Fig. 5B) or movement periods (Fig. 5C). These results are insensitive to the exact criteria used to select these events. Thus, for the subset of SWR events that exhibit substantial drift-diffusion dynamics, the speed of the drift is one order of magnitude faster than the running speed, generally consistent with earlier reports on replay [8, 48].

Through an unbiased characterization of the replay dynamics, our results provide a way to reconcile the classic observation that shows a much higher speed for replay compared to running period and the more recent results that claimed most SWR events exhibit slow or stationary dynamics. On the one hand, classic results on time-compressed replay sequences were based on a subset of SWR events (generally 10%-20% of all SWR events) that exhibit substantial sequential structure, similar to the events we identified based on drift-diffusion dynamics. Consistent with earlier results, our analyses confirmed that the events with high sequentiality indeed exhibit a speed that is approximately an order of magnitude faster than the animal's typical running speed. On the other hand, the results in [24] emphasized a different type of events with stationary dynamics. Our results show that, out of 60% of the events contain a segment of stationary dynamics, about one-third of these are purely stationary events. While the speed of these events is generally much lower, these events are from a different subset of SWRs compared to those selected in the traditional analysis of replay events. Most of the stationary events during awake immobility are local, i.e., close to the animal's current location. Thus, the seemingly contradicting results could be primarily due to the largely non-overlapping sets of SWR events used in the analysis.

Application 2: Does hippocampal replay in 1-D environment resemble Brownian motion?

Although early studies of replay emphasize the linear-sequential structure of replays, we found that the decoded trajectories of many SWR events have substantial variability compared to the path of the animals' behavioral trajectories. Intriguingly, a recent study [31] on hippocampal activity in a 2-D open arena argues that replays resemble Brownian motion (but see [25]). Though this a conceptually provocative result, it remains unclear whether similar results hold in other datasets and experimental settings, particularly that of 1-D environments, where replays are typically studied.

One of our results reported above seems to contradict the proposition that Brownian motion generally

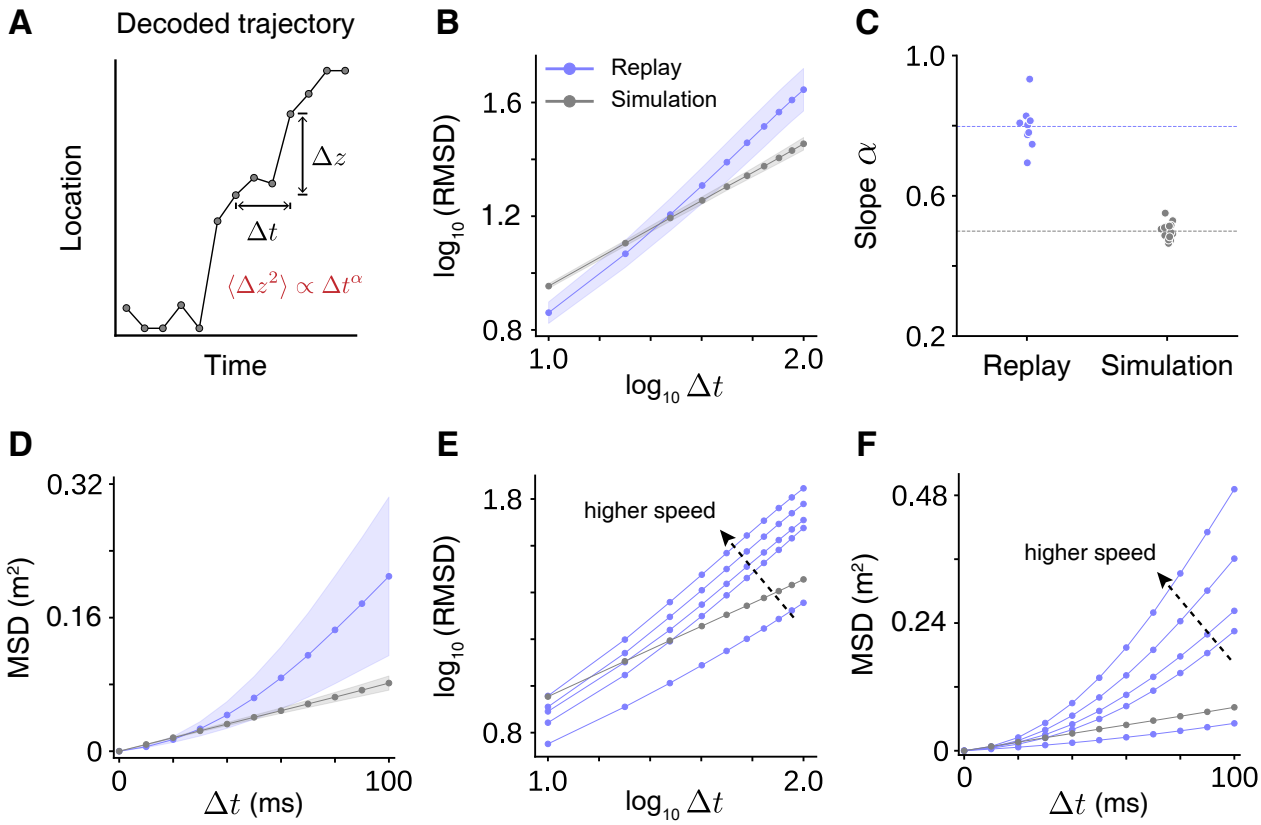

Figure 6: Replay trajectories differ from pure diffusion dynamics (Brownian motion), and are more consistent with drift-diffusion. **A**. Schematic diagram of the calculation of position displacement Δz, mean-square displacement (MSD) $\langle \Delta z^2 \rangle$, and root-mean-square displacement (RMSD) $\sqrt{\langle \Delta z^2 \rangle}$ for the time interval Δt from a replay trajectory. **B**. The log-log relationship between RMSD and Δt for replay (purple) and simulated (gray) trajectories. Note that replay trajectories were inferred under a fixed pure diffusion model ($\sigma = 0.8$ m/$\sqrt{\text{s}}$). Simulated replay trajectories were generated from the same pure diffusion model [25]. Solid lines and shaded bands indicate the mean and ± 1 standard deviation across sessions, respectively. **C**. Slopes of linear regression fit to log(RMSD)-log(Δt) data shown in **B**. Each dot represents an estimated slope from one session. Dashed horizontal lines indicate the mean across sessions: 0.79 for replay and 0.50 for simulation. **D**. The relationship between MSD and Δt for replay (purple) and simulated (gray) trajectories. Solid lines and shaded bands indicate the mean and ± 1 standard deviation across sessions. **E**. The log-log relationship between RMSD and Δt for replays with different estimated speed (purple) and simulations (gray). Replay events from all sessions were grouped based on five speed levels: $\lambda \in [0, 2)$, $[2, 4)$, $[4, 6)$, $[6, 8)$ and $[8, 10]$ m/s. The slopes of linear regression fits are 0.61, 0.85, 0.84, 0.88 and 0.90, respectively. Gray lines show the same analysis using simulated trajectories from the pure diffusion model, where the slope is 0.50. **F**. The relationship between MSD and Δt for replays with different estimated speed (purple) and simulations (gray). The definition of speed groups is identical to **E**. Gray lines show the same analysis using simulated trajectories from the pure diffusion model.

describes the dynamics of SWRs, that is, a considerable proportion of SWRs exhibit a substantial drift rate (Fig. 5B). However, it remains possible that the noise may cause the empirically estimated drift to substantially exceed zero even when sampling the trajectories from a pure diffusion process. To rule out this possibility, we generated a null distribution of the estimated speed from the diffusion process, by sampling and analyzing trajectories from diffusion models. We find that this null distribution is substantially different from the distribution obtained from the rat data (Fig. S3). In particular, the speeds obtained from the simulated diffusion process are substantially slower. These results suggest that the subset of drift-diffusion events identified by our method are unlikely to be characteristically generated by

Brownian motion.

To make the results directly comparable to prior work, we further analyzed the data using an approach generalized from [31], which examines how the root-mean-square displacement (RMSD) scales with the time interval (Fig. 6A). Their key idea is that the Brownian motion predicts that RMSD should increase linearly with the square root of time interval ($\sqrt{\Delta t}$), i.e., RMSD $\propto \sqrt{\Delta t}$. Empirically, Stella et al. [31] reported that the slope between these two quantities is reasonably close to 0.5 in a log-log plot. To evaluate this prediction independently, we generalize the previous analysis framework to more general drift-diffusion dynamics. Importantly, if the replay trajectory follows a drift-diffusion process, the mean-squared displacement (MSD) should be characterized by a quadratic rather than linear function, that is, MSD $= (\lambda\Delta t)^2 + \sigma^2\Delta t$. Under this view, Brownian motion is only expressed in the special case of zero drift, in which MSD $= \sigma^2\Delta t$.

In examining log-log plots of RMSD vs. Δt (limited to drift-diffusion events, namely events with drift-diffusion time lasting at least 50 ms), we observed a slope of approximately 0.8, substantially greater than 0.5. The results are consistent across nine sessions (mean=0.80, std=0.06) and deviate substantially from the predictions of Brownian motion (Fig. 6B,C). Importantly, to avoid biases toward finding systematic drifts, we used a fixed diffusion model (with $\sigma = 0.8$ m$/\sqrt{\mathrm{s}}$, zero drift) to decode the trajectories [25].

Our theoretical analysis reveals one important limitation of the approach based on the slope between RMSD and time interval in the log-log space. The issue is that, while the slope in the log-log space may be adequate to recover the exponent α for a power-law relation MSD $= C(\Delta t)^\alpha$, it is insufficient to test the more general prediction of drift-diffusion models in the form of MSD $= (\lambda\Delta t)^2 + \sigma^2\Delta t$. To overcome this limitation, we further examined how MSD scales with time interval in the original scale for drift-diffusion events. Consistent with the prediction of the drift-diffusion dynamics, we find that the relationship between MSD and time interval can be well described by a quadratic, rather than a linear, relationship (Fig. 6D). A further prediction of the drift-diffusion model is that the relationship between MSD and time interval should be systematically modulated by the drift rate. To test this, we sorted and separately analyzed the SWR events with different drift rates, finding an even stronger quadratic relationship for the events with higher speed (Fig. 6F). In the log-log plots, this yielded a slope that was closer to 1 for high speed events (Fig. 6E). These results provide strong empirical support for the drift-diffusion dynamics of replays.

A key advantage of our modeling approach is that it potentially enables identification and characterization of different types of replay dynamics. Indeed, our empirical results in Fig. 4 suggest that SWRs exhibit different types of dynamics. If these different types of events are indeed meaningful, one substantive prediction, generalizing the one suggested specifically for Brownian motion above, would be systematically different scaling relationships between MSD and time interval. To test this prediction, we additionally analyzed the subset of events that contain only stationary dynamics, as well as the subset of events that contain only uniform dynamics. Consistent with our prediction, for SWRs with stationary dynamics (Fig. S5A-C), the slope in the log-log plot is substantially smaller than 0.5 (mean=0.26,

std=0.13). Furthermore, for events with uniform dynamics (Fig. S5D-F), the slope is also smaller than 0.5 (mean=0.39, std=0.07). When all SWR events are combined, the resulting slope between RMSD and time interval in the log-log plots is 0.61, which is closer to, yet remains different from, the prediction of Brownian motion (Fig. S5G-I). Thus, these results strongly suggest that the scaling relationship between MSD and time is heterogeneous in SWR events with different types of dynamics.

Application 3: Is there evidence for preplay during SWRs?

It has been suggested that the hippocampal "preplay" sequences of neural activity before the animal has any experience in a given environment [47, 49]. While preplay potentially has important implications for the mechanisms and functions of hippocampus, the existence and prevalence of preplay remain controversial [44]. This is in part due to the difficulty of precisely characterizing the dynamics of the SWR events. Using our method, we sought to re-examine the evidence for preplay, and if preplay exists, how its dynamics compare with replay dynamics after spatial experience in a given environment.

We analyzed the SWRs during pre-run sleep (data from ref [30], a total of N=6,749 events from 4 rats). We first quantified the composition of dynamics (see Fig. 7A-B), and compared the results to those shown in Fig. 4 for replay. For pre-run sleep, drift-diffusion dynamics accounted for 5.8% of total SWR time. This is substantially lower than in post-run sleep (15.6%). At the event level, 5.6% of pre-run SWRs can be classified as drift-diffusion dynamics only and 6.3% contained at least one drift-diffusion segment. These numbers are again much smaller than those from post-run (12.2% and 18.9%, respectively). Thus, when aggregated across events, pre-run SWRs exhibit less sequential dynamics than post-run SWRs.

To assess whether the observed fraction of preplay events is significantly higher than chance levels, we conducted three shuffle analyses on preplay using data from pre-run sleep periods [44]. First, time-bin shuffle permutes the spike raster of each SWR across time bins, preserving per-neuron spike counts while disrupting within-event temporal structure. Second, neuron-identity shuffle randomly shuffles spike trains across neurons, preserving the spiking profile of each neuron but breaking the neuron-position mapping. Third, place-field rotation shuffle circularly shifts the place field of each neuron by an independent random offset. For comparison, we also performed these three shuffle analyses on post-run sleep SWRs. Taking advantage of our drift-diffusion framework, we used the duration of drift-diffusion dynamics within an event as a metric to characterize the sequentiality of an event (Fig. 7C-H). Across a range of thresholds, we computed the fraction of SWR events that exceeded each threshold to assess statistical significance. We find that the results diverge across the three shuffle methods. Based on either neuron-identity shuffle or place-field rotation shuffle, we did not find evidence for preplay, consistent with results reported in [44]. However, for time-bin shuffle, the fraction of events that exhibited sequentiality was higher than the chance level, consistent with [49]. These results may help reconcile the previous discrepancy on the existence of preplay from studies that used different shuffle procedures [44, 49]. It is also interesting to note that the time-bin shuffle has been suggested to be relatively less conservative than the neuron-identity and place-field shuffles [19], which also underscores the finding of significant preplay structure as surprising.

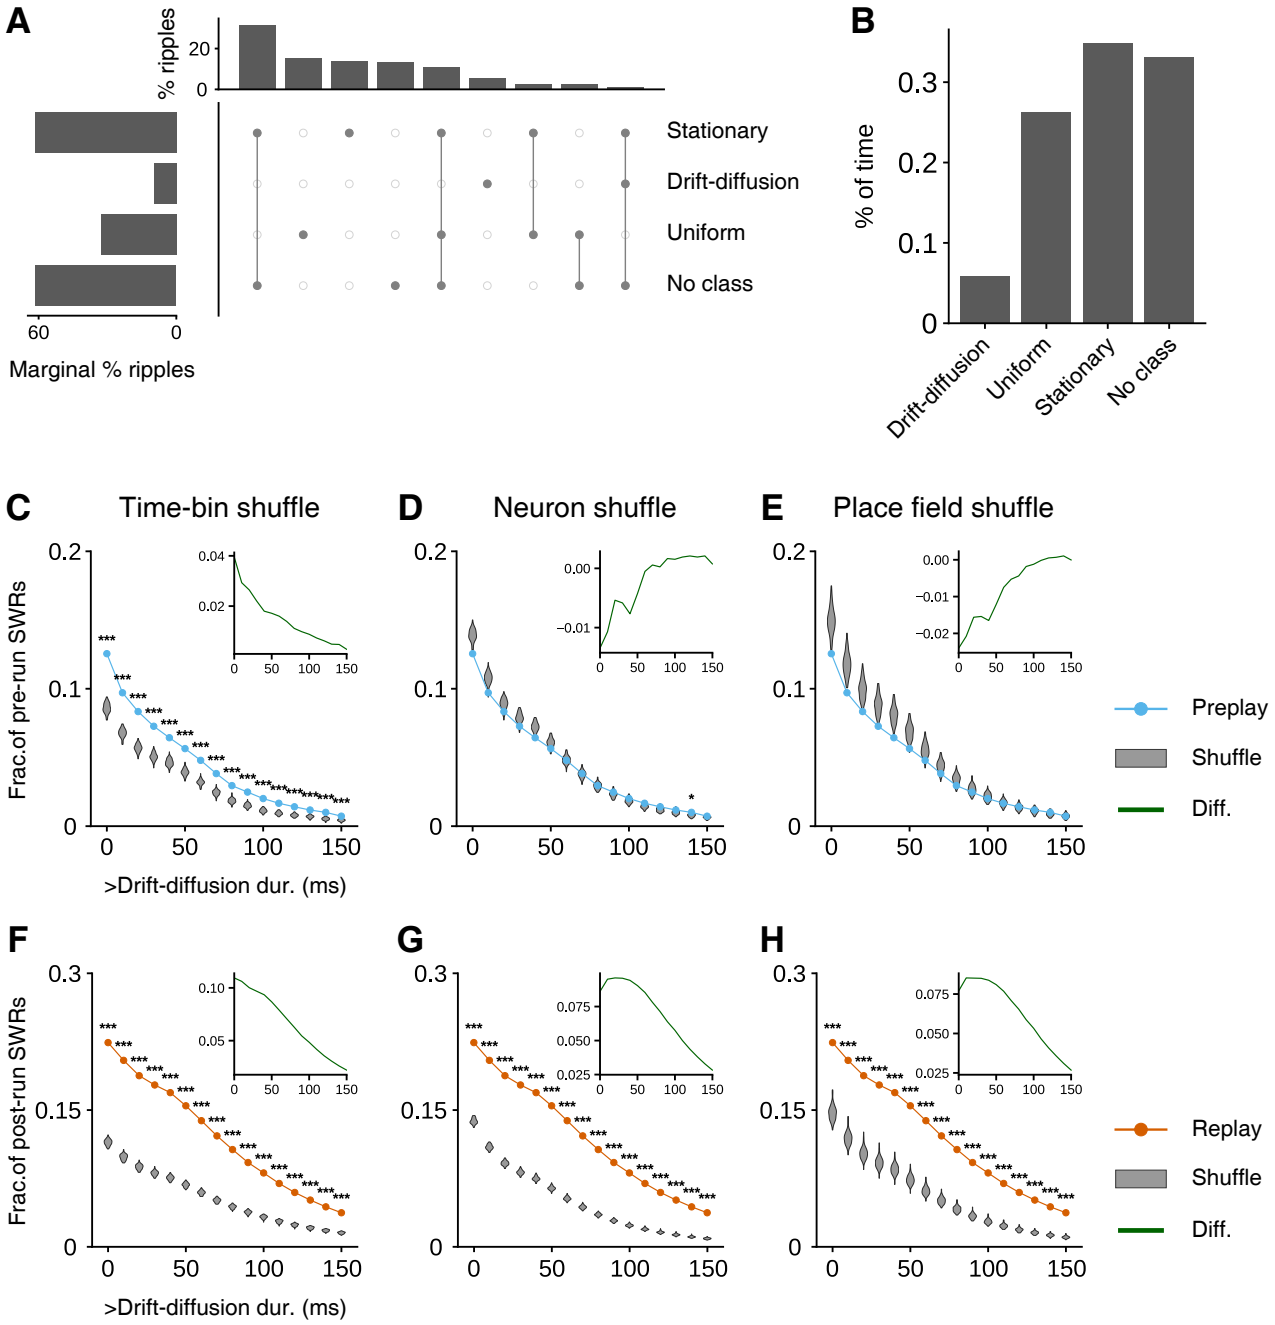

Figure 7: Preplay events show substantially weaker temporal structure than replay. A-B. Composition of spatiotemporal dynamics of pre-run SWRs. Plotting conventions follow Fig. 4. Compared to the post-run, pre-run SWRs showed a markedly lower occurrence of drift-diffusion dynamics and a higher occurrence of uniform dynamics. **C-H.** Each plot shows the fraction of SWRs that have at least the minimum drift-diffusion duration specified on the x-axis. Colored lines indicate the fraction of SWRs in data. Gray violins indicate the distribution of fractions from 120 shuffles per SWR. Asterisks mark the significance of the fraction in data being higher than the shuffle (one-sided permutation test, *** $p < 0.001$ and * $p < 0.05$). Inset shows the difference across thresholds between fraction in data and mean fraction in shuffle. Positive value indicates a higher fraction than expected under shuffle and negative value indicates lower. Rows correspond to two subsets of SWRs: pre-run SWRs (**C, D, E**) and post-run SWRs (**F, G, H**). Columns correspond to three shuffle controls: time-bin (**C, F**), neuron-identity (**D, G**), and place-field rotation (**E, H**).

When comparing pre-run and post-run SWRs, we consistently find that, across the entire threshold range, the fraction of events exceeding any given threshold was higher for post-run than pre-run SWRs. This suggests that at the level of individual events, pre-run SWRs also exhibited shorter drift-diffusion segments than post-run SWRs, mirroring the pooled result in Fig. 4. We also compared the distributions of drift-diffusion durations against their corresponding null distributions based on shuffle (Fig. S6). Post-run SWRs showed differences from all three controls with strong significance ($p < 0.001$ for each; one-sided permutation test). In contrast, pre-run SWRs differed only from the time-bin shuffle ($p < 0.001$; one-sided permutation test) and not significantly from the neuron-identity or place-field rotation shuffles. These results suggest that replay after experience can be detected robustly, while preplay is detectable only when neuron-position mappings are preserved, as in time-bin shuffle, but not in neuron-identity and place-field shuffles.

The results above suggest that hippocampal population activity during pre-run SWRs is coordinated between potential (unexperienced) spatial locations at a level that is higher than chance. We next examined whether such coordination happens at a timescale that is comparable to the classic replay events (typically around 100 ms). By taking the difference of the distribution of drift-diffusion durations between real data and shuffles (Fig. 7C-H inset), we obtained an estimate of the fraction of *bona fide* drift-diffusion events with relatively long durations. According to this metric, for post-run SWRs, the fraction of events that exhibit drift-diffusion longer than 100 ms is about 4.85% according to time-bin shuffle. For preplay, this fraction is 0.88%. Thus, hippocampal population activity during pre-run SWRs exhibits slightly higher sequentiality than time-bin shuffled data, yet long temporal sequences are particularly rare.

Together, these results provide evidence for pre-existing structures in place-cell population activity before rats have spatial experience in a given environment. However, the level of coordination in these structures is substantially weaker than that seen in replay. In particular, we found that long sequences during pre-run SWRs are rare, much weaker than replay sequences after spatial experience.

3 Discussion

We have developed a new framework, conceptually and statistically, for analyzing and understanding the dynamics of replays. This framework is formulated based on drift-diffusion dynamics while allowing the switching between several types of dynamics. Drift-diffusion models have been widely used to model the neural basis of decision-making [50–52]. Here we show that the drift-diffusion framework also provides a powerful framework to study the dynamics of the hippocampus. Compared to standard methods for analyzing replays, there are several important differences that make our method advantageous. First, using a generative modeling approach, our method enables principled probabilistic inference on the structures of replays that fully takes into account the uncertainty in estimation. Second, our framework explicitly models speed and quality of classic replay events (i.e., drift-diffusion events) with drift and diffusion parameters. Third, with the capability of switching between different types of dynamics, our

method is able to characterize the putatively rich dynamics in a broader set of hippocampal SWRs, which continues to remain an important open question. Evaluation and comparison of our method to others reveal that our method allows for more direct and unambiguous interpretation of replay dynamics.

Application of our modeling framework to a large-scale hippocampal dataset revealed several insights that helped resolve controversies about replays. First, our method leads to an unbiased characterization of the dynamics during SWR events. Consistent with the results reported in [24], we found that many SWR events contain segments of stationary dynamics. During awake immobility, the stationary segments of SWRs are mostly local, i.e., the decoded location is close to the animal's location on the track. One interpretation of this result is that place cells during these time steps in SWRs encode the animal's current location. Importantly, we also find that many SWR events do not appear to correspond to mixtures of stationary dynamics and random sampling [24] (Fig. 4A). In particular, for a subset of SWR events that exhibit consistent sequential structures, the speed is about 10 times faster than the animal's running speed, and these events account for 10-20% of the SWRs, which is consistent with the classic replay sequences during SWRs. Thus, our results partially reconcile the results in [24] and classic studies of replay. These results also highlight the rich and diverse types of neural dynamics during SWRs. The role that different types of neural dynamics play in learning and memory remains unclear–an important question that needs to be determined in future work.

Second, our results clarify the dynamics of replay events, at least in 1-D linear tracks. By carefully analyzing the scaling relationship between the MSD and time interval, we found that the two follow a quadratic relationship for events with high sequentiality. This result is highly consistent with the prediction of drift-diffusion dynamics, but not with the Brownian motion as proposed in [31]. SWR events with higher speed lead to a more substantial divergence from Brownian motion. We also found that different types of events exhibited different scaling relationships. For stationary events and uniform events, the relationship is sublinear. These results provide further support for the distinctness of different types of dynamics identified using our method. Our analysis suggests that it is more informative to analyze the scaling relationship between MSD and time interval in the original scale, rather than the previously proposed log-log scale [31]. While our results are based on data collected in 1-D environments, they motivate revisiting or additional exploration of replay dynamics in 2-D, or in additional environment and behavioral settings. In particular, it would be interesting to examine or re-examine the 2-D data to see whether a general quadratic relation also holds in 2-D for the replay events they identified. It would also be important to examine whether there is substantial heterogeneity in the dynamics of SWR events in 2-D. One possibility is that, for SWRs in 2-D environment, the dynamics are also heterogeneous as in 1-D, yet the fractions of SWR events with different types of dynamics differ from those in 1-D. When combining all events together, this may result in an overall slope close to 0.5 in the log-log scale. Regardless, our results clearly demonstrate that the dynamics of SWRs are diverse, and cannot be adequately described by Brownian motion in 1-D environments.

Third, we find evidence supporting the existence of weak pre-existing temporal structure before spatial experience in a given environment, consistent with the idea of preplay [47]. However, these pre-existing structures appear to be detected with particular analysis approaches–in particular, we found population

structure beyond randomization only through time-based shuffle. Furthermore, our results reveal that long drift-diffusion sequences that are comparable to classic replay events with a timescale of 100 ms are exceptionally rare (less than 1% of SWRs in pre-run sleep). Even with these results, caution is in order, as time-based shuffle may eliminate temporal structure in individual neurons, thus it may provide a weaker null distribution. Furthermore, given that the animals may have been exposed to environments with somewhat similar spatial structure, sequences observed during pre-run SWRs may still reflect a type of replay of analogous experiences [9, 53, 54]. Exactly what properties of the pre-existing hippocampal circuits lead to the higher-than-chance sequential structure during pre-run sleep remains to be determined. One possibility is that the structure during pre-run sleep is caused by the cellular correlation between "rigid" place cells that is preserved across environments [55].

Earlier studies have used state-space models to decode positions and reconstruct "mental" trajectories from hippocampal activity during running and sleep, although they did not explicitly characterize the spatiotemporal structure of replay [27–29, 36, 37]. More recently, state-space models have been applied to explicitly model dynamics during SWRs [24, 25]. In particular, [24] proposed a method for analyzing hippocampal replay by considering switching dynamics during decoding. There are two crucial differences between the method in [24] and the method presented here. First, the parameters of their method were manually specified rather than inferred from data. Second, their method does not model drift, in part as a consequence of the manual specification. Instead, it effectively solely relies on Brownian motion as its spatially "continuous" dynamics to capture the sequential structure of replay events. As we have shown, Brownian motion is inadequate for capturing sequential replay dynamics. These issues are liable to introduce biases in estimating dynamics and decoding errors in estimating positions. For example, the replay speeds could be underestimated, if the replay in neural data has a substantial drift, which does not conform to the manually specified structure in the model. In contrast, our method directly estimates two essential model parameters from individual replay events and is thereby more flexible. Another recent study [25] formulated HMM models to study transition dynamics in 2-D open-field environments. They found that the momentum of trajectories is important for fitting replay events. One important difference is that their model does not allow for the switching between different types of dynamics. Also, their model is a second-order HMM, our model is first-order. From what we have shown, a first-order switching HMM appears sufficient to capture the key dynamical patterns of replay. Moreover, its dynamics and parameters (e.g., drift and diffusion) map directly onto first-order transition statistics, leading to a more transparent interpretation of how replay trajectories evolve over time. Our first-order formulation also simplifies the probabilistic inference procedure and potentially improves parameter identifiability, leading to more stable and accurate characterizations of replay dynamics.

Extensions and additional applications of the proposed modeling framework are possible. First, our current implementation primarily deals with 1-D track or circular environment, two popular environments that have been used to study hippocampal replay. Other types of environments have also been used in studying replay, such as 2-D open-field arenas and 1-D environments with more complex geometry. Adaptation of our method to deal with these environments would require modifications to the probabilistic

transition matrix for latent positions. One simple generalization of our method to 2-D environments would be to assume that dynamics in 2-D factorize into two individual dimensions, and apply our method to each dimension separately. Second, another natural extension would be a switching drift-diffusion model that switches between multiple drift-diffusion processes with distinct parameters. However, in practice, the limited data per SWR event may prevent one from estimating the parameters reliably. In this paper, we therefore adopt a more parsimonious model by explicitly including two stereotyped dynamics (i.e., stationary and uniform sampling) alongside a single drift-diffusion component. This choice preserves tractability while maintaining flexibility to capture the spatiotemporal heterogeneity of individual SWRs. Crucially, by allowing stationary and uniform regimes in the model, the method can isolate SWR segments with substantial drift and spatial coherence. Thus, the duration of drift-diffusion segment for individual events provides a measure of the strength of sequential replay. Third, while we only focus on replay sequences during SWRs, our method is in principle applicable to the study of theta sequences in rodents during running [56, 57]. Theta sequences represent another important type of dynamical structure in the hippocampus [58–60] and entorhinal cortex [61], which may have equally important functional implications for the role of hippocampus or other areas in cognition. Previous analysis typically relied on analyzing the averaged slope of the theta sequences. Using our method, one may potentially reveal the sequentiality and variability of the dynamics within individual theta sequences.

A better understanding of replay dynamics in the brain may help inform the development of better learning algorithms in artificial intelligence. Algorithms based on replaying past experiences have been developed [62] and shown to improve performance across a variety of tasks in reinforcement learning [63, 64]. Yet most current AI algorithms still struggle in continual learning without catastrophic forgetting [65–67]. Developing brain-inspired algorithms to replay past experiences efficiently [68, 69], or to sample from a generative model built on past experiences (i.e., generative replay) [70–72], remains a promising path forward [73, 74]. More accurate and interpretable data-driven techniques, such as the one presented here, will likely be instrumental toward clarifying the exact role of replay in natural cognition, thereby informing brain-inspired replay mechanisms for continual learning in AI.

Hippocampal SWRs have direct implications in cognitive functions. It was reported that disrupting SWRs in rodents impairs spatial memory in an AD-developmental mouse model [75] or spatial learning in wild-type [76], whereas enhancing SWRs may improve subsequent task performance [77]. Recent work shows that in a mouse model of Alzheimer's disease, the reactivation of hippocampal activity during SWRs was disrupted [78]. Most prior work used simple metrics such as the frequency of SWRs. Using our method to analyze the detailed dynamics during different task conditions may reveal new insights into the relationship between the properties of SWRs and behavior. Finally, adapting our method to the online setting may enable the design of more targeted, dynamics-specific, SWR-based neurofeedback protocols for cognitive perturbation and enhancement [79, 80].

Acknowledgments

We would like to thank Brad Pfeiffer and David Foster for generously sharing their data. We thank Laura Colgin and Kenneth Kay for helpful discussions and their comments on an earlier version of the paper. This research is supported by NSF 2318065, and a Sloan Research Fellowship to X.-X. Wei. The funders had no role in designing and conducting this research, decision to publish, or preparation of the manuscript.

References

[1] Brad E Pfeiffer and David J Foster. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79, 2013.

[2] Kamran Diba and György Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241–1242, 2007.

[3] David J Foster and Matthew A Wilson. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680–683, 2006.

[4] R Ellen Ambrose, Brad E Pfeiffer, and David J Foster. Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron*, 91(5):1124–1136, 2016.

[5] Matthew A Wilson and Bruce L McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994.

[6] William E Skaggs and Bruce L McNaughton. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271(5257):1870–1873, 1996.

[7] Zoltán Nádasdy, Hajime Hirase, András Czurkó, Jozsef Csicsvari, and György Buzsáki. Replay and time compression of recurring spike sequences in the hippocampus. *Journal of neuroscience*, 19(21):9497–9507, 1999.

[8] Albert K Lee and Matthew A Wilson. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6):1183–1194, 2002.

[9] Mattias P Karlsson and Loren M Frank. Awake replay of remote experiences in the hippocampus. *Nature neuroscience*, 12(7):913–918, 2009.

[10] Adrien Peyrache, Mehdi Khamassi, Karim Benchenane, Sidney I Wiener, and Francesco P Battaglia. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature neuroscience*, 12(7):919–926, 2009.

[11] Justin D Shin, Wenbo Tang, and Shantanu P Jadhav. Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. *Neuron*, 104(6):1110–1125, 2019.

[12] Karola Kaefer, Michele Nardin, Karel Blahna, and Jozsef Csicsvari. Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron*, 106(1):154–165, 2020.

[13] Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature neuroscience*, 10(1):100–107, 2007.

[14] Yifan Yang, David A Leopold, Jeff H Duyn, and Xiao Liu. Hippocampal replay sequence governed by spontaneous brain-wide dynamics. *PNAS nexus*, 3(4):pgae078, 2024.

[15] Anna C Schapiro, Elizabeth A McDevitt, Timothy T Rogers, Sara C Mednick, and Kenneth A Norman. Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature communications*, 9(1):3920, 2018.

[16] Yitzhak Norman, Erin M Yeagle, Simon Khuvis, Michal Harel, Ashesh D Mehta, and Rafael Malach. Hippocampal sharp-wave ripples linked to visual episodic recollection in humans. *Science*, 365(6454):eaax1030, 2019.

[17] Yunzhe Liu, Raymond J Dolan, Zeb Kurth-Nelson, and Timothy EJ Behrens. Human replay spontaneously reorganizes experience. *Cell*, 178(3):640–652, 2019.

[18] Yunzhe Liu, Marcelo G Mattar, Timothy EJ Behrens, Nathaniel D Daw, and Raymond J Dolan. Experience replay is associated with efficient nonlocal learning. *Science*, 372(6544):eabf1357, 2021.

[19] David J Foster. Replay comes of age. *Annual review of neuroscience*, 40(1):581–602, 2017.

[20] Laura Lee Colgin. Five decades of hippocampal place cells and eeg rhythms in behaving rats. *Journal of Neuroscience*, 40(1):54–60, 2020.

[21] Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. The hippocampus as a predictive map. *Nature neuroscience*, 20(11):1643–1653, 2017.

[22] Zhe Sage Chen and Matthew A Wilson. How our understanding of memory replay evolves. *Journal of Neurophysiology*, 129(3):552–580, 2023.

[23] David Tingley and Adrien Peyrache. On the methods for reactivation and replay analysis. *Philosophical Transactions of the Royal Society B*, 375(1799):20190231, 2020.

[24] Eric L Denovellis, Anna K Gillespie, Michael E Coulter, Marielena Sosa, Jason E Chung, Uri T Eden, and Loren M Frank. Hippocampal replay of experience at real-world speeds. *Elife*, 10:e64505, 2021.

[25] Emma L Krause and Jan Drugowitsch. A large majority of awake hippocampal sharp-wave ripples feature spatial trajectories with momentum. *Neuron*, 110(4):722–733, 2022.

[26] Daniel C McNamee, Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. Flexible modulation of sequence generation in the entorhinal–hippocampal system. *Nature neuroscience*, 24(6):851–862, 2021.

[27] Riccardo Barbieri, Matthew A Wilson, Loren M Frank, and Emery N Brown. An analysis of hippocampal spatio-temporal representations using a bayesian algorithm for neural spike train decoding. *IEEE transactions on neural systems and rehabilitation engineering*, 13(2):131–136, 2005.

[28] Marc Box, Matt W Jones, and Nick Whiteley. A hidden markov model for decoding and the analysis of replay in spike trains. *Journal of computational neuroscience*, 41(3):339–366, 2016.

[29] Kourosh Maboudi, Etienne Ackermann, Laurel Watkins de Jong, Brad E Pfeiffer, David Foster, Kamran Diba, and Caleb Kemere. Uncovering temporal structure in hippocampal output patterns. *Elife*, 7:e34467, 2018.

[30] Brad E Pfeiffer and David J Foster. Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science*, 349(6244):180–183, 2015.

[31] Federico Stella, Peter Baracskay, Joseph O'Neill, and Jozsef Csicsvari. Hippocampal reactivation of random trajectories resembling brownian diffusion. *Neuron*, 102(2):450–461, 2019.

[32] Kechen Zhang, Iris Ginzburg, Bruce L McNaughton, and Terrence J Sejnowski. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of neurophysiology*, 79(2):1017–1044, 1998.

[33] Xiaojing Wu and David J Foster. Hippocampal replay captures the unique topological structure of a novel environment. *Journal of Neuroscience*, 34(19):6459–6469, 2014.

[34] Leonard E Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *The annals of mathematical statistics*, 41(1):164–171, 1970.

[35] Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 2002.

[36] Zhe Chen, Fabian Kloosterman, Emery N Brown, and Matthew A Wilson. Uncovering spatial topology represented by rat hippocampal population neuronal codes. *Journal of computational neuroscience*, 33(2):227–255, 2012.

[37] Zhe Chen. An overview of bayesian methods for neural spike train analysis. *Computational intelligence and neuroscience*, 2013(1):251905, 2013.

[38] Scott W Linderman, Matthew J Johnson, Matthew A Wilson, and Zhe Chen. A bayesian nonparametric approach for uncovering rat hippocampal population codes during spatial navigation. *Journal of neuroscience methods*, 263:36–47, 2016.

[39] Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial intelligence and statistics*, pages 914–922. PMLR, 2017.

[40] Joshua Glaser, Matthew Whiteway, John P Cunningham, Liam Paninski, and Scott Linderman. Recurrent switching dynamical systems models for multiple interacting neural populations. *Advances in neural information processing systems*, 33:14867–14878, 2020.

[41] Felix Pei, Joel Ye, David Zoltowski, Anqi Wu, Raeed H Chowdhury, Hansem Sohn, Joseph E O'Doherty, Krishna V Shenoy, Matthew T Kaufman, Mark Churchland, et al. Neural latents benchmark'21: evaluating latent variable models of neural population activity. *arXiv preprint arXiv:2109.04463*, 2021.

[42] Carles Balsells-Rodas, Yixin Wang, and Yingzhen Li. On the identifiability of switching dynamical systems. *arXiv preprint arXiv:2305.15925*, 2023.

[43] Yonghong An, Yingyao Hu, Johns Hopkins, and Matt Shum. Identifiability and inference of hidden markov models. Technical report, Technical report, 2013.

[44] Delia Silva, Ting Feng, and David J Foster. Trajectory events across hippocampal place cells require previous experience. *Nature neuroscience*, 18(12):1772–1779, 2015.

[45] Masahiro Takigawa, Marta Huelin Gorriz, Margot Tirole, and Daniel Bendor. Evaluating hippocampal replay without a ground truth. *ELife*, 13:e85635, 2024.

[46] Alexander Lex, Nils Gehlenborg, Hendrik Strobelt, Romain Vuillemot, and Hanspeter Pfister. Upset: visualization of intersecting sets. *IEEE transactions on visualization and computer graphics*, 20(12):1983–1992, 2014.

[47] George Dragoi and Susumu Tonegawa. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401, 2011.

[48] Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, 2009.

[49] Usman Farooq, Jeremie Sibille, Kefei Liu, and George Dragoi. Strengthened temporal coordination within pre-existing sequential cell assemblies supports trajectory replay. *Neuron*, 103(4):719–733, 2019.

[50] Michael N Shadlen and William T Newsome. Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of neurophysiology*, 86(4):1916–1936, 2001.

[51] Jamie D Roitman and Michael N Shadlen. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, 22(21):9475–9489, 2002.

[52] Joshua I Gold and Michael N Shadlen. The neural basis of decision making. *Annu. Rev. Neurosci.*, 30(1):535–574, 2007.

[53] Wenbo Tang, Justin D Shin, Loren M Frank, and Shantanu P Jadhav. Hippocampal-prefrontal reactivation during learning is stronger in awake compared with sleep states. *Journal of Neuroscience*, 37(49):11789–11805, 2017.

[54] Justin D Shin, Wenbo Tang, and Shantanu P Jadhav. Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. *Neuron*, 104(6):1110–1125, 2019.

[55] Andres D Grosmark and György Buzsáki. Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, 351(6280):1440–1443, 2016.

[56] George Dragoi and György Buzsáki. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron*, 50(1):145–157, 2006.

[57] David J Foster and Matthew A Wilson. Hippocampal theta sequences. *Hippocampus*, 17(11):1093–1099, 2007.

[58] Chenguang Zheng, Ernie Hwaun, Carlos A Loza, and Laura Lee Colgin. Hippocampal place cell sequences differ during correct and error trials in a spatial memory task. *Nature communications*, 12(1):3373, 2021.

[59] Andrew M Wikenheiser and A David Redish. Hippocampal theta sequences reflect current goals. *Nature neuroscience*, 18(2):289–294, 2015.

[60] Kenneth Kay, Jason E Chung, Marielena Sosa, Jonathan S Schor, Mattias P Karlsson, Margaret C Larkin, Daniel F Liu, and Loren M Frank. Constant sub-second cycling between representations of possible futures in the hippocampus. *Cell*, 180(3):552–567, 2020.

[61] Abraham Z Vollan, Richard J Gardner, May-Britt Moser, and Edvard I Moser. Left–right-alternating theta sweeps in entorhinal–hippocampal maps of space. *Nature*, 639(8056):995–1005, 2025.

[62] Long-Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3):293–321, 1992.

[63] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[64] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[65] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019.

[66] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021.

[67] Gido M Van de Ven, Tinne Tuytelaars, and Andreas S Tolias. Three types of incremental learning. *Nature Machine Intelligence*, 4(12):1185–1197, 2022.

[68] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.

[69] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in neural information processing systems*, 32, 2019.

[70] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017.

[71] Gido M Van de Ven, Hava T Siegelmann, and Andreas S Tolias. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 11(1):4069, 2020.

[72] Rui Gao and Weiwei Liu. Ddgr: Continual learning with deep diffusion-based generative replay. In *International Conference on Machine Learning*, pages 10744–10763. PMLR, 2023.

[73] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020.

[74] Lennart Wittkuhn, Samson Chien, Sam Hall-McMaster, and Nicolas W Schuck. Replay in minds and machines. *Neuroscience & Biobehavioral Reviews*, 129:367–388, 2021.

[75] Emily A Jones, Anna K Gillespie, Seo Yeon Yoon, Loren M Frank, and Yadong Huang. Early hippocampal sharp-wave ripple deficits predict later learning and memory impairments in an alzheimer's disease mouse model. *Cell reports*, 29(8):2123–2133, 2019.

[76] Valérie Ego-Stengel and Matthew A Wilson. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20(1):1–10, 2010.

[77] Antonio Fernández-Ruiz, Azahara Oliva, Eliezyer Fermino de Oliveira, Florbela Rocha-Almeida, David Tingley, and György Buzsáki. Long-duration hippocampal sharp wave ripples improve memory. *Science*, 364(6445):1082–1086, 2019.

[78] Sarah Shipley, Marco P Abrate, Robin Hayman, Dennis Chan, and Caswell Barry. Disordered hippocampal reactivations predict spatial memory deficits in a mouse model of alzheimer's disease. *bioRxiv*, pages 2024–11, 2024.

[79] Igor Gridchyn, Philipp Schoenenberger, Joseph O'Neill, and Jozsef Csicsvari. Assembly-specific disruption of hippocampal replay leads to selective memory deficit. *Neuron*, 106(2):291–300, 2020.

[80] Michael E Coulter, Anna K Gillespie, Joshua Chu, Eric L Denovellis, Trevor Thai K Nguyen, Daniel F Liu, Katherine Wadhwani, Baibhav Sharma, Kevin Wang, Xinyi Deng, et al. Closed-loop modulation of remote hippocampal representations with neurofeedback. *Neuron*, 113(6):949–961, 2025.

4 Methods

4.1 Description of the datasets

We used electrophysiological recordings in Pfeiffer and Foster (2015) [30]. In these experiments, four adult Long-Evans rats (Janni, Harpy, Imp and Ettin) performed back-and-forth runs on a one-dimensional linear track. A total of nine recording sessions were analyzed: four sessions from Janni, two from Harpy, one from Imp, and two from Ettin. Seven of these sessions followed the standard three-epoch structure: pre-run sleep in an isolated box, active running on the track, and post-run sleep in the same box. In the remaining two sessions, after an initial pre-run sleep epoch, the animals underwent three alternating run and post-run sleep epochs. Each run epoch lasted for approximately 40 min, and each pre- or post-run sleep epoch lasted for 40-60 min. The linear track was 1.8 m in length and 6.35 cm in width, expanding to 12.7 cm at each end, and was enclosed by 3.8 cm high walls. At both ends of the track, reward wells were filled out of view with 200–250 μL of chocolate-flavored liquid to encourage traversal. During each session, dorsal CA1 place cells were recorded as the animals navigated the track. Animal position and head direction were simultaneously monitored at 60 Hz using an overhead camera. More details about the experiments can be found in [30].

4.2 Place field estimation

Place fields were estimated using the neural responses while the animal was moving at speeds larger than 5 cm/s. The linear track was discretized into 2 cm spatial bins. For each neuron i, the neural activity vector was constructed by counting the number of spikes that occurred within each spatial bin. This vector was then divided by a spatial occupancy, computed by accumulating the time the animal spent in each bin. The resulting firing rate was smoothed using a Gaussian kernel with a standard deviation of 5 cm. The smoothed firing rate vector, denoted $f_i(z)$, was taken as the place field of neuron i. A neuron was classified as a place cell if the peak firing rate of its place field exceeded 1 spike/s. Only neurons identified as place cells were included in subsequent analysis.

4.3 Sharp-wave ripple detection

Sharp-wave ripples (SWRs) were detected following the procedure of Karlsson and Frank (2009) [9]. Raw local field potential (LFP) data were first band-pass filtered between 150 and 250 Hz to isolate the ripple-band oscillations. The instantaneous amplitude of the band signal was then obtained by taking the absolute value of its Hilbert transform and smoothing with a Gaussian kernel ($\sigma = 4$ ms). Candidate ripple events were identified as periods during which the smoothed amplitude exceeded the session-wide mean by three standard deviations. Each event's start and end times were extended backward and forward, respectively, to the time when the amplitude first crossed above and then fell below the mean. Overlapping detections across tetrodes or channels were merged into single events.

Ripple event boundaries were further refined by excluding edge time bins with no spiking activity. Only events lasting no less than 50 ms were retained.

4.4 State-space modeling

4.4.1 Transition dynamics

Temporal transitions of replayed positions were modeled with several categories of dynamics, including drift-diffusion, stationary and uniform processes.

Drift-diffusion. We start from the continuous-time differential equation for drift-diffusion process

$$\mathrm{d}z = \lambda\,\mathrm{d}t + \sigma\,\mathrm{d}W$$

where λ and σ are the drift and diffusion parameters, respectively, and W is a Wiener process with $\mathbb{E}[\mathrm{d}W] = 0$ and $\mathrm{Var}[\mathrm{d}W] = \mathrm{d}t$. Discretizing this stochastic differential equation with time-bin width $\mathrm{d}t$ gives a Gaussian transition kernel,

$$p(z_t \mid z_{t-1}) = \mathcal{N}(z_t;\, z_{t-1} + \lambda \mathrm{d}t,\, \sigma^2 \mathrm{d}t)$$

where $\mathcal{N}(z;\, \mu,\, \sigma^2)$ denotes the normal density of z with mean μ and variance σ^2.

Stationary. Similarly, we write the continuous-time differential equation for stationary dynamics

$$\mathrm{d}z = 0\,\mathrm{d}t$$

meaning the latent position is held fixed in time. Discretizing with bin width $\mathrm{d}t$ yields the trivial transition kernel:

$$p(z_t \mid z_{t-1}) = \delta(z_t - z_{t-1})$$

where $\delta(\cdot)$ is the Dirac delta.

Uniform. Here we assume complete absence of temporal structure. Each latent state is resampled independently at every time step. There is no meaningful continuous-time limit, so in discrete time

$$p(z_t \mid z_{t-1}) = \frac{1}{M}$$

Thus, every next state is equally likely, irrespective of the previous state.

Under certain parameter limits, the drift-diffusion dynamics reduces to other two candidate dynamics. In the limit $\lambda \to 0$ and $\sigma \to 0$, the Gaussian transition collapses to a point-mass at the previous state, reproducing the stationary model. As $\sigma \to \infty$, the Gaussian becomes flat over the track, and the transition matrix approaches the uniform model. In practice, during SWR events dominated by a single dynamical regime, these extreme parameter settings lead to near-equal marginal posterior probabilities for the drift-diffusion versus stationary or uniform categories.

4.4.2 Switching HMMs

An individual SWR event may switch between multiple regimes rather than following a single type of dynamics [24]. We developed a latent-position HMM model based on a two-level Markov-switching framework. The full joint distribution in this model can be described as

$$p(x_{1:T},\, z_{1:T},\, c_{1:T}) = \underbrace{\left[\prod_{t=1}^{T} p(x_t \mid z_t)\right]}_{\text{emissions}} \underbrace{\left[p(z_1)\prod_{t=2}^{T} p(z_t \mid z_{t-1},\, c_t)\right]}_{\text{position dynamics}} \underbrace{\left[p(c_1)\prod_{t=2}^{T} p(c_t \mid c_{t-1})\right]}_{\text{regime switching}}$$

We model the emission probability $p(x_{1:T} \mid z_{1:T})$ as conditionally independent Poisson observations given the latent location. Writing $x_{t,i}$ for the spike count of neuron i at step t, and $f_i(z)$ for its place field at location z, the emission likelihood factorizes as

$$p(x_{1:T} \mid z_{1:T}) = \prod_{t=1}^{T} p(x_t \mid z_t) = \prod_{t=1}^{T}\prod_{i=1}^{N} \text{Poisson}(x_{t,i} \mid \tau\,\gamma\, f_i(z_t))$$

where

$$\text{Poisson}(x \mid \lambda) = \frac{\lambda^x}{x!} e^{-\lambda}$$

Here $\tau = 10\,\mathrm{ms}$ is the bin width, and γ is a global gain factor that accounts for the increased firing rates during SWR events relative to locomotion.

To allow individual SWR events to exhibit mixtures of our three candidate dynamics, we augment the latent-position HMM with a discrete regime chain $c_t \in \{1, 2, 3\}$, where states 1, 2 and 3 index drift-diffusion, stationary and uniform dynamics, respectively. We write

$$p(c_1 = i) = \pi_i^{(c)}$$

where $\pi^{(c)} = \left(\pi_1^{(c)}, \pi_2^{(c)}, \pi_3^{(c)}\right)$ is the initial distribution over regimes, and

$$p(c_t = j \mid c_{t-1} = i) = \Pi_{i,j}$$

where Π is the regime-transition matrix. We further define the transition matrix as

$$\Pi = \begin{pmatrix} q & \frac{1-q}{2} & \frac{1-q}{2} \\ \frac{1-q}{2} & q & \frac{1-q}{2} \\ \frac{1-q}{2} & \frac{1-q}{2} & q \end{pmatrix}$$

This formulation assumes equal self-transition probability q across all regimes and uniform switching probability between regimes. The parameter q is learned from the data on an event-by-event basis.

Given the current dynamical regime c_t, the latent position z_t then evolves according to the corresponding transition probability kernel:

$$p(z_t \mid z_{t-1},\, c_t) = A^{(c_t)}_{z_{t-1},\, z_t}$$

with $A^{(1)}$, $A^{(2)}$, and $A^{(3)}$ defined in the above section for drift-diffusion, stationary and uniform dynamics. The joint probability over $(c_{1:T},\, z_{1:T})$ thus factorizes as

$$p(c_{1:T},\, z_{1:T}) = \left[p(z_1)\prod_{t=2}^{T} p(z_t \mid z_{t-1},\, c_t)\right]\left[p(c_1)\prod_{t=2}^{T} p(c_t \mid c_{t-1})\right]$$

Optimization of the model. For each SWR, we fit the three free parameters

$$\theta = \{\lambda,\, \sigma,\, q\}$$

by maximizing the marginal likelihood of the observed spikes under the switching HMM,

$$\hat{\theta} = \arg\max_{\lambda,\,\sigma,\,q}\; p(x_{1:T} \mid \theta) = \arg\max_{\lambda,\,\sigma,\,q} \sum_{c_{1:T}} \sum_{z_{1:T}} p(x_{1:T},\, z_{1:T},\, c_{1:T} \mid \theta)$$

We compute $p(x_{1:T} \mid \theta)$ via a forward pass over the two-level Markov chain $(c_t,\, z_t)$. We define the forward variable

$$\alpha_t(c_t, z_t) \;\equiv\; p(x_{1:t},\, c_t,\, z_t \mid \theta)$$

which satisfies the following initialization at $t = 1$:

$$\alpha_1(c_1, z_1) = p(c_1)\, p(z_1)\, p(x_1 \mid z_1)$$

and recursion $(t = 2, \ldots, T)$:

$$\alpha_t(c_t, z_t) = p(x_t \mid z_t) \sum_{c_{t-1}=1}^{C} \sum_{z_{t-1}=1}^{M} \alpha_{t-1}(c_{t-1}, z_{t-1})\, p(c_t \mid c_{t-1})\, p(z_t \mid z_{t-1},\, c_t)$$

Finally, the marginal likelihood is the total mass of the forward variable at $t = T$:

$$p(x_{1:T} \mid \theta) = \sum_{c_T=1}^{C} \sum_{z_T=1}^{M} \alpha_T(c_T, z_T)$$

In practice the computation happens in the log scale, thus yielding the overall log-likelihood. The negative log-likelihood serves as the loss function for optimization. Optimization is implemented with `scipy.optimize.minimize` using the L-BFGS-B algorithm, enforcing the box constraints

$$-15 < \lambda < 15 \text{ m/s}, \quad 0.1 < \sigma < 1 \text{ m}/\sqrt{\text{s}}, \quad 0.50 < q < 0.99$$

We use default convergence criteria provided by SciPy for termination: the relative objective decrease $\leq 2.2 \times 10^{-9}$, or the projected-gradient ℓ_∞-norm $\leq 10^{-5}$, or upon reaching the iteration limits 15000.

Inference. We compute the joint posterior over regimes and positions at each time t,

$$p(c_t, z_t \mid x_{1:T})$$

via a standard forward–backward pass. The forward recursion (as discussed in optimization) yields

$$\alpha_t(c_t, z_t) \;=\; p(x_{1:t},\, c_t,\, z_t \mid \theta)$$

and the backward recursion defines

$$\beta_t(c_t, z_t) \;=\; p(x_{t+1:T} \mid c_t,\, z_t,\, \theta)$$

with

$$\beta_T(c_T, z_T) = 1$$

and for $t = T-1, \ldots, 1$:

$$\beta_t(c_t, z_t) = \sum_{c_{t+1}=1}^{C} \sum_{z_{t+1}=1}^{M} p(c_{t+1} \mid c_t)\, p(z_{t+1} \mid z_t,\, c_{t+1})\, p(x_{t+1} \mid z_{t+1})\, \beta_{t+1}(c_{t+1}, z_{t+1})$$

The smoothed joint distribution then can be computed as

$$p(c_t, z_t \mid x_{1:T}) \;\propto\; \alpha_t(c_t, z_t)\, \beta_t(c_t, z_t)$$

Marginal posteriors can be obtained by marginalization:

$$p(c_t \mid x_{1:T}) = \sum_{z_t=1}^{M} p(c_t, z_t \mid x_{1:T}), \qquad p(z_t \mid x_{1:T}) = \sum_{c_t=1}^{C} p(c_t, z_t \mid x_{1:T})$$

from which we extract both a point estimate of the replayed trajectory and a segmentation into dynamical regimes.

Simulation. Our switching HMM is a generative model. We can simulate both latent regimes, latent positions, and spike-count observations given ground-truth parameters

$$\theta_{\mathrm{gt}} = \{\lambda_{\mathrm{gt}},\, \sigma_{\mathrm{gt}},\, p_{\mathrm{gt}}\}$$

and place fields $\{f_i(z)\}$. For a sequence of length T, we proceed as follows:

Initialize:

$$c_1^{\mathrm{sim}} \sim p(c_1), \quad z_1^{\mathrm{sim}} \sim p(z_1).$$

Latent-state evolution $(t = 2, \ldots, T)$:

$$c_t^{\mathrm{sim}} \sim p\big(c_t \mid c_{t-1}^{\mathrm{sim}}; p_{\mathrm{gt}}\big), \quad z_t^{\mathrm{sim}} \sim p\big(z_t \mid z_{t-1}^{\mathrm{sim}},\, c_t^{\mathrm{sim}}; \lambda_{\mathrm{gt}}, \sigma_{\mathrm{gt}}\big).$$

Observation sampling:

$$x_t^{\mathrm{sim}} \sim \prod_{i=1}^{N} \mathrm{Poisson}\big(x_{t,i} \mid \tau\, \gamma\, f_i(z_t^{\mathrm{sim}})\big), \quad t = 1, \ldots, T.$$

This yields synthetic spike-count data $x_{1:T}^{\mathrm{sim}}$ together with their true latent trajectories $\{c_t^{\mathrm{sim}},\, z_t^{\mathrm{sim}}\}$. We then can apply the same forward–backward inference and gradient-based parameter optimization to verify the ability to recover θ_{gt}.

4.5 Classic Bayesian decoding and replay metrics

We next describe the procedures for applying standard Bayesian decoding to study replays. SWRs were first binned into non-overlapping 10 ms windows, giving spike-count vector $\mathbf{x}_t = (x_1(t), \ldots, x_N(t))$. At time t, we computed the posterior over M spatial bins as

$$p(z \mid \mathbf{x}_t) = \frac{p(z)\, p(\mathbf{x}_t \mid z)}{\sum_{z'=1}^{M} p(z')\, p(\mathbf{x}_t \mid z')} \quad \text{with} \quad p(z) = \frac{1}{M}$$

The likelihood was modeled as independent Poisson spiking across cells:

$$p(\mathbf{x}_t \mid z) = \prod_{i=1}^{N} \frac{(\tau\, f_i(z))^{x_i(t)}\, e^{-\tau\, f_i(z)}}{x_i(t)!}$$

where $\tau = 0.01$ s is the time-bin width and $f_i(z)$ is the place fields of neuron i at location z. The denominator ensures $\sum_z p(z \mid \mathbf{x}_t) = 1$.

We then extracted a decoded trajectory by taking the maximum-a-posteriori estimate at each t,

$$\hat{z}_t = \arg\max_z p(z \mid \mathbf{x}_t)$$

and concatenating $\{\hat{z}_t\}$ across the SWR.

Absolute correlation. We computed the absolute weighted Pearson correlation between time bin index t and position bin z_k under the posterior $p(z_k \mid \mathbf{x}_t)$. We first compute the weighted means

$$\bar{t} = \frac{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}\, t}{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}}, \quad \bar{z} = \frac{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}\, z_k}{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}}$$

Next, we form covariance

$$\text{cov}(t, z) = \frac{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}\, (t - \bar{t})\, (z_k - \bar{z})}{\sum_{t,k} w_{tk}}$$

and weighted variances

$$\text{var}(t) = \frac{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}\, (t - \bar{t})^2}{\sum_{t,k} w_{tk}}, \quad \text{var}(z) = \frac{\sum_{t=1}^{T} \sum_{k=1}^{M} w_{tk}\, (z_k - \bar{z})^2}{\sum_{t,k} w_{tk}}$$

The Pearson correlation is then

$$\rho = \frac{\text{cov}(t, z)}{\sqrt{\text{var}(t)\, \text{var}(z)}}$$

and we report the absolute correlation

$$|\rho| \in [0, 1]$$

where values near 1 indicate a sequential trajectory and values near 0 indicate little or no monotonic relationship between decoded position and time.

Max jump distance. We computed the maximum spatial jump between two time bins within each SWR. From the MAP trajectory $\{\hat{z}_t\}_{t=1}^{T}$, we first formed the sequence of jump magnitudes

$\Delta_t = |\hat{z}_t - \hat{z}_{t-1}|, \quad t = 2, \ldots, T$. The max jump distance is then defined as $D_{\text{max}} = \max_{2 \le t \le T} \Delta_t$. High values of D_{max} indicate large noise in the decoded trajectory, whereas low values reflect more spatially continuous sequences.

Replay speed. We fit a simple linear regression to the MAP trajectory $\{\hat{z}_t\}_{t=1}^{T}$ as a function of time bin index t: $\hat{z}_t = \beta_0 + \beta_1 t + \varepsilon_t$. We then define replay speed as the absolute value of the fitted slope $|\beta_1|$, which has units of spatial bins per time bin. To express this in physical units (e.g., cm/s), we multiply it by the spatial bin size and divide it by the time bin width.

4.6 Synthetic data experiments

Parameter identifiability. We ran two simulation experiments to assess the identifiability of drift-diffusion parameters and the accuracy of regime sequence recovery in our switching HMM. In both experiments, we constructed a synthetic population of 300 place cells whose place fields are defined as Gaussian bumps $f_i(z) = 20\,\text{Hz}\,\exp\left[-\frac{(z-\mu_i)^2}{2\,(0.1\,\text{m})^2}\right]$. The preferred positions μ_i uniformly tile a 4 m one-dimensional linear track.

Drift-diffusion regime. To test recovery of drift-diffusion parameter, we defined a partial 2-D grid of ground-truth drift and diffusion values. We swept $\lambda_{\text{gt}} \in \{0, 2, 4, 6, 8\}\,\text{m/s} \times \sigma_{\text{gt}} \in \{0.2, 0.3\}\,\text{m}/\sqrt{\text{s}}$ and $\sigma_{\text{gt}} \in \{0.1, 0.2, 0.3, 0.4, 0.5\}\,\text{m}/\sqrt{\text{s}} \times \lambda_{\text{gt}} \in \{0, 2\}\,\text{m/s}$. For each $(\lambda_{\text{gt}}, \sigma_{\text{gt}})$ pair on the grid, we simulated 100 independent SWRs as follows. First, generate a 300 ms latent trajectory $\{z_t\}$ by evolving the drift-diffusion process with specified λ_{gt} and σ_{gt}. Next, use the place-cell population to sample Poisson spike counts $\{x_t\}$ in 10 ms bins at the trajectory positions. We then ran our inference pipeline on each SWR, recovered estimates $\hat{\lambda}$ and $\hat{\sigma}$ and compared them to the ground truth.

Switching of dynamics (dynamical regimes). To test recovery of both regime structure and drift-diffusion parameters, we expanded our ground truth to: (a) every ordered pair of candidate dynamics (drift-diffusion, stationary, uniform), and (b) for any segment governed by drift-diffusion, each $(\lambda_{\text{gt}}, \sigma_{\text{gt}})$ combination from the same grid as in drift-diffusion regime experiment. For each of these regime-pair $\times$ parameter conditions, we simulated 100 independent SWRs as follows. First, specify the two-segment regime sequence (e.g., drift-diffusion $\rightarrow$ uniform). Next, generated a latent path for each 300 ms segment. Use the drift-diffusion dynamics with the given λ_{gt} and σ_{gt} when applicable, or the stationary/uniform dynamics otherwise. Concatenate the two segments into a continuous 600 ms trajectory. Along this trajectory, sample Poisson spike counts from place-cell population in 10 ms bins. By applying the inference pipeline to each simulated SWR, we recovered the latent regime sequence $\{\hat{c}_t\}$ and, when drift–diffusion segment presents, the parameter estimates $\hat{\lambda}$ and $\hat{\sigma}$. We then compared $\{\hat{c}_t\}$, $\hat{\lambda}$ and $\hat{\sigma}$ against the ground-truth.

Comparative analysis between drift-diffusion parameters and prior replay metrics. We simulated SWR events under the pure drift-diffusion dynamics using the same settings as in identifiability experiments. We ran two sets of simulations: (a) Fix the $\sigma = 0.4\,\text{m}/\sqrt{\text{s}}$ and simulate 2000 SWRs for each $\lambda_{\text{gt}} \in \{0, 2, 4, 6, 8, 10\}$ m/s. (b) Fixed the $\lambda = 2$ m/s and simulate 2000 SWRs for each

$\sigma_{\mathrm{gt}} \in \{0.2, 1, 2, 4, 6, 8\}$ m/$\sqrt{\mathrm{s}}$. We applied classic Bayesian decoding to calculate the location posterior and MAP trajectory. From these decoded trajectories, we then computed absolute correlation, maximum jump distance and replay speed as outlined above, and compared them to the ground-truth λ_{gt} and σ_{gt}.

4.7 Additional details of the figures

Speed calculation. We computed three complementary measures of speed in Fig. 5. *Per-lap running speed*. Individual laps on the linear track were identified by crossings of the track midpoint (90 cm). For an up-crossing of the midpoint, the start of the lap was defined as the most recent up-crossing of 65 cm, and the end as the first new up-crossing of 115 cm. For down-crossings, the similar down-crossings of 115 cm and 65 cm were used. After the position-time trace of each lap was extracted, we fit a simple linear regression to the data within that lap and took the per-lap running speed as the absolute value of slope. *Instantaneous behavioral speed*. Within each running epoch, we estimated instantaneous velocity by two-step centered differencing of the position time series, $v(t) = (\text{pos}(t+1) - \text{pos}(t-1))/2\Delta t$, where $\Delta t = 10$ ms is the time bin. The instantaneous behavioral speed was taken as the absolute value $|v(t)|$. *Replay speed*. We used the absolute value of inferred drift-rate parameter $|\lambda|$ (in m/s) as replay speed.

Shuffle analysis. For Fig. 7, null replay distributions were generated using three types of shuffle procedures. *Time-bin shuffle*. The time bins of spike raster were randomly permuted within each SWR. *Neuron-identity shuffle*. The neuron indices of spike raster were randomly permuted within each SWR. *Place-field shuffle*. Place fields of each neuron were circularly shifted by a random offset. Each shuffle procedure was applied to every SWR and repeated 120 times to construct a null distribution. Shuffled SWRs were then processed and analyzed using the same pipeline used for real replay events.

4.8 Data and code availability

This study uses data from Pfeiffer and Foster (2015) [30]. Please contact Brad Pfeiffer and David Foster to access the analyzed datasets. The code used for data analysis, simulations, and figure generation is archived on a GitHub repository (`https://github.com/ZhongxuanWu/ReplaySwitchingHMM`). Any further information needed to reproduce the results can be obtained from the lead contact upon request.

Supplemental Information

A Additional figures and discussions

A.1 Parameters in drift-diffusion framework are more interpretable than prior metrics

Previous methods first infer the posterior probability over space using the place fields inferred during running [1], then compute three metrics for interpretation (summarized in [2]): (i) (absolute) correlation between time and decoded trajectory; (ii) replay speed inferred from the slope of decoded trajectory; and (iii) maximum jump distance (Fig. S1A). In contrast, our modeling framework uses a generative approach and describes a replay event using the drift and diffusion parameters, which unambiguously characterize the speed and the sequentiality (or quality) of the replay respectively. We next aim to determine the relationship between the key parameters in prior methods and those in ours.

We generate simulated neural population recording data with different combinations of drift and diffusion parameters, and then apply the traditional analysis method based on Bayesian decoding to extract the three metrics mentioned above, i.e., absolute correlation, speed of replay, maximum jump distance (Fig. S1A). Fig. S1B shows the results when fixing the diffusion parameter while allowing the drift parameter to vary, while Fig. S1C shows the results when fixing the drift parameter while allowing the diffusion parameter to vary. These results suggest that the absolute correlation is positively correlated with drift and negatively correlated with diffusion. For the speed inferred from the Bayesian approach, it is positively correlated with both the drift and diffusion parameter. Thus, the speed can be overestimated when the diffusion parameter is large. Furthermore, the maximum jump distance [2] mostly reflects the magnitude of the diffusion parameter, i.e., large diffusion leads to bigger maximum jump distance, yet it also increases systematically with the drift. This suggests that maximum jump distance reflects a combination of drift and diffusion, in particular when the diffusion parameter is small. Together, these results suggest that, while the three key metrics used in prior studies may capture some structures in the replay dynamics as conceived when they were developed, each of them reflects the combined effects of the speed and the quality of a candidate replay event. Thus, practically they may not allow for an unambiguous interpretation of the replay dynamics. In contrast, our approach clearly distinguishes the speed and quality of a replay event using separate drift and diffusion parameters, and it allows one to dissect the contribution of the two in a principled single-step inference framework.

A.2 Discussion of shuffle controls

Across shuffle controls, we observed a specific discrepancy for the time-bin shuffle (Fig. 7 and Fig. S7). First, the fraction of SWRs exceeding each drift-diffusion duration threshold for time-bin shuffled data was consistently lower than that for the other two controls, whereas the neuron-identity and place-field rotation shuffles yielded similar null fractions. Second, the null fractions from time-bin shuffle did not align

across pre- and post-run data. At every threshold a larger fraction of post-run shuffled events exceeded the criterion than pre-run shuffled events did. This is in contrast to the overlapping pre/post null fractions produced by the neuron-identity and place-field rotation shuffles.

Note that a comparable divergence between time-bin and other shuffles was also reported by Takigawa et al. [3]. The reason is possibly that only the time-bin shuffle preserves spike–position correlations [4, 5]. But this interpretation is provisional, and the mechanism requires further investigation.

A.3 Discussion of bidirectional place fields

On a linear track, place cells may exhibit direction-dependent place fields. In our analyses, we estimate place fields of each neuron by pooling spikes and occupancy from both running directions. For clarity, we refer to these pooled estimates as bidirectional place fields throughout and use them in the main analyses. We chose to combine the data from the two directions for the following reasons. First, the peaks of individual place fields are highly consistent across directions (Fig. S8 and Fig. S10). Second, using the bidirectional fields, the animal's position during run periods can be decoded accurately (Fig. S9 and Fig. S11). Therefore, if replay recapitulates place cell activity during running and forms coherent trajectories, the bidirectional place fields should be appropriate for decoding the content of the replay events. Finally, our choice matches the preprocessing in the original report on this dataset [6], which did not split place fields by running directions.

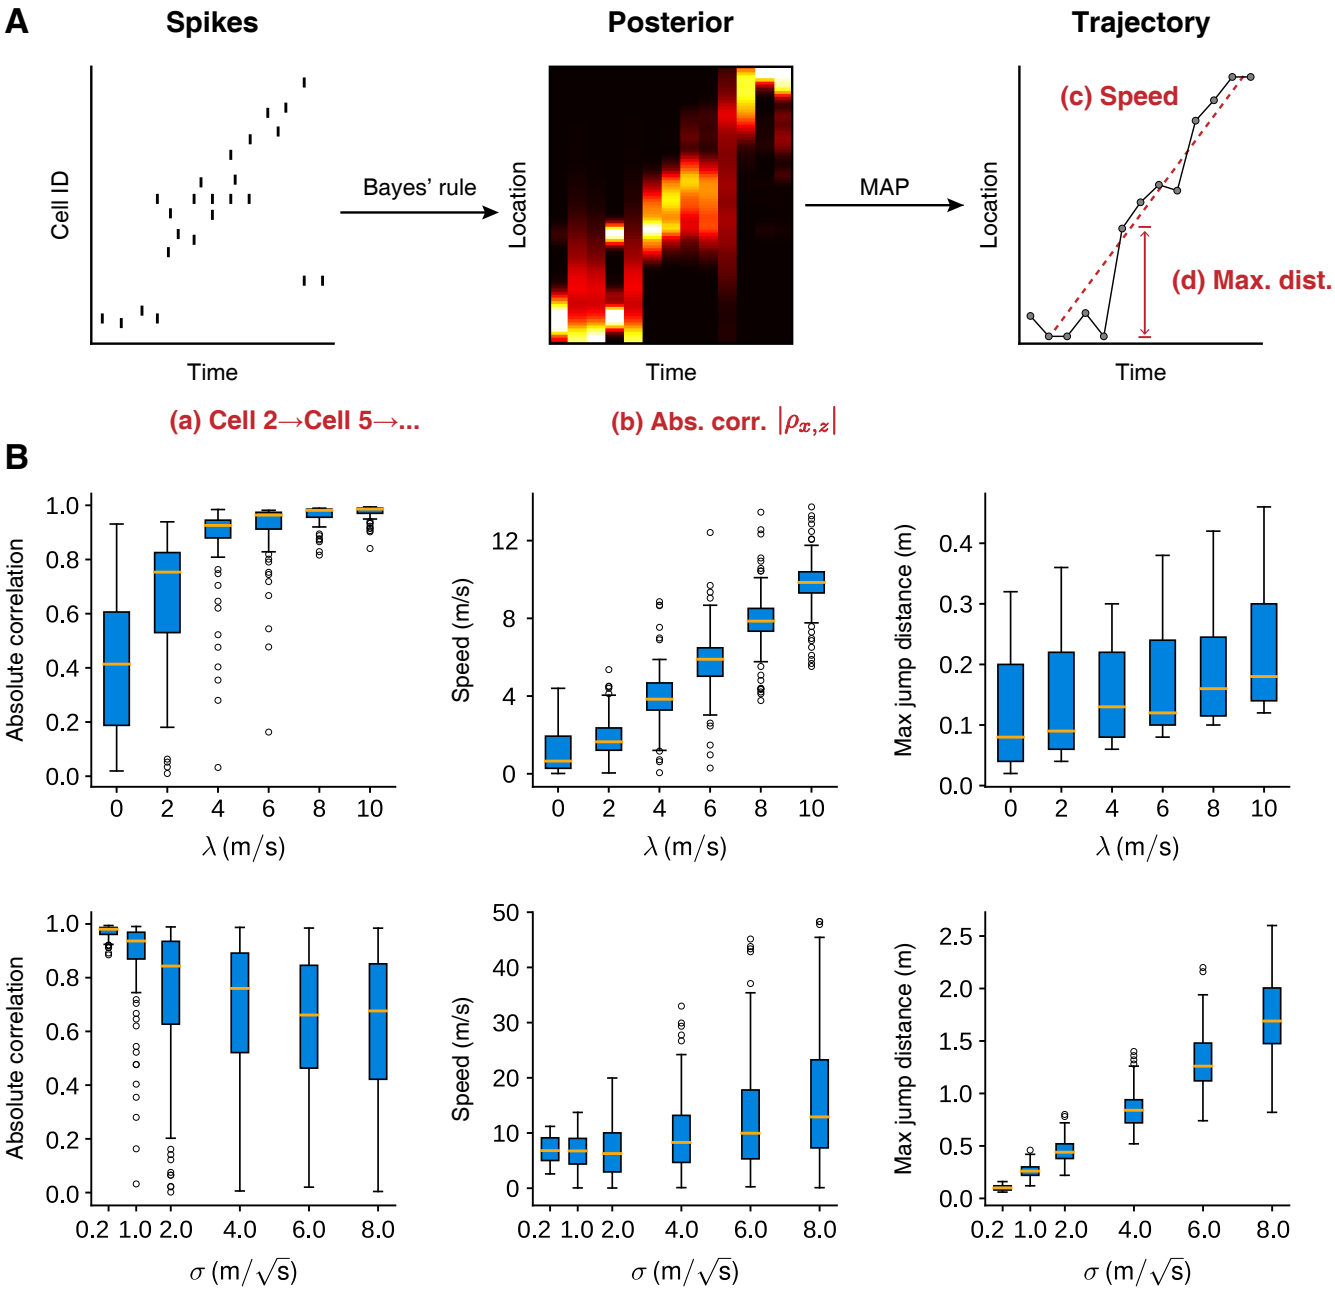

Figure S1: Relationship to previous replay analysis framework. A. Previous two-step replay analysis. Bayesian decoding (in black). Left: Spike-raster matrix for a single SWR. Middle: Bayesian location posterior $p(z_t|x_t) \propto p(x_t|z_t)p(z_t)$, computed using each neuron's place field and a uniform prior $p(z)$. Right: decoded trajectory obtained as the MAP estimate $\hat{z}_t = \arg\max_z p(z_t|x_t)$. Metric calculation (in red). Temporal firing sequence, compared against the animal's running firing template. Absolute correlation $|\rho_{t,z}|$: weighted Pearson correlation between time bin t and position z under the posterior, reflecting sequence linearity. Replay speed: absolute slope $|\beta_1|$ of a linear regression fit $\hat{z}_t = \beta_0 + \beta_1 t$. Max jump distance: $D_{\max} = \max_t |\hat{z}_t - \hat{z}_{t-1}|$, the largest spatial jump between consecutive bins. **B**. Boxplots illustrate how the three classic Bayesian replay metrics depend on the drift λ (top row) and diffusion parameter σ (bottom row). In each plot, one parameter is held constant while the other varies. For λ sweeps, $\sigma = 0.4$ m/$\sqrt{s}$ and for σ sweeps, $\lambda = 2$ m/s. Boxes show the median and interquartile range, and whiskers extend to $1.5\times$ IQR.

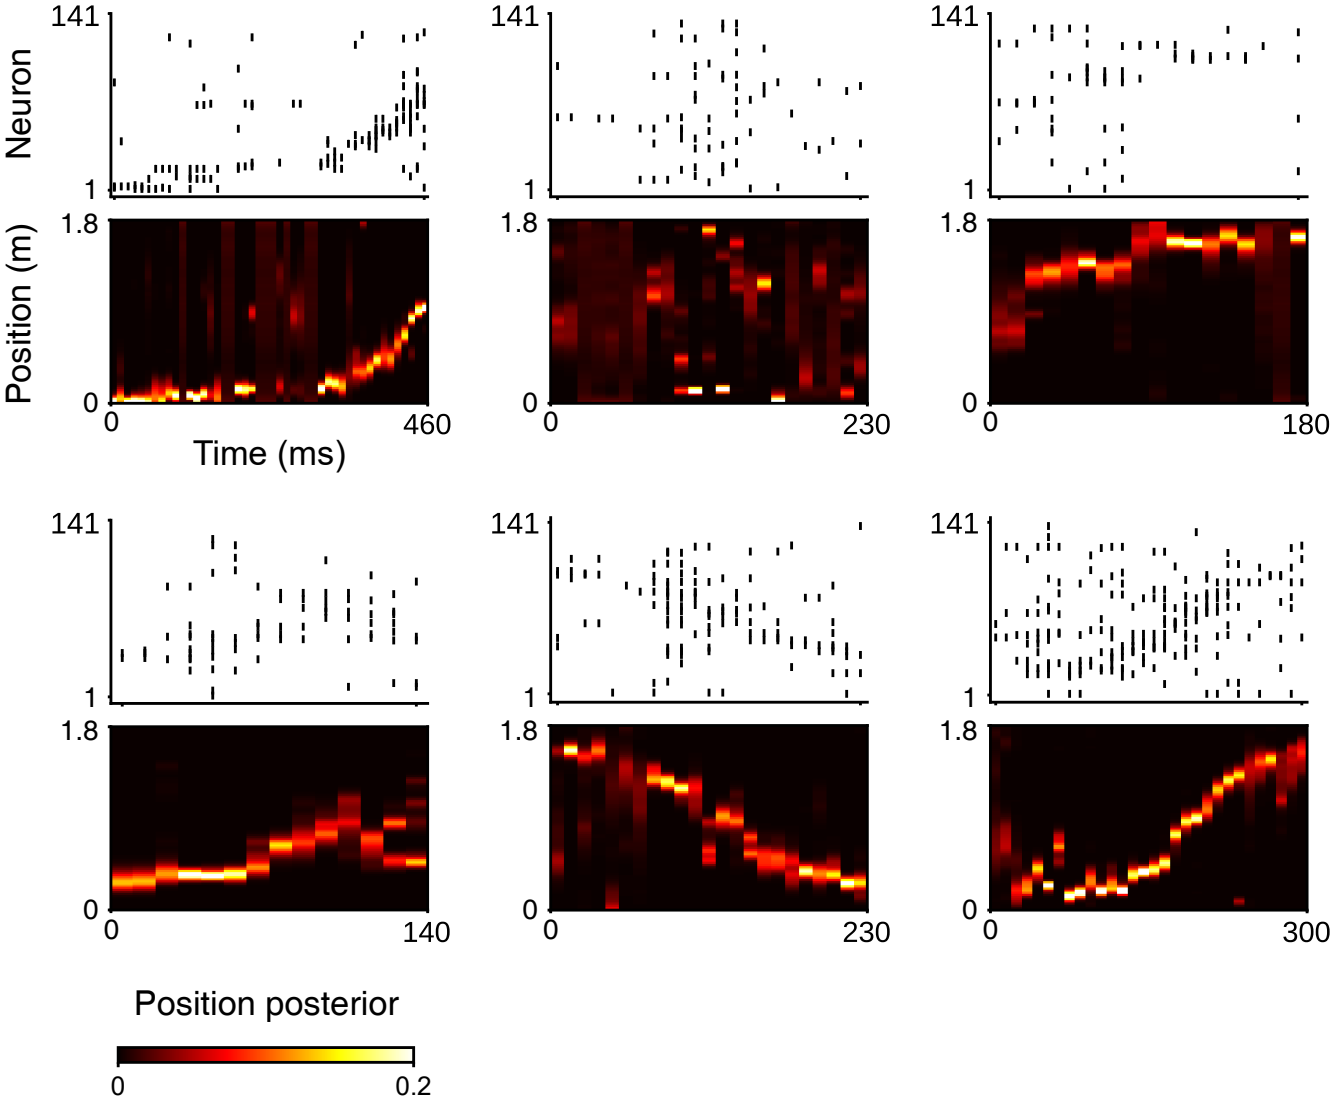

Figure S2: Spike rasters and position posteriors from classic Bayesian decoding for the example events. These events correspond to those shown in Fig. 5A. For each event, the spike raster (top, neurons sorted by place field peak location) and decoded position posterior (bottom) are shown. Posteriors were computed with a Poisson Bayesian decoder using 10 ms time bins and no temporal smoothing. Horizontal axis: time (ms). Vertical axis: linear-track position (m). Color bar: posterior probability.

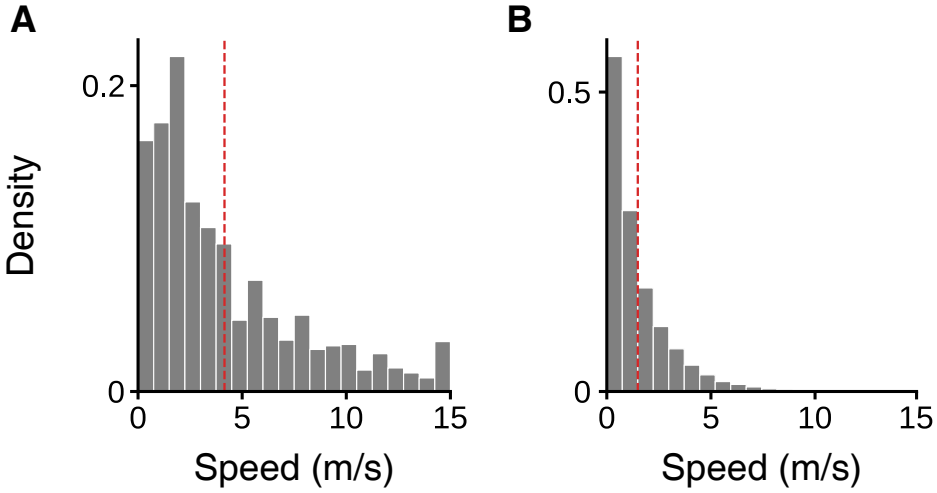

Figure S3: Replay-speed distribution observed in data is inconsistent with a pure Brownian motion null. A. Histogram of replay speed, estimated by the inferred drift rate λ for SWRs in which the marginal posterior of drift-diffusion regime exceeded 0.7 for at least 50 ms. Panel repeats the one shown in Fig. 5 for comparison. **B.** Histogram under a pure Brownian motion null. For each SWR in A, we fit a Brownian motion model by MLE to estimate the diffusion parameter and simulate 10 trajectories of matching duration. For each trajectory, we compute its speed as the absolute value of slope of a linear regression fit to position and time. The pooled distribution of speeds is shown. Dashed vertical lines show the mean of each distribution: 4.1 m/s (replay) and 1.5 m/s (null).

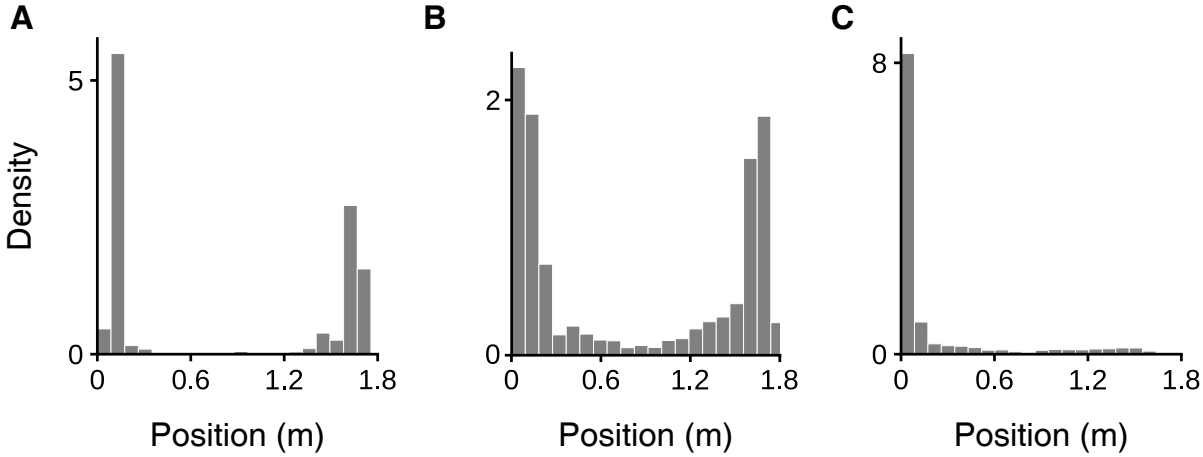

Figure S4: The distribution of latent position during stationary periods. Analyses combined periods that our model labeled as stationary within awake-rest SWR events (total duration $T = 88.99$ sec across 9 sessions). **A.** Histogram of the actual position of rats on the linear track during awake SWR periods. **B.** Histogram of the decoded latent position. **C.** Histogram of the absolute difference between actual and decoded latent positions.

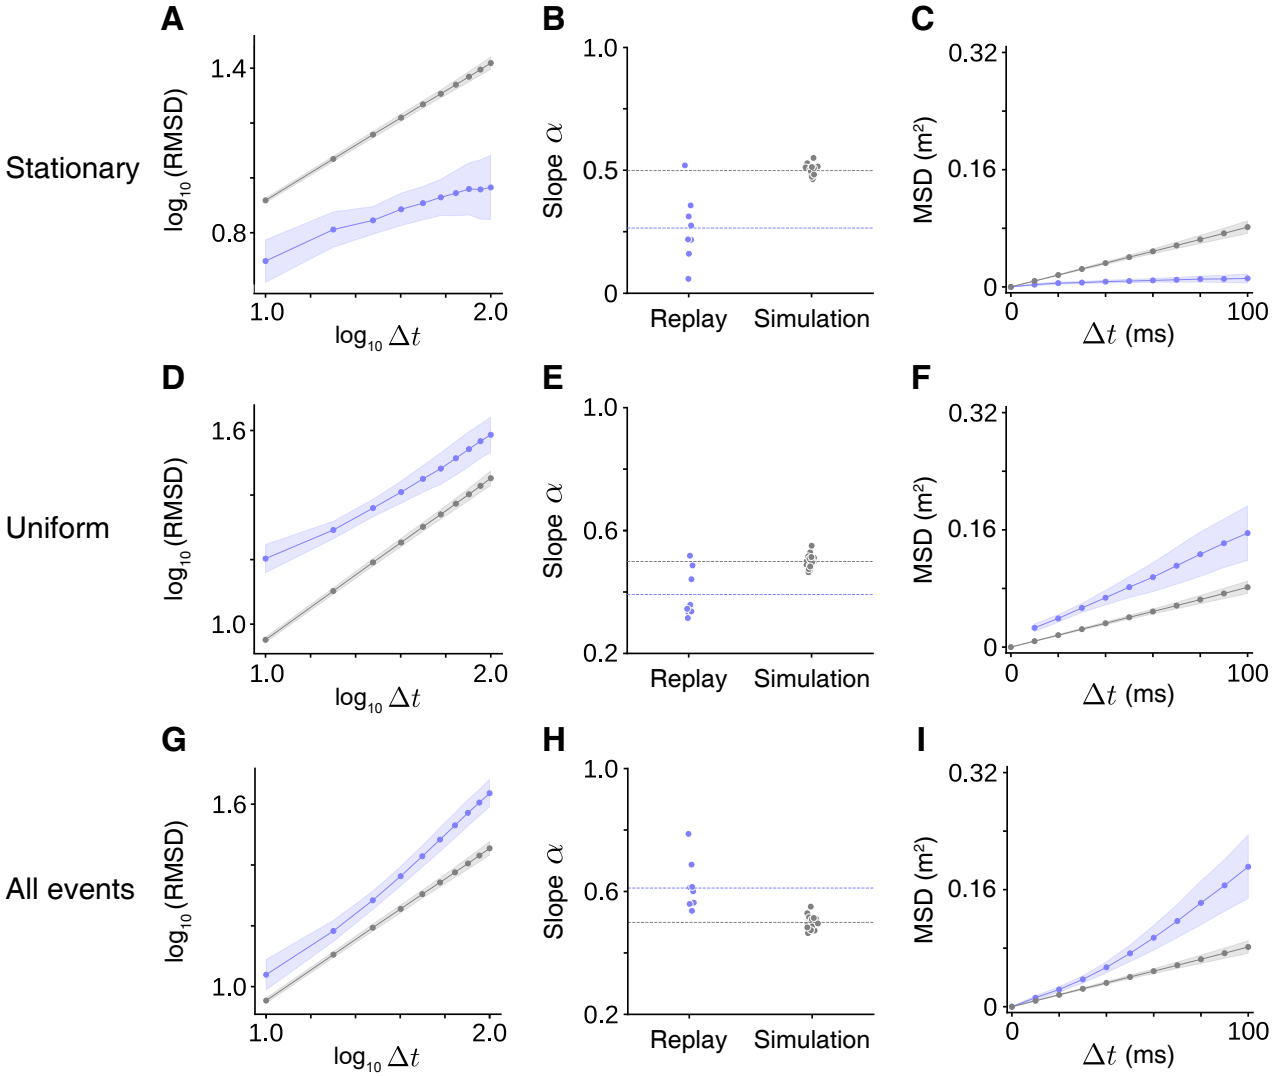

Figure S5: The scaling relationships of RMSD-Δt and MSD-Δt are distinct for stationary, uniform, and all SWR events. **A-C.** Plotting conventions follow Fig. 6B-D. Analyses repeated to SWR events with pure stationary dynamics. The mean across sessions in **B**: 0.26 for SWRs with stationary dynamics; 0.50 for simulated data from Brownian motion. **D-F.** Analyses repeated to SWR events with pure uniform dynamics. The mean across sessions in **E**: 0.39 for replay; 0.50 for simulation. For **F**, the real-data point at $\Delta t = 0$ is omitted. The reason is that for uniform dynamics, the MSD is strictly positive for any $\Delta t \geq 0$ but the value at $\Delta t = 0$ is degenerate and cannot be estimated reliably. **G-I**. Analyses repeated to all SWR events. The mean across sessions in **H**: 0.61 for replay; 0.50 for simulation. One session was excluded from the analyses in **A-F** because fewer than 50 SWR events met the inclusion criteria.

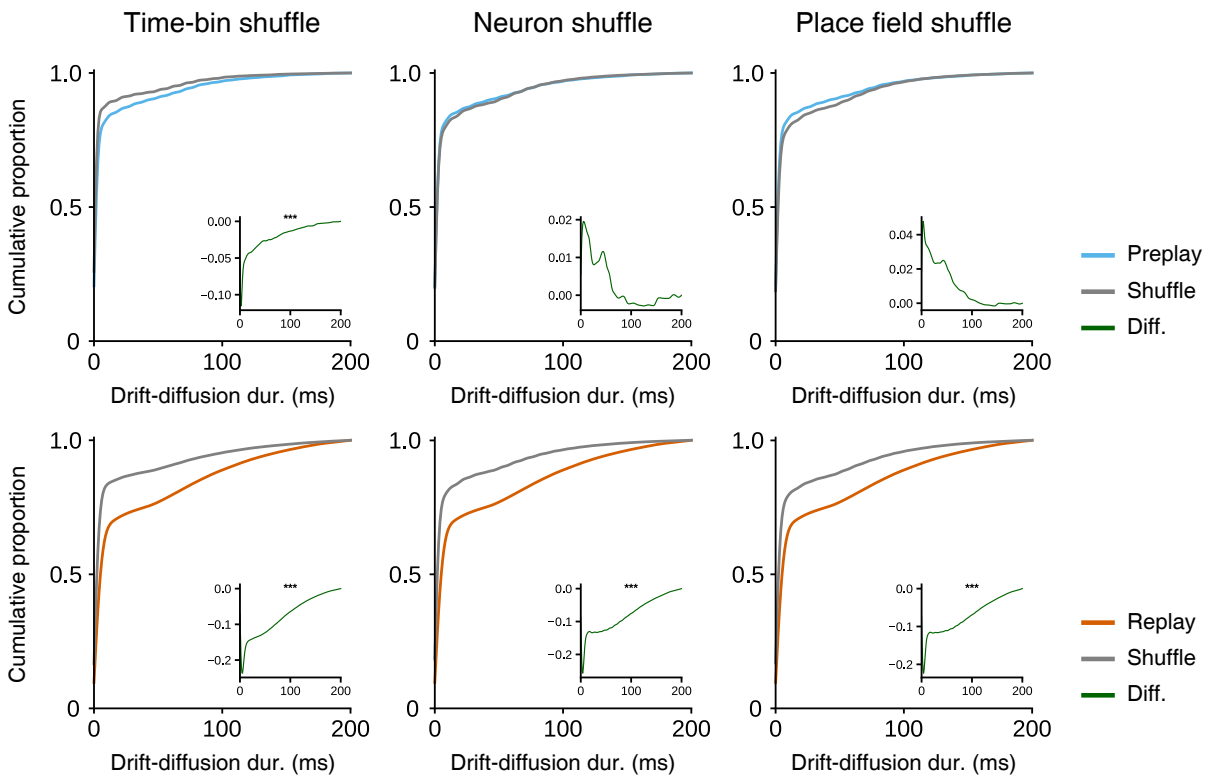

Figure S6: Comparison of drift-diffusion duration in preplay and replay SWRs to corresponding shuffle. Cumulative proportions of drift-diffusion durations for pre-run (top row, blue) and post-run (bottom row, orange) SWRs, compared to shuffled SWRs. Columns correspond to three shuffle controls—time-bin (left), neuron-identity (middle), and place-field rotation (right) Inset: the point-wise difference between the two cumulative proportions. Asterisks marking a significant divergence (one-sided Kolmogorov–Smirnov test, *** $p < 0.001$).

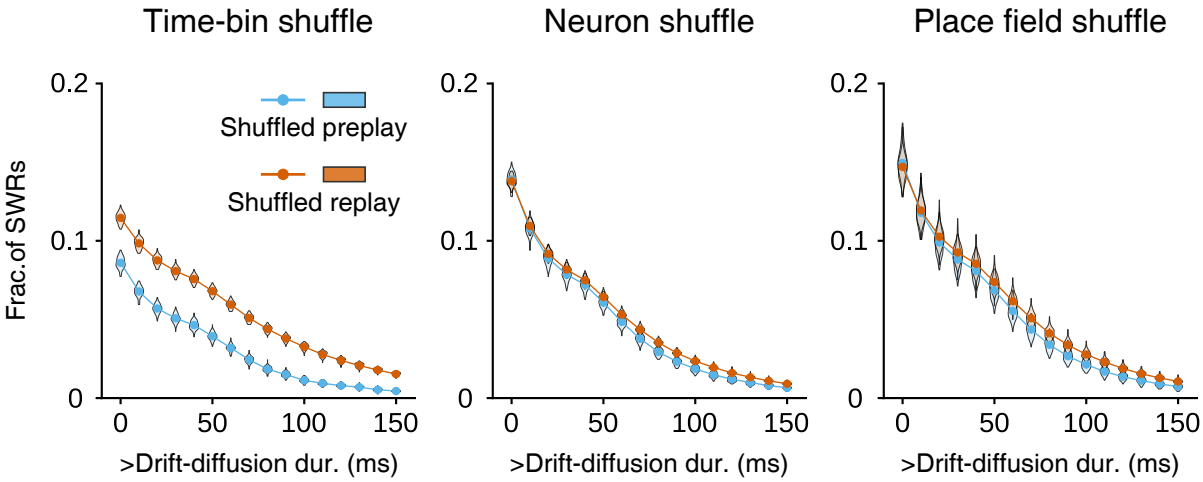

Figure S7: Comparison of shuffled preplay versus shuffled replay SWRs. The fraction of shuffled preplay (blue) and shuffled replay (orange) SWRs exceeding each drift-diffusion duration threshold (calculation as in Fig. 7C-H). Solid lines plot the mean fractions. Violins plot the distribution of the fractions across 120 shuffles. Columns show three shuffle methods: time-bin (left), neuron-identity (middle), and place-field rotation (right).

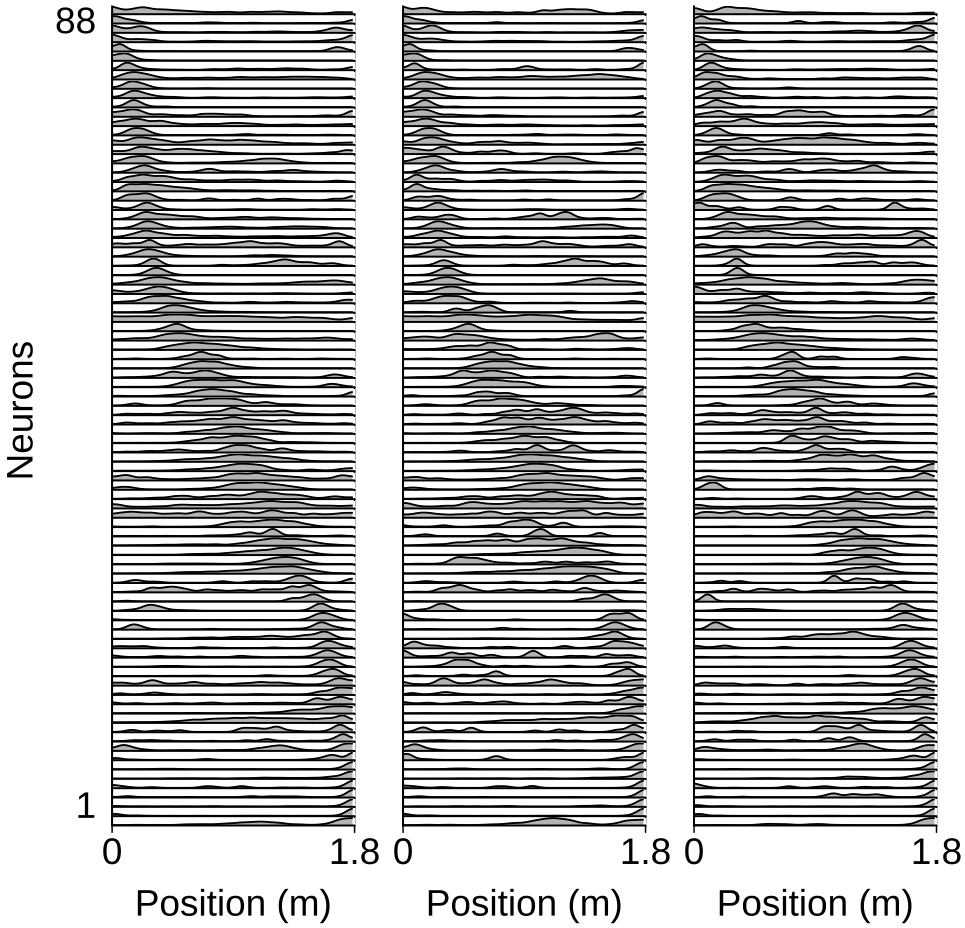

Figure S8: Place fields in an example session of rat Janni. Left: bidirectional place fields obtained by pooling spikes and occupancy from both running directions. Middle: place fields estimated from left-to-right traversals only. Right: place fields estimated from right-to-left traversals only. Each row is the place fields of one neuron. Place fields are normalized to the peak firing rate. Neurons are ordered by the peak position of the bidirectional place fields, and the same order is used in all panels.

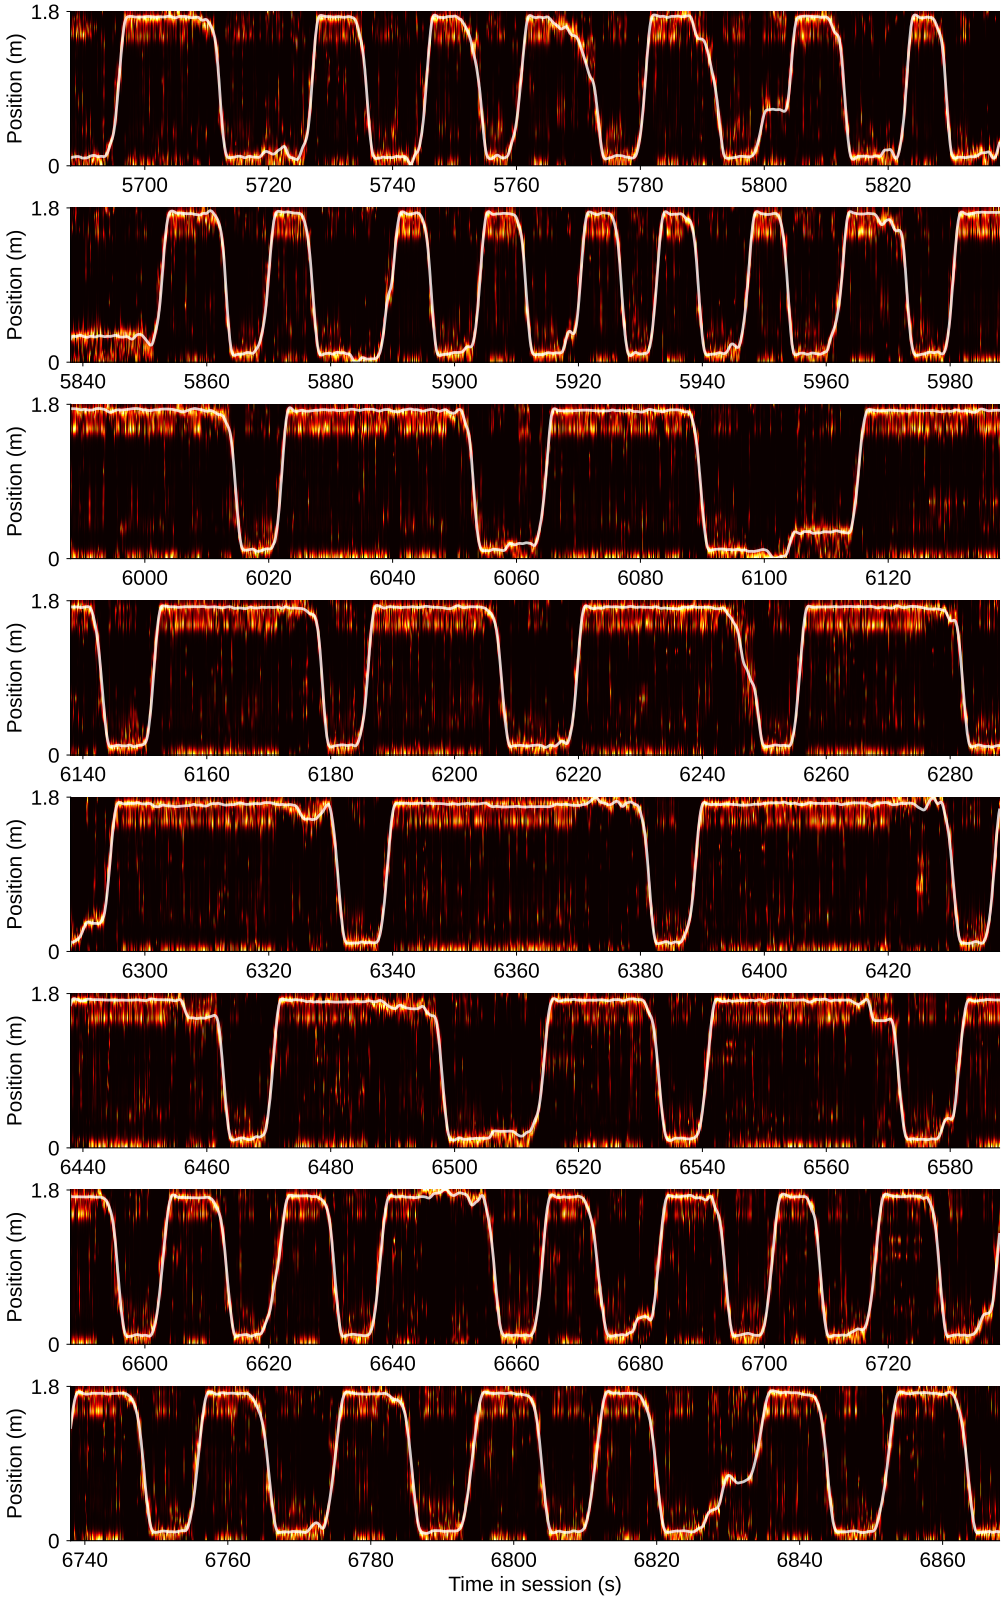

Figure S9: Decoding run position using bidirectional place fields. Heatmap shows the position posterior along the linear track (brighter = more probable). This is estimated by a memoryless Bayesian decoder [1] applied to place cell activity during the run epoch. The white trace shows the tracked position of the animal. Panels depict consecutive time segments from the same session. Same session as in Fig. S8.

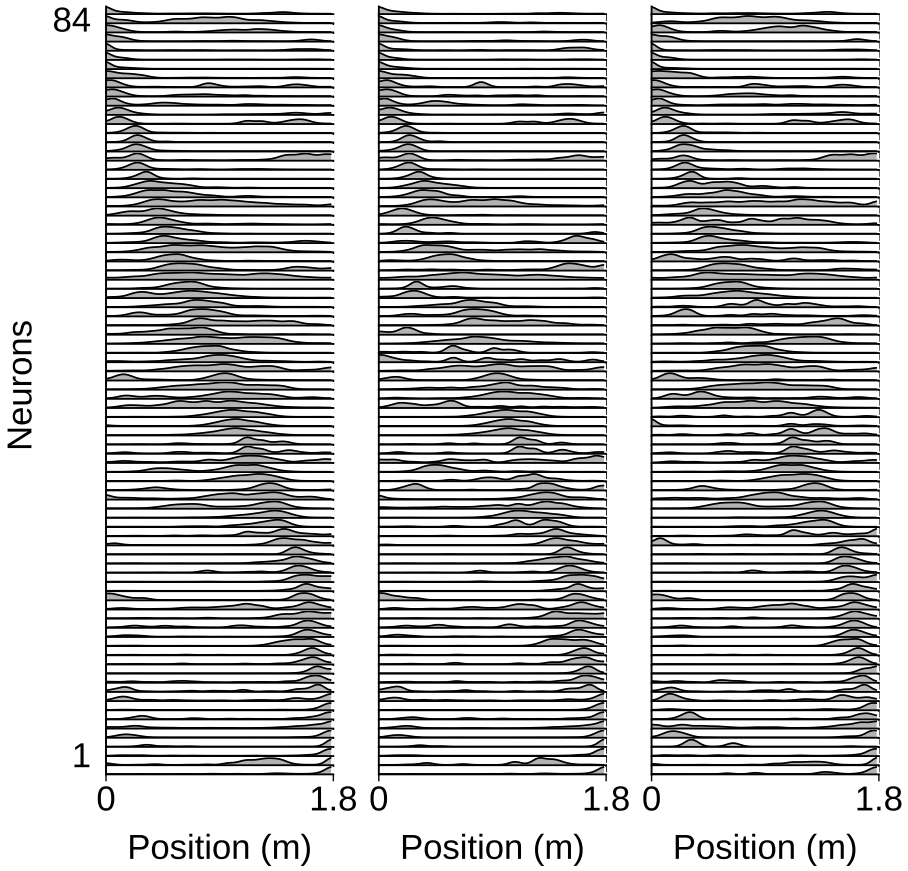

Figure S10: Place fields in an example session of rat Harpy. Plotting conventions follow Fig. S8.

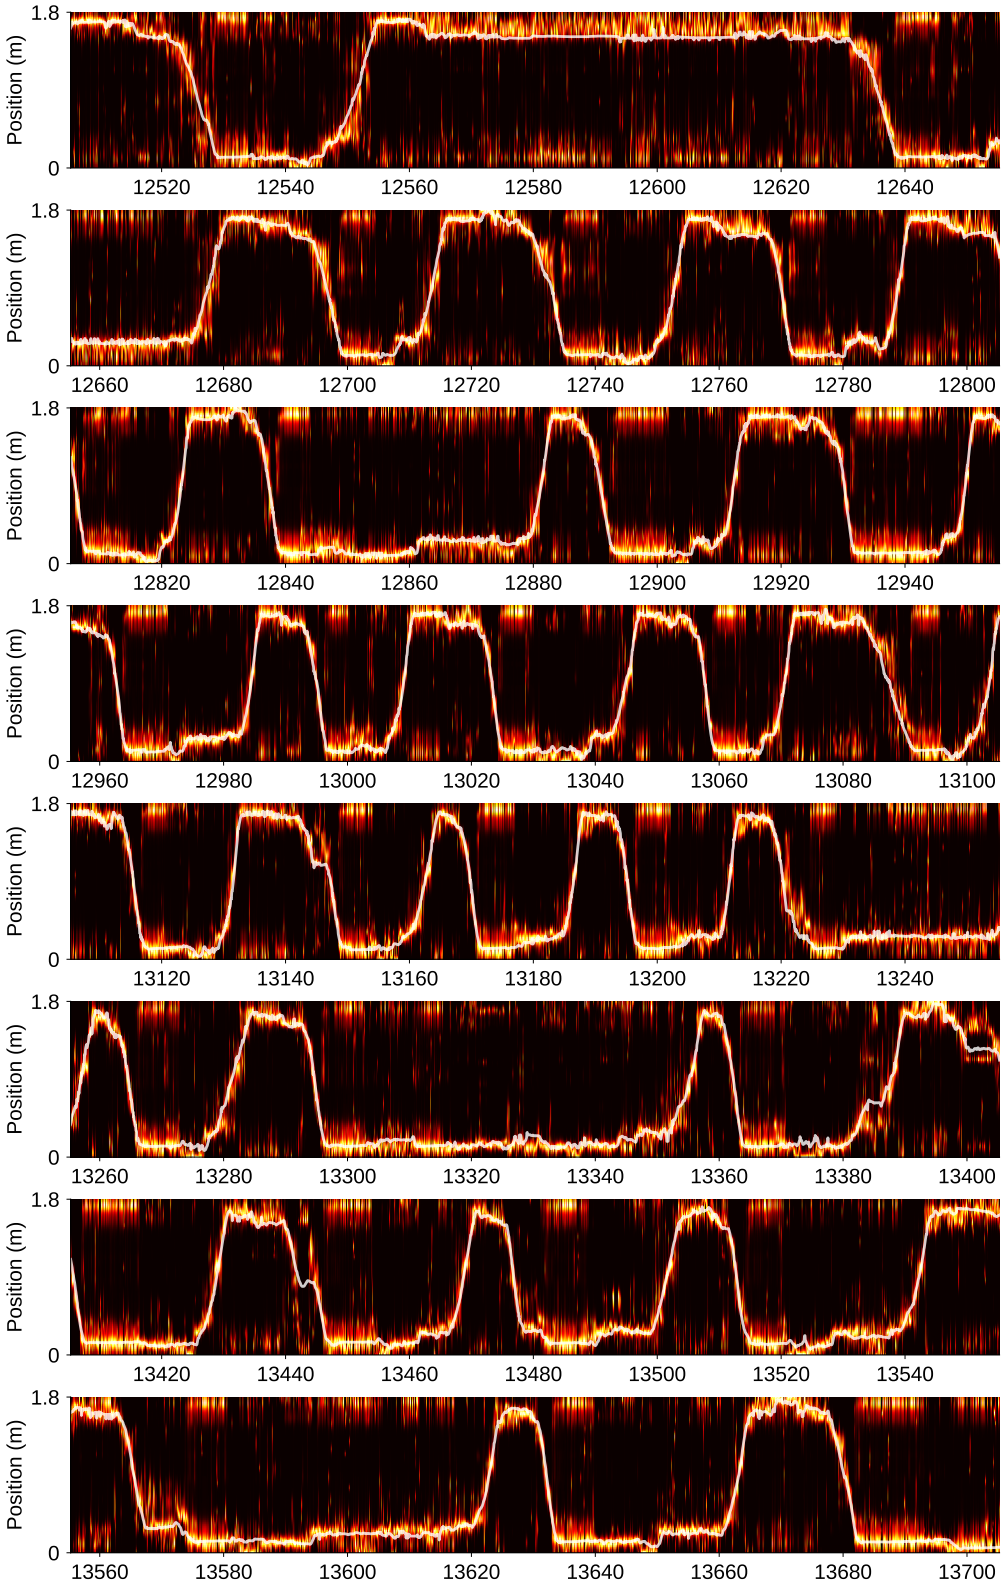

(continues in the next page)

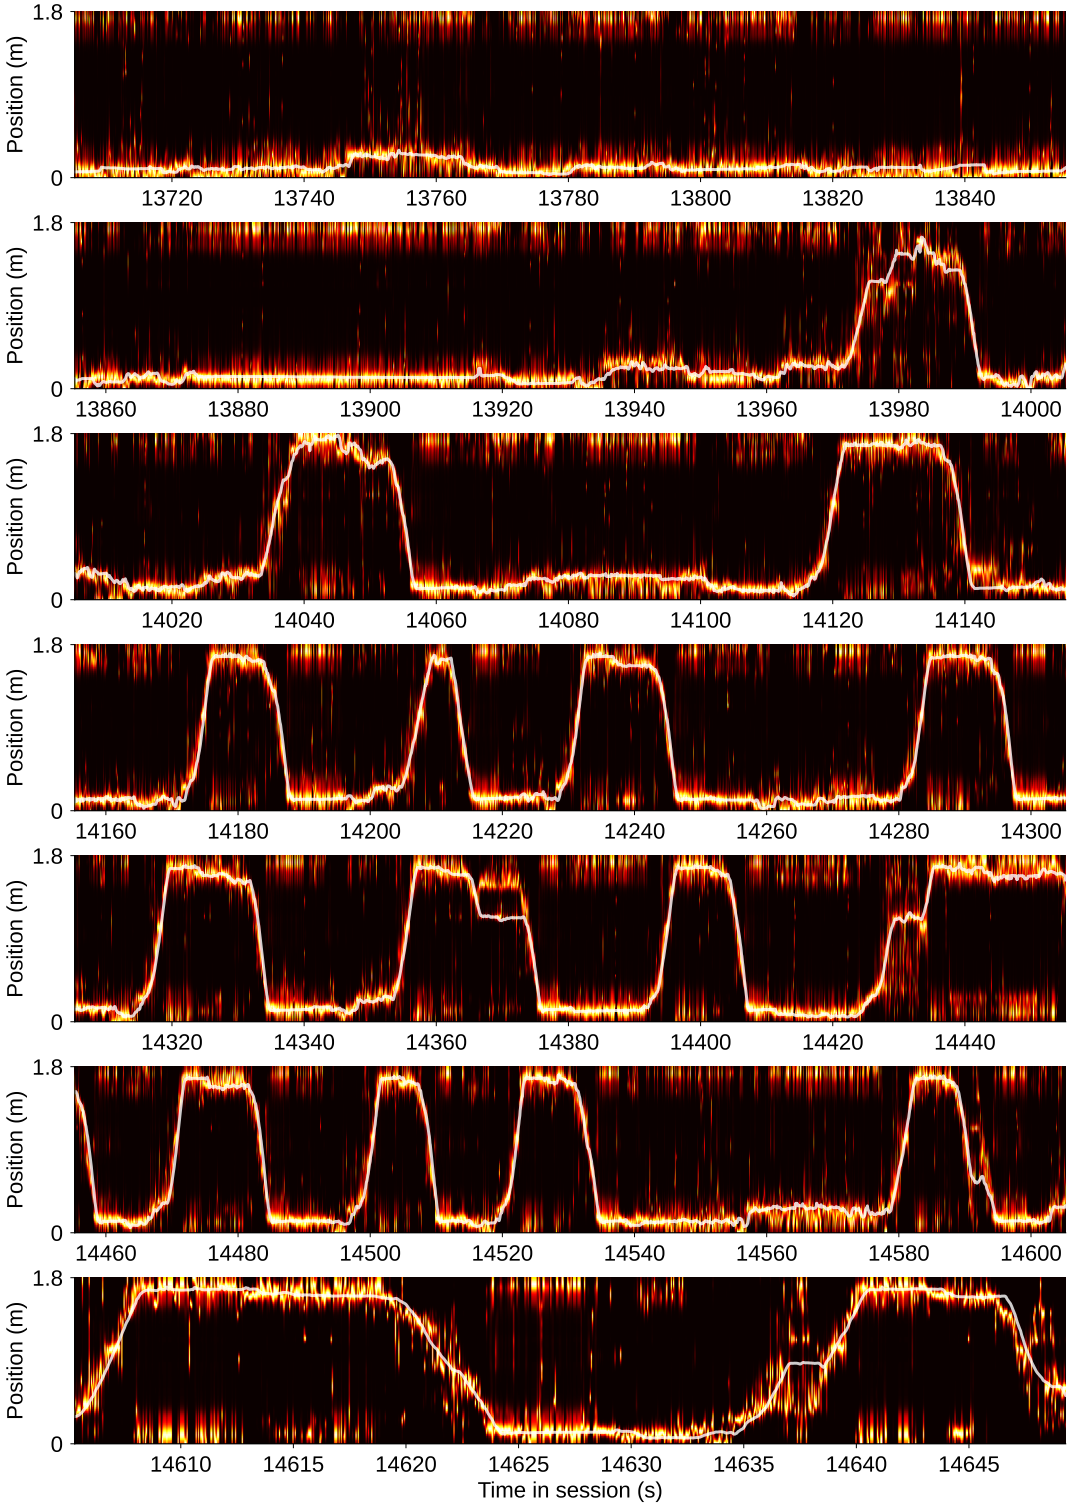

Figure S11: Decoding run position using bidirectional place fields. Plotting conventions follow Fig. S9. Same session as in Fig. S10.

B Mathematical notes on the model

B.1 Derivation of posterior of the switching HMMs

We consider a regime-switching hidden Markov model. At each discrete time point t, an observed variable $x_t \in \mathcal{X}$ is generated by a latent Markov chain $z_t \in \{1, \ldots, S\}$, whose transition probabilities are modulated by a discrete regime variable $c_t \in \{1, \ldots, C\}$. The joint evolution is specified by

$$p(x_t, z_t, c_t \mid z_{t-1}, c_{t-1}) = p(x_t \mid z_t)\, p(z_t \mid z_{t-1}, c_t, c_{t-1})\, \Pr(c_t \mid c_{t-1})$$

where $p(x_t \mid z_t)$ is the emission distribution, $p(z_t \mid z_{t-1}, c_t, c_{t-1})$ is the state-transition kernel conditional on current and previous regime, and $\Pr(c_t \mid c_{t-1})$ is the regime-switching matrix.

Forward pass. We consider the forward message,

$$\alpha_t(z_t, c_t) \equiv p(z_t, c_t \mid x_{1:t})$$

By Bayes' rule,

$$p(z_t, c_t \mid x_{1:t}) = \frac{p(x_t, z_t, c_t \mid x_{1:t-1})}{p(x_t \mid x_{1:t-1})} = \frac{p(x_t \mid z_t, c_t, x_{1:t-1})\, p(z_t, c_t \mid x_{1:t-1})}{p(x_t \mid x_{1:t-1})}$$

The Markovian assumption and the emission independence give $p(x_t \mid z_t, c_t, x_{1:t-1}) = p(x_t \mid z_t, c_t) = p(x_t \mid z_t)$. Thus,

$$p(z_t, c_t \mid x_{1:t}) = \frac{p(x_t \mid z_t)\, p(z_t, c_t \mid x_{1:t-1})}{p(x_t \mid x_{1:t-1})} \propto p(x_t \mid z_t)\, p(z_t, c_t \mid x_{1:t-1})$$

Using the Chapman-Kolmogorov equation, we can express the one-step prior $p(z_t, c_t \mid x_{1:t-1})$ in terms of α_{t-1}. By marginalizing over the previous latent pair (z_{t-1}, c_{t-1}),

$$\begin{aligned} p(z_t, c_t \mid x_{1:t-1}) &= \sum_{z_{t-1}} \sum_{c_{t-1}} p(z_t, c_t \mid z_{t-1}, c_{t-1})\, p(z_{t-1}, c_{t-1} \mid x_{1:t-1}) \\ &= \sum_{z_{t-1}, c_{t-1}} p(z_t \mid z_{t-1}, c_t, c_{t-1})\, \Pr(c_t \mid c_{t-1})\, \alpha_{t-1}(z_{t-1}, c_{t-1}) \end{aligned}$$

Combining the prior into the Bayes proportionality gives the usual recursion:

$$\boxed{\alpha_t(z_t, c_t) \propto p(x_t \mid z_t) \sum_{z_{t-1}, c_{t-1}} p(z_t \mid z_{t-1}, c_t, c_{t-1})\, \Pr(c_t \mid c_{t-1})\, \alpha_{t-1}(z_{t-1}, c_{t-1})}$$

The initialization of $\alpha_t(z_t, c_t)$ can be specified as $\alpha_1(z_1, c_1) \propto p(x_1 \mid z_1)\, \pi_z(z_1 \mid c_1)\, \pi_c(c_1)$, where π_c and $\pi_z(\cdot \mid c_1)$ are the regime and state priors. The normalizing constant is $p(x_t \mid x_{1:t-1}) = \sum_{z_t, c_t} p(x_t \mid z_t) \sum_{z_{t-1}, c_{t-1}} p(z_t \mid z_{t-1}, c_t, c_{t-1})\, \Pr(c_t \mid c_{t-1})\, \alpha_{t-1}(z_{t-1}, c_{t-1})$, so that $\sum_{z_t, c_t} \alpha_t(z_t, c_t) = 1$.

Backward pass. We consider the backward message,

$$\beta_t(z_t, c_t) \equiv p(x_{t+1:T} \mid z_t, c_t)$$

By marginalizing over the next latent pair (z_{t+1}, c_{t+1}),

$$\begin{aligned}\beta_t(z_t, c_t) &= p(x_{t+1:T} \mid z_t, c_t) \\ &= \sum_{c_{t+1}=1}^{C} \sum_{z_{t+1}=1}^{S} p(x_{t+1:T}, z_{t+1}, c_{t+1} \mid z_t, c_t)\end{aligned}$$

Apply the chain rule to factorize,

$$\begin{aligned}p(x_{t+1:T}, z_{t+1}, c_{t+1} \mid z_t, c_t) &= p(x_{t+2:T} \mid x_{t+1}, z_{t+1}, c_{t+1}, z_t, c_t) \\ &\quad \times p(x_{t+1} \mid z_{t+1}, c_{t+1}, z_t, c_t) \\ &\quad \times p(z_{t+1}, c_{t+1} \mid z_t, c_t)\end{aligned}$$

The future observations $x_{t+2:T}$ are conditionally independent of (z_t, c_t, x_{t+1}) given (z_{t+1}, c_{t+1}),

$$p(x_{t+2:T} \mid x_{t+1}, z_{t+1}, c_{t+1}, z_t, c_t) = p(x_{t+2:T} \mid z_{t+1}, c_{t+1}) = \beta_{t+1}(z_{t+1}, c_{t+1})$$

The emission depends only on z_{t+1}, not on c_{t+1} or past states,

$$p(x_{t+1} \mid z_{t+1}, c_{t+1}, z_t, c_t) = p(x_{t+1} \mid z_{t+1})$$

The joint transition factorizes as

$$p(z_{t+1}, c_{t+1} \mid z_t, c_t) = \Pr(c_{t+1} \mid c_t)\, p(z_{t+1} \mid z_t, c_{t+1}, c_t)$$

Combine and obtain the recursion,

$$\boxed{\beta_t(z_t, c_t) = \sum_{c_{t+1}} \sum_{z_{t+1}} \Pr(c_{t+1} \mid c_t)\, p(z_{t+1} \mid z_t, c_{t+1}, c_t)\, p(x_{t+1} \mid z_{t+1})\, \beta_{t+1}(z_{t+1}, c_{t+1})}$$

The initialization is $\beta_T(z_T, c_T) = 1$ for all (z_T, c_T).

Model posterior. At time t, the posterior is defined as

$$\gamma_t(z_t, c_t) \equiv p(z_t, c_t \mid x_{1:T})$$

Express it as the full joint,

$$p(z_t, c_t \mid x_{1:T}) = \frac{p(z_t, c_t, x_{1:T})}{p(x_{1:T})} = \frac{p(x_{t+1:T} \mid z_t, c_t, x_{1:t})\, p(z_t, c_t, x_{1:t})}{p(x_{1:T})}$$

By the Markov structure,

$$p(x_{t+1:T} \mid z_t, c_t, x_{1:t}) = p(x_{t+1:T} \mid z_t, c_t) \equiv \beta_t(z_t, c_t)$$

Likewise,

$$p(z_t, c_t, x_{1:t}) = p(z_t, c_t \mid x_{1:t})\, p(x_{1:t}) = \alpha_t(z_t, c_t)\, p(x_{1:t})$$

Putting these together,

$$p(z_t, c_t \mid x_{1:T}) = \frac{\beta_t(z_t, c_t)\ [\alpha_t(z_t, c_t)\, p(x_{1:t})]}{p(x_{1:T})} \propto \alpha_t(z_t, c_t)\, \beta_t(z_t, c_t)$$

Or

$$\boxed{\gamma_t(z_t, c_t) \propto \alpha_t(z_t, c_t)\, \beta_t(z_t, c_t)}$$

where the constant of proportionality $p(x_{1:t})/p(x_{1:T})$ is fixed by $\sum_{z_t, c_t} \gamma_t(z_t, c_t) = 1$.

B.2 Discussion on the estimation bias of drift and diffusion parameters

For simplicity, we perform analysis based on HMM, and briefly discuss switching HMM later. Let the 1-D track be the closed interval with grid points $z_t \in \{1, \ldots, S\}$ and spacing Δz. Time is discretized in steps of size Δt. The transition matrix is parameterized by a drift-diffusion step:

$$K_{ij}(\lambda, \sigma) \propto \mathcal{N}\left(z^j; z^i + \lambda \Delta t, \sigma^2 \Delta t\right) \Delta z$$

where $\mathcal{N}(\cdot; \mu, \nu)$ is the Gaussian density, λ is the drift parameter, and σ is the diffusion scale. Observations $\mathbf{x}_{0:T}$ are replay spike trains. The emission probabilities $p(\mathbf{x}_t \mid z_t = z^j) = B_j(\mathbf{x}_t)$ are specified by place fields fit from running data. The marginal likelihood $L(\lambda, \sigma)$ is evaluated with the forward algorithm, and we maximize log-likelihood $\ell(\lambda, \sigma) = \log L(\lambda, \sigma)$ to obtain the estimated parameters. For notation in derivatives, we use the usual marginal probabilities obtained via forward-backward algorithm:

$$\gamma_t(i) = \Pr(z_t = z^i \mid \mathbf{x}_{0:T}), \qquad \xi_t(i, j) = \Pr(z_t = z^i, z_{t+1} = z^j \mid \mathbf{x}_{0:T})$$

$\hat{\lambda}$ **is unbiased**. Let $\ell(\theta) = \log p(\mathbf{x}_{0:T} \mid \theta)$ with $\theta = (\lambda, \sigma)$. We start from,

$$\ell(\theta) = \log \sum_{z_{0:T}} p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta)$$

Differentiate using $\partial_\lambda p = p\, \partial_\lambda \log p$:

$$\begin{aligned}
\frac{\partial \ell}{\partial \lambda} &= \frac{1}{p(\mathbf{x}_{0:T} \mid \theta)} \frac{\partial}{\partial \lambda} \sum_{z_{0:T}} p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) \\
&= \frac{1}{p(\mathbf{x}_{0:T} \mid \theta)} \sum_{z_{0:T}} p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) \frac{\partial}{\partial \lambda} \log p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) \\
&= \sum_{z_{0:T}} \underbrace{\frac{p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta)}{p(\mathbf{x}_{0:T} \mid \theta)}}_{= p(z_{0:T} \mid \mathbf{x}_{0:T}, \theta)} \frac{\partial}{\partial \lambda} \log p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) \\
&= \mathbb{E}_{p(z_{0:T} \mid \mathbf{x}_{0:T}, \theta)} \left[\frac{\partial}{\partial \lambda} \log p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) \right]
\end{aligned}$$

The complete-data joint distribution factors as,

$$p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) = \underbrace{\pi_{z_0}}_{\text{init}} \left(\prod_{t=0}^{T-1} \underbrace{K_{z_t z_{t+1}}(\lambda, \sigma)}_{\text{transitions}} \right) \left(\prod_{t=0}^{T} \underbrace{B_{z_t}(\mathbf{x}_t)}_{\text{emissions}} \right)$$

Taking logs,

$$\log p(\mathbf{x}, z \mid \theta) = \log \pi_{z_0} + \sum_{t=0}^{T-1} \log K_{z_t z_{t+1}}(\lambda, \sigma) + \sum_{t=0}^{T} \log B_{z_t}(\mathbf{x}_t)$$

When differentiating w.r.t. λ,

$$\frac{\partial}{\partial \lambda} \log p(\mathbf{x}, z \mid \theta) = \sum_{t=0}^{T-1} \frac{\partial}{\partial \lambda} \log K_{z_t z_{t+1}}(\lambda, \sigma)$$

Plug this into the expectation:

$$\frac{\partial \ell}{\partial \lambda} = \sum_{t=0}^{T-1} \mathbb{E}_{p(z|\mathbf{x},\theta)} \left[\frac{\partial}{\partial \lambda} \log K_{z_t z_{t+1}}(\lambda, \sigma) \right]$$

Write the expectation using the pairwise marginals $\xi_t(i,j) = \Pr(z_t = i,\, z_{t+1} = j \mid \mathbf{x}_{0:T}, \theta)$ obtained by forward–backward:

$$\frac{\partial \ell}{\partial \lambda} = \sum_{t=0}^{T-1} \sum_{i,j} \xi_t(i,j) \, \frac{\partial}{\partial \lambda} \log K_{ij}(\lambda, \sigma)$$

For a Gaussian transition matrix,

$$K_{ij}(\lambda, \sigma) \;\approx\; \underbrace{\frac{\Delta z}{\sqrt{2\pi\sigma^2 \Delta t}}}_{\text{const. in } \lambda} \exp\left[-\frac{\left(z^j - (z^i + \lambda \Delta t)\right)^2}{2\sigma^2 \Delta t} \right]$$

Hence,

$$\log K_{ij}(\lambda, \sigma) = \text{const} - \frac{\left(z^j - (z^i + \lambda \Delta t)\right)^2}{2\sigma^2 \Delta t}$$

Differentiate w.r.t. λ. Let $\mu_i(\lambda) = z^i + \lambda \Delta t$ and $v = \sigma^2 \Delta t$.

$$\begin{aligned}
\frac{\partial}{\partial \lambda} \log K_{ij} &= \frac{\partial}{\partial \mu} \left[-\frac{(z^j - \mu)^2}{2v} \right] \Bigg|_{\mu = \mu_i} \cdot \frac{\partial \mu_i}{\partial \lambda} \\
&= \frac{(z^j - \mu_i)}{v} \cdot \Delta t \\
&= \frac{z^j - z^i - \lambda \Delta t}{\sigma^2}
\end{aligned}$$

Therefore,

$$\frac{\partial \ell}{\partial \lambda} = \frac{1}{\sigma^2} \sum_{t=0}^{T-1} \sum_{i,j} \xi_t(i,j) \left(z^j - z^i - \lambda \Delta t \right)$$

Setting the score function to 0 gives,

$$\hat{\lambda} = \frac{1}{T \Delta t} \sum_{t=0}^{T-1} \sum_{i,j} \xi_t(i,j) \, (z^j - z^i)$$

Using $\xi_t(i,j) = \Pr(z_t = z^i, z_{t+1} = z^j \mid \mathbf{x}_{0:T}, \theta)$,

$$\sum_{i,j} \xi_t(i,j) \, (z^j - z^i) = \mathbb{E}[z_{t+1} - z_t \mid \mathbf{x}_{0:T}, \theta]$$

If we evaluate the posteriors at the true parameters (λ^*, σ^*), by the law of iterated expectations,

$$\mathbb{E}_{\mathbf{x}} \Big[\mathbb{E}[z_{t+1} - z_t \mid \mathbf{x}, \lambda^*, \sigma^*] \Big] = \mathbb{E}[z_{t+1} - z_t] = \lambda^* \Delta t$$

Hence,

$$\mathbb{E}[\hat{\lambda}] = \frac{1}{T \Delta t} \sum_t \lambda^* \Delta t = \lambda^*$$

This means the MLE estimate has zero bias.

$\hat{\sigma}$ **is biased downward**. With $\theta = (\lambda, \sigma)$ and $\ell(\theta) = \log p(\mathbf{x}_{0:T} \mid \theta)$,

$$\frac{\partial \ell}{\partial \sigma^2} = \mathbb{E}_{p(z_{0:T}|\mathbf{x}_{0:T},\theta)} \left[\frac{\partial}{\partial \sigma^2} \log p(\mathbf{x}_{0:T}, z_{0:T} \mid \theta) \right]$$

Using the factorization,

$$\log p(\mathbf{x}, z \mid \theta) = \log \pi_{z_0} + \sum_{t=0}^{T-1} \log K_{z_t z_{t+1}}(\lambda, \sigma) + \sum_{t=0}^{T} \log B_{z_t}(\mathbf{x}_t)$$

Only the transition terms depend on σ. Hence,

$$\frac{\partial \ell}{\partial \sigma^2} = \sum_{t=0}^{T-1} \mathbb{E}_{p(z|\mathbf{x},\theta)} \left[\frac{\partial}{\partial \sigma^2} \log K_{z_t z_{t+1}}(\lambda, \sigma) \right] = \sum_{t=0}^{T-1} \sum_{i,j} \xi_t(i,j) \frac{\partial}{\partial \sigma^2} \log K_{ij}(\lambda, \sigma)$$

With Gaussian transition kernel,

$$\log K_{ij}(\lambda, \sigma) = -\tfrac{1}{2} \log(\sigma^2 \Delta t) - \frac{(z^j - z^i - \lambda \Delta t)^2}{2\sigma^2 \Delta t}$$

Let $A_{ij} := z^j - z^i - \lambda \Delta t$. Then differentiate $\log K_{ij}$ w.r.t. σ^2,

$$\frac{\partial}{\partial \sigma^2} \log K_{ij} = -\frac{1}{2\sigma^2} + \frac{A_{ij}^2}{2\sigma^4 \Delta t}$$

Set the score to zero and solve,

$$0 = \sum_{t,i,j} \xi_t(i,j) \left(-\frac{1}{2\sigma^2} + \frac{A_{ij}^2}{2\Delta t\, \sigma^4} \right)$$

$$0 = \sum_{t,i,j} \xi_t(i,j) \left(-\sigma^2 + \frac{A_{ij}^2}{\Delta t} \right)$$

$$\sigma^2 \sum_{t,i,j} \xi_t(i,j) = \sum_{t,i,j} \xi_t(i,j) \, \frac{A_{ij}^2}{\Delta t}$$

Hence we have the variance MLE estimator,

$$\hat{\sigma}^2 = \frac{\sum_{t,i,j} \xi_t(i,j) \, (z^j - z^i - \lambda \Delta t)^2}{\sum_{t,i,j} \xi_t(i,j) \, \Delta t}$$

Using $\sum_{i,j} \xi_t(i,j) = 1$,

$$\hat{\sigma}^2 = \frac{1}{T \Delta t} \sum_{t=0}^{T-1} \mathbb{E}\Big[(z_{t+1} - z_t - \hat{\lambda} \Delta t)^2 \,\Big|\, \mathbf{x} \Big]$$

If λ were known and $\hat{\lambda}$ is equal to λ^*, then by the law of iterated expectations,

$$\mathbb{E}[\hat{\sigma}^2] = \frac{1}{T \Delta t} \sum_{t} \mathbb{E}\big[(z_{t+1} - z_t - \lambda^* \Delta t)^2 \big] = \sigma^2$$

If $\hat{\lambda}$ is estimated from the same data, we evaluate at the true parameters $(\lambda^*, (\sigma^*)^2)$.

Define increments,

$$y_t \equiv z_{t+1} - z_t \sim \mathcal{N}(\lambda^* \Delta t,\, (\sigma^*)^2 \Delta t), \quad \bar{y} = \frac{1}{T}\sum_{t=1}^{T} y_t$$

For simplicity, we further approximate $\hat{\lambda}$ as $\hat{\lambda} = \bar{y}/\Delta t$ under the assumption that emissions are informative and $\xi_t(i,j)$ are sharply peaked. We then calculate moments,

$$\mathrm{Var}(\hat{\lambda}) = \frac{1}{(\Delta t)^2}\,\mathrm{Var}\left(\frac{1}{T}\sum_t y_t\right) = \frac{1}{(\Delta t)^2}\cdot\frac{1}{T}\,\mathrm{Var}(y_t) = \frac{(\sigma^*)^2}{T\Delta t}$$

$$\mathrm{Cov}(\hat{\lambda}, y_t) = \frac{1}{\Delta t}\,\mathrm{Cov}(\bar{y}, y_t) = \frac{1}{\Delta t}\cdot\frac{1}{T}\,\mathrm{Var}(y_t) = \frac{(\sigma^*)^2}{T}$$

By the law of iterated expectations,

$$\mathbb{E}[\hat{\sigma}^2] = \frac{1}{T\Delta t}\sum_{t=1}^{T}\mathbb{E}\Big[(y_t - \hat{\lambda}\Delta t)^2\Big]$$

and

$$\begin{aligned}
\mathbb{E}\Big[(y_t - \hat{\lambda}\Delta t)^2\Big] &= \mathbb{E}\big[(y_t - \lambda^*\Delta t)^2\big] + \Delta t^2\,\mathbb{E}\Big[(\hat{\lambda} - \lambda^*)^2\Big] - 2\Delta t\,\mathbb{E}\Big[(\hat{\lambda} - \lambda^*)(y_t - \lambda^*\Delta t)\Big] \\
&= (\sigma^*)^2\Delta t + \Delta t^2\,\frac{(\sigma^*)^2}{T\Delta t} - 2\Delta t\cdot\frac{(\sigma^*)^2}{T} \\
&= (\sigma^*)^2\Delta t\left(1 - \frac{1}{T}\right)
\end{aligned}$$

Thus,

$$\mathbb{E}\big[\hat{\sigma}^2\big] = \frac{T-1}{T}\,(\sigma^*)^2$$

This is known as Bessel effect and can explain a noticeable downward bias when each replay event contains only a handful of time bins.

Extension to switching HMM. We consider a switching HMM with a discrete regime c_t that controls the transition kernel, and only one regime (noted as $s = \star$) has drift-diffusion parameters (λ, σ) while the others are fixed (i.e., stationary or uniform). Extending the above bias analysis should be relatively straightforward.

In the above HMM setup, the score for λ and σ come entirely from the transition terms and use pairwise marginals,

$$\xi_t(i,j) = \Pr(z_t = i, z_{t+1} = j \mid x)$$

In the switching HMM, the log-joint includes the regime chain and the transition kernel is $K_{ij}^{(c_t)}$. The score for $\theta = (\lambda, \sigma)$ only accrues when $c_t = \star$. Practically, we can replace $\xi_t(i,j)$ by regime-conditioned pairwise marginals,

$$\xi_t^{(\star)}(i,j) = \Pr(c_t = \star,\, z_t = i,\, z_{t+1} = j \mid \mathbf{x}, \theta)$$

and ignore transitions from other regimes in the derivative, because their kernels are not relevant to (λ, σ).

A subsequent change in derivation is that HMM formulas become weighted by the posterior probability that the $\star$ regime is active at time t:

$$\hat{\lambda} = \frac{\sum_t \sum_{i,j} \xi_t^{(\star)}(i,j)\,(z_j - z_i)}{\Delta t \ \sum_t w_t^{(\star)}} \qquad \hat{\sigma}^2 = \frac{\sum_t \sum_{i,j} \xi_t^{(\star)}(i,j)\,\left(z_j - z_i - \hat{\lambda}\Delta t\right)^2}{\Delta t \ \sum_t w_t^{(\star)}}$$

with the regime-posterior weight,

$$w_t^{(\star)} = \sum_{i,j} \xi_t^{(\star)}(i,j) = \Pr(c_t = \star \mid \mathbf{x}).$$

Similar to HMM case, for $\hat{\lambda}$, under the correctly specified switching model and evaluate the posteriors at the true parameters $(\lambda^*, (\sigma^*)^2)$,

$$\mathbb{E}[\hat{\lambda}] = \lambda^*$$

For $\hat{\sigma}^2$, the total time step T in above HMM expressions should be replaced by Kish's effective sample size,

$$N_{\text{eff}}^{(\star)} \equiv \frac{\left(\sum_t w_t^{(\star)}\right)^2}{\sum_t {w_t^{(\star)}}^2}$$

The expectation of $(\sigma^*)^2$ estimator becomes:

$$\mathbb{E}[\hat{\sigma}^2] = \frac{N_{\text{eff}}^{(\star)} - 1}{N_{\text{eff}}^{(\star)}}\,(\sigma^*)^2$$

I.e., $\hat{\lambda}$ is unbiased and the finite-sample downward bias for $\hat{\sigma}^2$, with T replaced by $N_{\text{eff}}^{(\star)}$.

References

[1] Kechen Zhang, Iris Ginzburg, Bruce L McNaughton, and Terrence J Sejnowski. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of neurophysiology*, 79(2):1017–1044, 1998.

[2] Delia Silva, Ting Feng, and David J Foster. Trajectory events across hippocampal place cells require previous experience. *Nature neuroscience*, 18(12):1772–1779, 2015.

[3] Masahiro Takigawa, Marta Huelin Gorriz, Margot Tirole, and Daniel Bendor. Evaluating hippocampal replay without a ground truth. *ELife*, 13:e85635, 2024.

[4] Usman Farooq, Jeremie Sibille, Kefei Liu, and George Dragoi. Strengthened temporal coordination within pre-existing sequential cell assemblies supports trajectory replay. *Neuron*, 103(4):719–733, 2019.

[5] Matthijs AA van der Meer, Caleb Kemere, and Kamran Diba. Progress and issues in second-order analysis of hippocampal replay. *Philosophical Transactions of the Royal Society B*, 375(1799):20190238, 2020.

[6] Brad E Pfeiffer and David J Foster. Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science*, 349(6244):180–183, 2015.