

A large majority of awake hippocampal sharp-wave ripples feature spatial trajectories with momentum

Highlights

- Place cells replay spatial trajectories during sharp-wave ripples
- A probabilistic approach reveals spatiotemporal dynamics of replayed trajectories
- The approach reveals that almost all sharp-wave ripples encode spatial trajectories
- The dynamics of most replayed trajectories feature ballistic motion, or “momentum”

Authors

Emma L. Krause, Jan Drugowitsch

Correspondence

jan_drugowitsch@hms.harvard.edu

In brief

Krause and Drugowitsch use a probabilistic approach to characterize the spatial dynamics encoded by hippocampal place cell activity of rodents during sharp-wave ripple events and show that these dynamics feature momentum, like the animal’s movement.



Article

A large majority of awake hippocampal sharp-wave ripples feature spatial trajectories with momentum

Emma L. Krause¹ and Jan Drugowitsch^{1,2,*}¹Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA²Lead contact*Correspondence: jan_drugowitsch@hms.harvard.edu<https://doi.org/10.1016/j.neuron.2021.11.014>**SUMMARY**

During periods of rest, hippocampal place cells feature bursts of activity called sharp-wave ripples (SWRs). Heuristic approaches have revealed that a small fraction of SWRs appear to “simulate” trajectories through the environment, called awake hippocampal replay. However, the functional role of a majority of these SWRs remains unclear. We find, using Bayesian model comparison of state-space models to characterize the spatiotemporal dynamics embedded in SWRs, that almost all SWRs of foraging rodents simulate such trajectories. Furthermore, these trajectories feature momentum, or inertia in their velocities, that mirrors the animals’ natural movement, in contrast to replay events during sleep, which lack such momentum. Last, we show that past analyses of replayed trajectories for navigational planning were biased by the heuristic SWR sub-selection. Our findings thus identify the dominant function of awake SWRs as simulating trajectories with momentum and provide a principled foundation for future work on their computational function.

INTRODUCTION

Planning through mental simulations, or the anticipation of future action-outcome sequences, is a powerful mechanism for improving action selection (Sutton and Barto, 1998). Strikingly, rodents performing spatial navigation tasks appear to perform such mental simulations (Buzsáki, 2015; Carr et al., 2011). While the animal is moving, hippocampal place cells exhibit spatially localized firing patterns, such that a population of place cells represents the animal’s current location in the environment (O’Keefe and Dostrovsky, 1971; Figures 1A and 1B). During a fraction of sharp-wave ripples (SWRs), which are population bursts of activity associated with brief pauses in the animal’s movement, these place cells appear to shift to generating “simulated” trajectories through the environment (Davidson et al., 2009; Diba and Buzsáki, 2007; Tingley and Peyrache, 2020; Figures 1A–1C). These “awake hippocampal replay events” have been proposed to support a range of computational functions, such as memory storage (Wilson and McNaughton, 1994), recall (Gillespie et al., 2021; Shin et al., 2019; Xu et al., 2019), and planning (Jadhav et al., 2012; Pfeiffer and Foster, 2013; Singer et al., 2013). However, determining the precise computational function of these events has been challenging (Joo and Frank, 2018; Mattar and Daw, 2018), and it remains unclear why they constitute only a small fraction of SWRs (Tingley and Peyrache, 2020). A critical foundation for assessing the computational role of SWRs is a comprehensive and systematic characterization of the spatiotemporal dynamics of the trajectories they encode and how these dynamics relate to natural movement in the environment.

Two main challenges have hampered a systematic characterization of the spatiotemporal dynamics of replay events. First, most studies of hippocampal replay are performed in one-dimensional (1D) maze environments, which greatly constrains the possible set of observed dynamics to linear trajectories. This obscures subtle task-specific details in the trajectories’ dynamics and makes it hard to identify certain features in the underlying trajectories, such as their relation to natural movement (Stella et al., 2019). Second, established methods for identifying replay trajectories within SWRs use heuristics whose implicit assumptions cause them to identify only a subset of especially salient trajectories (Tingley and Peyrache, 2020). Commonly, they declare SWRs as replay events only if the position sequence decoded by maximum likelihood satisfies a restrictive set of criteria. For example, trajectories might need to be linear trajectories in 1D environments (Davidson et al., 2009; Diba and Buzsáki, 2007) or have minimum start-to-end and maximum consecutive position distances in two-dimensional (2D) environments (Pfeiffer and Foster, 2013, 2015). Particularly in 2D environments, which more closely resemble the rodents’ natural habitat, this leads to discarding more than two-thirds of SWRs as non-trajectory events (Pfeiffer and Foster, 2015; Stella et al., 2019). Does this imply that the discarded SWRs in fact do not encode trajectories but instead signal other events, or that these SWRs just do not conform to the assumed classification criteria? Furthermore, if the discarded SWRs indeed encode trajectories, to what degree is their characterization and subsequent analysis biased by discarding them?

To address these questions, we move away from traditional heuristics of replay event classification and instead take a



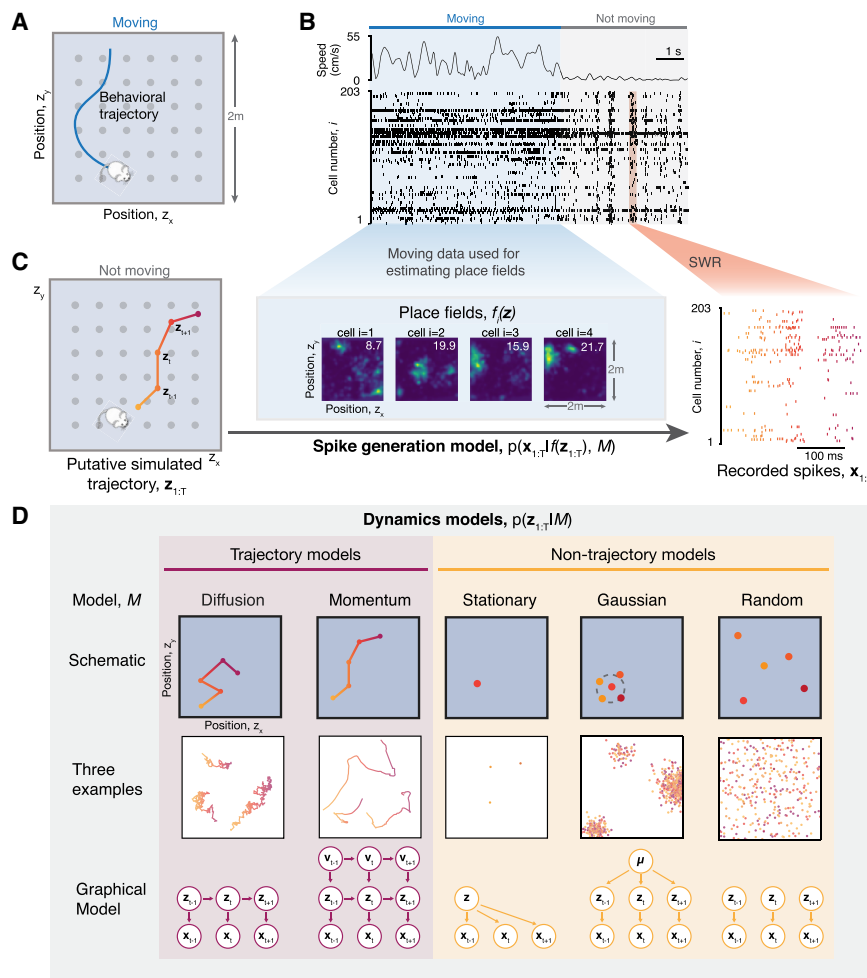


Figure 1. State-space models for characterizing the spatiotemporal structure of SWRs

(A) Rats foraged in a 2 m x 2 m open-field arena for food hidden in wells (shaded gray circles). The blue trace illustrates a schematic behavioral trajectory of the rat.

(B) Representative sample from data recording. Top trace shows the velocity of the animal, with an initial period in which it is moving (blue shading), followed by a period in which it is not moving (gray shading). The raster plot shows associated spiking activity for each cell, i , over time t . SWRs occur within the period in which the rat is not moving (orange shading, example SWR).

(C) Schematic of assumed relationship between the putative simulated trajectory, $z_{1:T}$, encoded by the example SWR (left panel; z_x/z_y denote encoded x/y positions) and the recorded spikes, $x_{1:T}$ (right panel, SWR from B expanded in time; color gradient from yellow to purple indicates time). The spike generation model (middle panel) uses the place fields (four example cells; top-right number is maximum firing rate [spikes/s] estimated from spiking data during movement) to predict spike counts for each spatial position during SWRs.

(D) Considered dynamics models for characterizing the spatiotemporal structure encoded by an SWR. Dynamics models are grouped into trajectory models, which assume a continuously evolving trajectory, and non-trajectory models, which do not. We show for each dynamics model, M , a schematic depicting the assumed dynamics (top row), three example simulated trajectories generated from the dynamics of the model (middle row), and the underlying graphical model showing the statistical relationship of the variables involved (bottom row; see STAR Methods for detailed description).

probabilistic modeling approach. This approach rests on defining explicit, dynamical models for trajectories as continuous sequences of position and compares them with alternative models of non-continuous position sequences. Applied to neural tetrode recordings of rats foraging in an open-field area (Pfeiffer and Foster, 2013), our approach revealed that almost all of the observed SWRs encode spatial trajectories. Furthermore, these trajectories appeared to feature movement dynamics with momentum, that is, with inertia on their velocities. Thus, they feature dynamics beyond simple Brownian motion and comparable with how rodents actually move through the environment. This has several consequences. First, the finding that almost all SWRs encode trajectories reveals that simulating spatial trajectories is, in fact, a dominant function of SWRs. Second, the finding of momentum embedded in these trajectories implies that the mechanism generating these events needs to include a notion of velocity, generating the potential inclusion of multiple brain networks (see Discussion). Third, finding consistent momentum stands in contrast to similar replay events during sleep, which appear to lack such momentum and instead follow simpler Brownian motion (Stella et al., 2019). Thus, awake and sleep replay events might differ in both their engaged neural mecha-

nisms and their functional role. Fourth, almost all SWRs encoding trajectories implies that previous work analyzing these trajectories might have been biased by discarding the majority of them, as we demonstrate in the context of navigational planning further below. Last, our approach is generally less sensitive to noise than the established heuristics and thus should lead to cleaner, less biased, decoded trajectories that can inform further work on their role in navigational planning.

Although previous work has applied probabilistic methods to the analysis of replay events, it has done so in a different context. Location decoding from place cell population activity at isolated time points, for example, commonly relies on maximum likelihood approaches (Johnson and Redish, 2007; Zhang et al., 1998). Other past studies have applied probabilistic methods to identifying replay events but did not model the spatial dynamics of the trajectories (Linderman et al., 2016; Maboudi et al., 2018). Recent work (Denovellis et al., 2021) used a related probabilistic formulation of the position sequence encoded within an SWR but restricted itself to W-maze environments and did not distinguish between diffusion and momentum dynamics. We instead perform a systematic characterization of the spatial dynamics of replay trajectories in unconstrained 2D

environments and use a rich class of considered dynamics models for such environments that yields insights into the dynamics underlying SWRs.

RESULTS

State-space models characterize the spatiotemporal structure of SWRs

We analyzed the spatiotemporal structure of SWRs in the dataset of Pfeiffer and Foster (2013, 2015), which consists of tetrode recordings from hippocampal CA1 collected while rats foraged around a 2 m × 2 m open-field environment for a hidden food reward (Figures 1A and 1B). The rats' ability to freely roam an open field, unconstrained by the topology of a maze, was essential to our main objective of characterizing the replay events' spatiotemporal structure. Each session consists of 80–263 single units from recording sessions of 37–66 min ($n = 8$ sessions, four rats for 2 sessions each). For our analysis we used the 2,956 SWRs (Figure 1B) that Pfeiffer and Foster (2015) identified by their ripple-band power of the local field potential (mean \pm SD = 372 ± 42 SWRs per session; see STAR Methods).

We used Bayesian model comparison applied to state-space models to distinguish between different spatiotemporal dynamics encoded by each SWR. Our state-space models assume that the place cell activity sequence underlying each SWR encodes a sequence of positions (i.e., the latent state sequence) in the open-field environment and allow us to distinguish different dynamics of these position sequences from recorded spike data (Figure 1C). Their two components are (1) a spiking model that determines, for each latent position in the environment, the resulting place cell activity encoding this position and (2) a dynamics model, describing the presumed dynamical structure of each sequence of latent positions. We represent the latent position \mathbf{z}_t at each time point t by discretizing the environment into a 50×50 grid (4×4 cm bins). We do not have access to the underlying true sequence of positions over time, $\mathbf{z}_{1:T} = \mathbf{z}_1, \dots, \mathbf{z}_T$ (assuming here a sequence of length T), but instead observe the sequence of recorded spikes, $\mathbf{x}_{1:T} = \mathbf{x}_1, \dots, \mathbf{x}_T$, where \mathbf{x}_t denotes the vector of spike counts emitted by each place cell in the t th time bin. As in previous work (Davidson et al., 2009; Pfeiffer and Foster, 2015), we assume that place cell activity encodes latent position during SWRs as they do during periods of active movement. Thus, we estimated each cell i 's place field $f_i(\mathbf{z})$ from its spiking activity during movement and in turn assumed that the observed spikes during SWRs were, for some latent position \mathbf{z} , generated by draws from a Poisson distribution, independent across cells, with spike rate $f_i(\mathbf{z})$ for place cell i (see STAR Methods).

We then define a set of candidate dynamics models that describe the hypothesized spatiotemporal structure of how the latent position, \mathbf{z}_t , evolves through time and space (Figure 1D). Each dynamics model M is fully described by the probability $p(\mathbf{z}_{1:T}|M)$ that it assigns to a specific position sequence $\mathbf{z}_{1:T}$. The models differ in the probability they assign to a sequence, with the highest probability assigned to position sequences that are most compatible with the model's assumed dynamics. To capture a reasonable range of potential spatiotemporal dynamics of position sequences, we defined five models that

each make different assumptions about these dynamics (Figure 1D). This is equivalent to defining a set of hidden Markov models with different priors on the transition function (Figure 1D, bottom row). The first model describes a random walk through space: a one-step Markov model in which the position at any time point depends only on the previous position. Its dynamics generate trajectories that resemble Brownian motion through the environment, so we refer to this model as the “diffusion” model. Notably, its trajectories lack momentum and so, unlike natural movement, do not follow smooth movement through space (Figure 1D, middle row). This is rectified by our diffusion with “momentum” model, whose sequence of velocities, $\mathbf{v}_{1:T} = \mathbf{v}_1, \dots, \mathbf{v}_T$ (i.e., consecutive position changes), rather than positions, performs a continuous random walk. The position sequences associated with both models form continuously evolving trajectories, such that we refer to them collectively as “trajectory” models.

To distinguish SWRs that encode trajectories from those that do not, we defined three additional models that do not hypothesize temporal continuity of the encoded position sequences. The first, “stationary” model assumes that the latent position remains constant over time within a single SWR. The second, stationary “Gaussian” model relaxes this rather stringent assumption by allowing the latent position to be drawn at each point in time from a Gaussian with mean and variance that remain fixed within individual SWRs but can vary across them. Last, the third, “random” model assumes that the latent positions are drawn independently across time within each SWR and uniformly on the discretized environment. As none of these alternative models assumed the position encoded by SWRs to evolve along a continuous trajectory, we refer to them collectively as “non-trajectory” models.

We applied model comparison separately to each SWR (Figure 2). That is, we did not assume all SWRs to encode the same latent state dynamics but instead asked, for each SWR in the dataset separately, how likely each model was to have generated the recorded sequence of spikes. To apply the model comparison, we chose to bin observed spikes into non-overlapping time bins of 3 ms but found comparable results for other choices of the time bin size (Figure S1A). Combining the spike generation model with the different dynamics models, and averaging over all possible latent position sequences, allowed us to compute the likelihood $p(\mathbf{x}_{1:T}|M)$ of observing the given spike sequence for each dynamics model (see STAR Methods). Assuming each dynamics model is a priori equally likely, and applying Bayes' rule to this likelihood, in turn provides a posterior distribution, $p(M|\mathbf{x}_{1:T})$, of how likely each of the dynamics models is for this particular SWR (see STAR Methods). This posterior both tells us which model described the spiking data best—the model with the highest likelihood $p(\mathbf{x}_{1:T}|M)$ —and also gives us a measure of relative certainty across the dynamics models (Figure 2B). Importantly, this procedure implicitly penalized more complex dynamics models (MacKay, 1995). For example, the latent position sequence arising from a stationary model would also be compatible with a stationary Gaussian model with a small variance. However, the higher complexity of the latter makes our model comparison prefer the former, if compatible with the observed spike sequence. This applies to all dynamics models

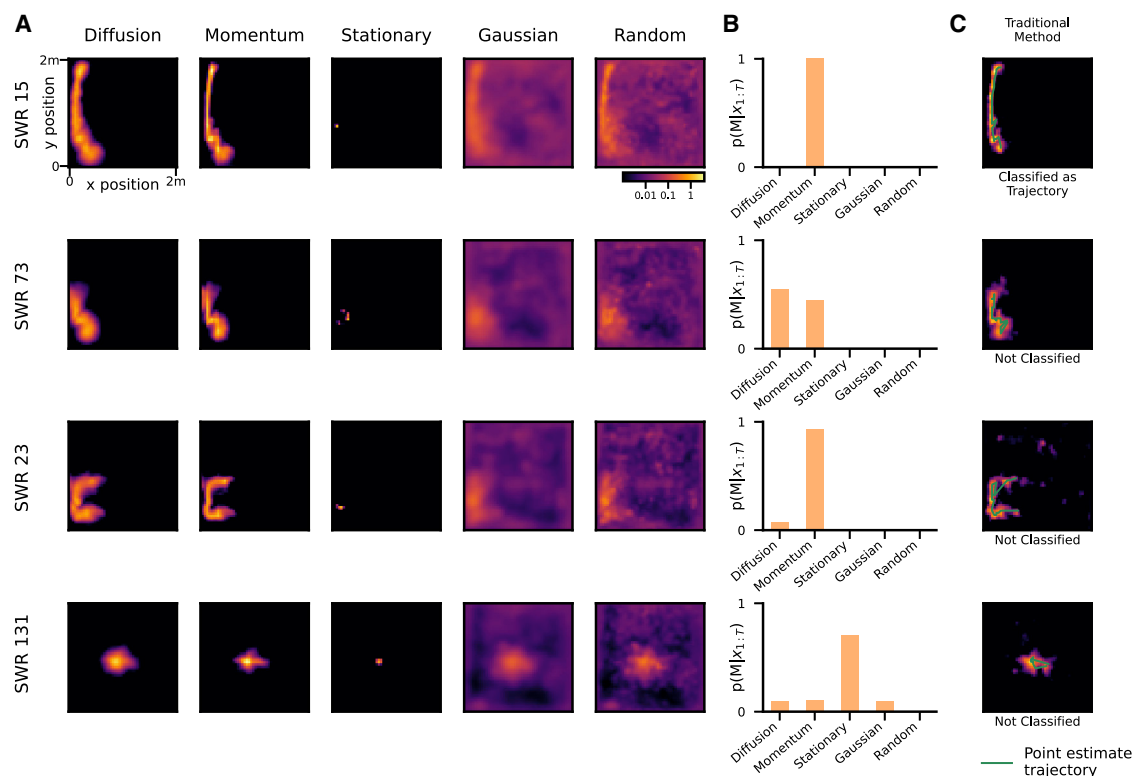


Figure 2. State-space models reveal the spatiotemporal dynamics within individual SWRs

The examples shown were chosen to illustrate different types of latent dynamics, most of which remained unclassified by the traditional method for analyzing replay events.

(A) Heatmaps visualize the decoded position under each dynamics model by $p(\mathbf{z}_t|x_{1:T}, M)$, summed over all time, $t = 1 \dots T$, for visualization. They illustrate how our approach combines the uncertain position information encoded by place cells with the stochastic position dynamics assumed by the different models. Each column shows the decoded position under one dynamics model, and each row is a representative SWR.

(B) The relative likelihood of each model to generate the recorded spikes within the SWRs shown in (A).

(C) Comparison with the traditional method for replay classification (Pfeiffer and Foster, 2015): the heatmap visualizes decoded posterior position and extracted trajectory using the traditional method for trajectory classification for each SWR. The green line indicates the extracted trajectory that is subjected to classification criteria, and the label under the heatmap indicates if the SWR was classified as a trajectory.

that can mimic other dynamics models in certain parameter regimes (see STAR Methods and Figures S7D and S7E for such model relationships); if two models explain the spike sequence equally well, our approach generally prefers the more constrained, less complex dynamics model.

Most awake SWRs feature trajectories with momentum

Applying our model comparison to all SWRs across all sessions confirmed that almost all (91.7%; see Figure S2) SWRs that were classified as trajectories by Pfeiffer and Foster (2015) using the previous, heuristic approach (which we refer to as the “traditional method” in the following) were best described by one of our trajectory models. For example, the traditional method classified SWR 15 in Figure 2 as a replay event and was best fit by our momentum model. However, we also found many previously unclassified SWRs to be best described by one of our trajectory models. For example, SWRs 23 and 73 in Figure 2 are best described by the momentum and diffusion models, respectively, though they failed the previous classification criteria because of a too small start-to-end distance of the extracted trajectory.

Indeed, the traditional method only classified 23.7% of the 2,956 SWRs as trajectories, which we will refer to as “previously classified” trajectories (see STAR Methods). Furthermore, few SWRs were best described by one of the non-trajectory models, as, for example, the stationary model for SWR 131.

We next asked how likely each dynamics model was on average to have generated individual SWRs within each session. We inferred this distribution over dynamics model by random-effects model comparison (Penny et al., 2010). In contrast to the more standard fixed-effects Bayesian model comparison that assumes all SWRs to follow the same dynamics model and infers the most likely one, random-effects model comparison assumes SWRs to be drawn from a distribution over dynamics models, making it less prone to outliers in per-SWR model likelihoods $p(\mathbf{x}_{1:T}|M)$ (Stephan et al., 2009; Figure S1B). Across sessions, random-effects analysis (Figure 3A) revealed that $84.9\% \pm 2.9\%$ (mean \pm SD across sessions) of SWRs are generated by dynamics containing momentum (exceedance probability of momentum model ≈ 1 for all sessions; the exceedance probability is the probability that the

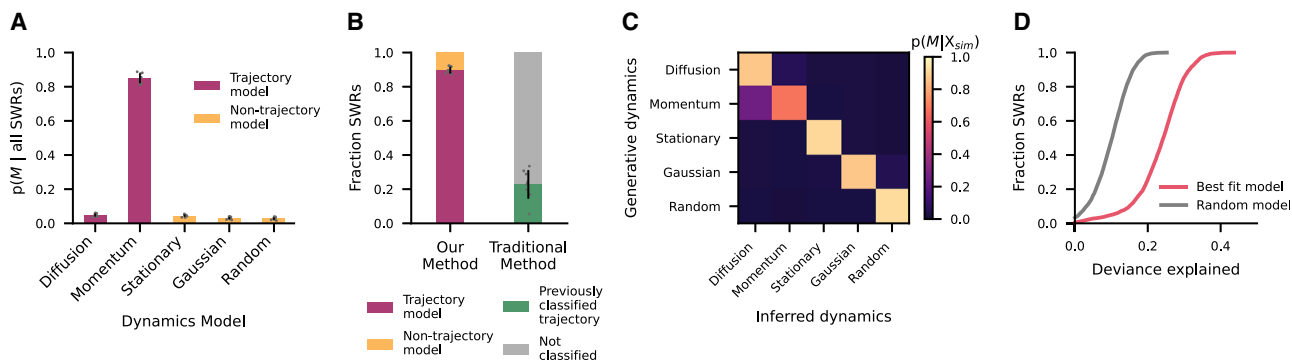


Figure 3. The vast majority of SWRs feature trajectories with momentum

(A) Inferred distribution of dynamics models underlying the observed SWRs (mean \pm SD across sessions), computed using random-effects model comparison (see main text).

(B) Comparison of our method with the traditional method, grouping the diffusion and momentum dynamics models into trajectory models and the remaining models into non-trajectory models (mean \pm SD across sessions; gray dots, individual sessions).

(C) Confusion matrix summarizing model recovery results on simulated data. Each row shows the inferred distribution of dynamics models, $p(M | \text{all simulated SWRs } X_{sim})$, given the set of simulated spiking data generated under a different dynamics model M .

(D) Cumulative histogram of deviance explained for all SWRs by the best-fitting model (pink line) and the random model (gray line). Zero percent deviance explained: predicting same average spike firing rate across all neurons/SWRs; 100% deviance explained: predicting correct spike count in each time bin.

See also [Figures S1](#) and [S2](#).

momentum model is more likely than all other models), suggesting that these SWRs include a notion of velocity in the dynamics underlying the position evolution. Additionally, $4.8\% \pm 0.9\%$ of SWRs are generated by diffusion dynamics without momentum. The remaining $10.3\% \pm 2.0\%$ are consistent with non-trajectory dynamics (stationary, $4.5\% \pm 0.7\%$; stationary Gaussian, $3.0\% \pm 0.7\%$; random, $2.8\% \pm 0.8\%$). Overall, this implies that $89.7\% \pm 2.0\%$ of SWR were identified to encode trajectories and thus contain temporal structure ([Figure 3B](#)). This stands in stark contrast to the $22.8\% \pm 8.4\%$ of SWRs that are classified as replay events by the traditional method ([Figure 3B](#)). We found similar results when applying the model comparison to high-synchrony events (HSEs; [STAR Methods](#)) rather than SWRs ([Figures S1D–S1G](#)), confirming that these results are not specific to the method used to identify neural population events of interest.

To ensure that our method did not erroneously identify temporal structure where there was none, we performed two additional checks. First, we applied the same model comparison to data in which we either separately randomly shuffled the neuron identities for each SWR or randomly shifted each neuron's place field. Both perturbations cause the momentum model to cease dominating the model comparison ([Figures S1H](#) and [S1I](#)). Second, we generated simulated spiking data under the dynamics of each of our five models and checked how reliably our method could recover the model that we used to generate the data. To match the statistics of true data, we used parameters similar to those recovered from data and generated spikes using the estimated place fields (see [STAR Methods](#)). We then asked how often the simulated place cell activity is identified to be generated by the model of the true underlying dynamics rather than any of the other models ([Figure 3C](#)). This revealed a reliable discrimination between trajectory and non-trajectory models (F score = 0.94),

thus making the spurious detection of temporal structure unlikely. Within trajectory models, the second check also illustrates the implicit model complexity penalization of Bayesian model comparison: even though the momentum model can in some parameter regimes capture diffusion dynamics (see [STAR Methods](#); [Figure S7E](#)), almost all simulated diffusion dynamics trajectories are attributed to the simpler diffusion model rather than the more complex momentum model ([Figure 3C](#), top row). Furthermore, SWRs following momentum dynamics are more likely misclassified as diffusion (25.5%) than the reverse (6.0%). Thus, some of the SWRs we have identified as diffusion without momentum in our data might in fact feature momentum. Overall, this might have led to underestimating the fraction of SWRs featuring momentum dynamics.

Last, we asked how much of the variance in spiking activity our best-fitting model was able to capture. We found the best-fitting model for each SWR to explain a substantial fraction of this variance (deviance explained mean \pm SD across all SWRs = 0.234 ± 0.070 ; [Figure 3D](#)), significantly more than the random model (mean \pm SD = 0.099 ± 0.050 greater than random model; paired t test, $t[2,883] = 339.9$, two-sided $p < 1 \times 10^{-6}$). Considering that 100% of deviance explained would require perfect knowledge of place cell tuning and encoded trajectory, rather than estimating them from noisy data, the observed fraction of variance is well within the expected range and in line with the literature ([Driscoll et al., 2017](#); [Minderer et al., 2019](#)). Furthermore, perturbing the data by shuffling neuron identities or randomly shifting the place fields caused the deviance explained to drop below that of the random model ([Figure S1J](#)). Thus, the identified latent dynamics structured the observed activity in a meaningful way. Taken together, our findings reveal that almost all SWRs contain dynamics with continuous spatial structure and consequently should be included in the analysis and interpretation of the computational role of replay events.

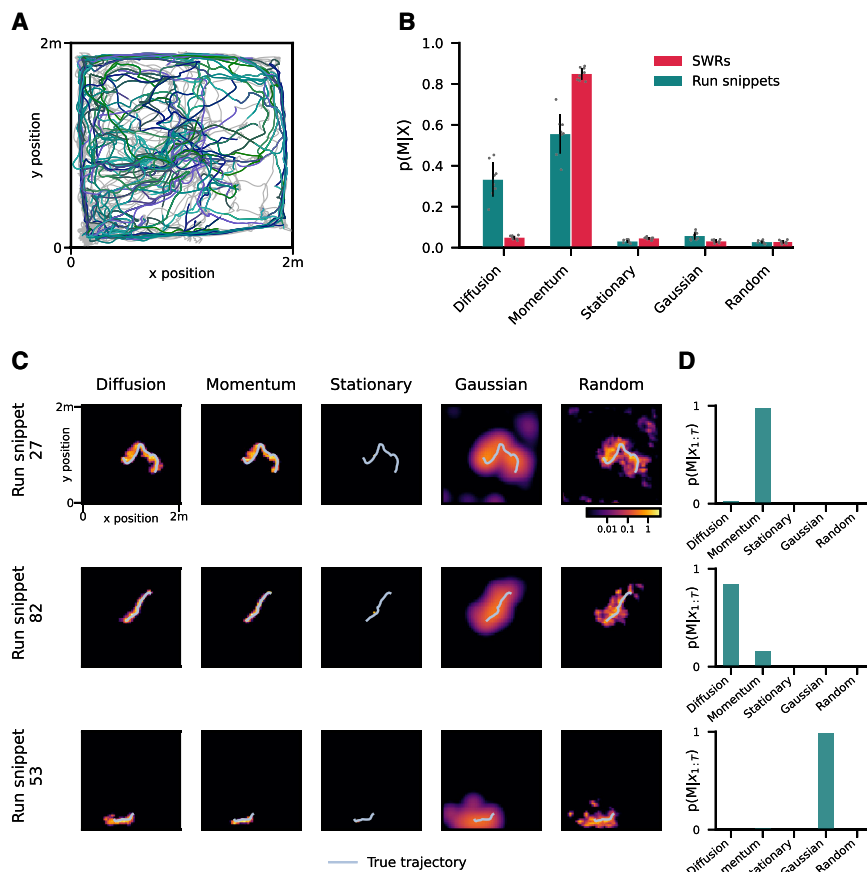


Figure 4. Place cell activity during behavior is also best described by momentum dynamics

(A) Run snippets selected for model comparison analysis for an example session. The gray trace indicates the animal's behavior throughout the entire example session. Each colored trace represents a selected run snippet (see STAR Methods; Figure S3).

(B) Inferred distribution of dynamics models underlying the generation of neural data during behavioral run snippets (teal; exceedance probability of momentum/diffusion models = $0.85 \pm 0.26/0.15 \pm 0.26$; mean \pm SD across all sessions). The same distribution is shown for SWRs (pink; same as Figure 3A) for comparison.

(C) Examples of individual run snippets. Heatmaps as in Figure 2A, with true trajectory overlaid (light blue line).

(D) The relative likelihood of each model to generate the recorded spikes within the example run snippets shown in a.

See also Figure S3.

Neural activity during SWRs resembles neural activity during behavior

If replay events are indeed involved in processing past, or planning for future, behavior, we speculate the position dynamics they encode to mimic those of real behavior. To test this, we compared the dynamics inferred from these replay events to dynamics inferred from place cell data of moving animals. We did so by applying our analysis to randomly selected snippets of place cell activity during periods of movement (Figure 4A). For a fair comparison with replay events, we matched the distribution of distances traversed in these run snippets to the distribution of distances traversed within individual SWRs (see STAR Methods; Figure S3).

Applying our random-effects model comparison to place cell activity during movement yields a similar distribution across dynamics models as for SWRs (Figure 4B); most snippets appeared to feature momentum dynamics, some diffusion dynamics, and few non-trajectory dynamics (mean \pm SD: diffusion, $33.2\% \pm 8.3\%$; momentum, $55.5\% \pm 9.6\%$; stationary, $3.0\% \pm 0.8\%$; stationary Gaussian, $5.7\% \pm 1.8\%$; random, $2.7\% \pm 0.7\%$; example run snippets in Figures 4C and 4D). The larger time bin size underlying the run snippet analysis made trajectories with momentum appear more similar to diffusion, resulting in an additional boost to diffusion dynamics compared with SWRs. Nonetheless, the overall qualitative match to the dynamics model distribution recov-

ered from SWRs, in particular for trajectory versus non-trajectory models, indicates that the statistical structure of spiking activity during SWRs is similar to that during real movement. Given that the rats' actual movement is known to feature momentum (Stella et al., 2019), it might appear surprising that the place cell activity associated with any of the run snippets appeared to feature non-momentum dynamics. However, it only confirms what we have already shown in simulations (Figure 3C): the ambiguity inherent in limited and noisy spiking data causes our model comparison to mistake trajectories with momentum as only featuring diffusion and rarely even as non-trajectories. Thus, it is also bound to underestimate the proportion of SWRs with momentum dynamics, in particular for short SWRs with few spikes (Figure S1C). As a consequence, we expect an even larger majority of SWRs than estimated to contain spatial trajectories with momentum.

Despite the good match in recovered model distributions across SWRs and run snippets, we observed quantitative differences in the underlying neural activity. SWRs are associated with bursts in population activity and thus feature more spikes per second across neurons than in similar time periods while the animal is moving (Nádasy et al., 1999; Figure S3). Even after lengthening movement snippets to match the SWRs spike counts, the trajectories decoded from SWRs traverse further distances than those decoded during movement (Figure S3). For the aforementioned model comparison, we further increased the movement snippet durations to achieve matching trajectory distance distributions. Overall, SWRs encoded trajectories with more spikes within the same time period and traversed larger distances with the same number of spikes compared with place cell activity during movement.

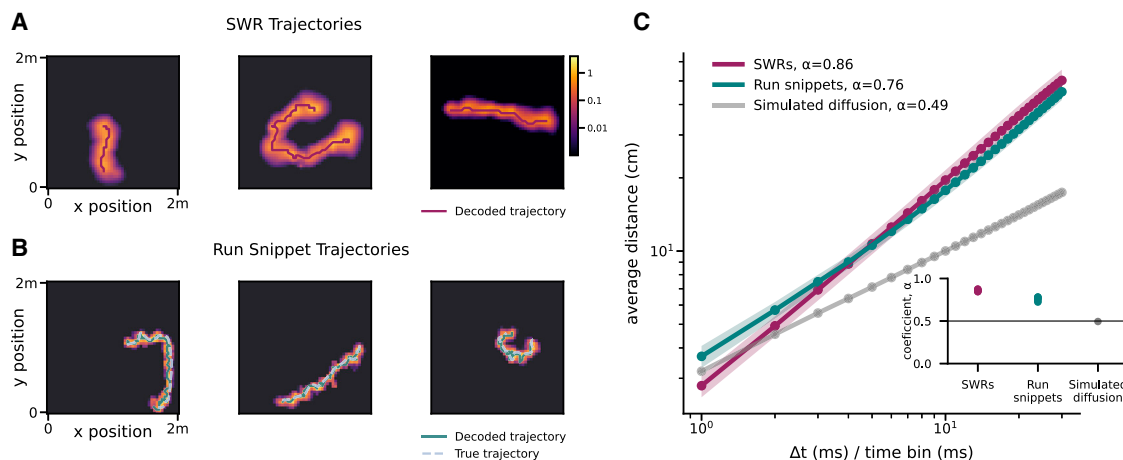


Figure 5. Trajectories decoded from awake SWRs and behavior are not consistent with Brownian motion

(A) Example most likely trajectories (purple line) decoded from SWRs (see STAR Methods), overlaid on a heatmap of the decoded positions under the diffusion dynamics model, summed over time (see Figure 2A for details).

(B) Example most likely trajectories (solid teal line) decoded from neural activity of run snippets, overlaid on a heatmap of the decoded positions under the diffusion dynamics model. The decoded trajectory lines up well with the animal's behavior trajectory (light blue dotted line).

(C) Log-log plot of time window index and mean distance from starting point for trajectories decoded from SWRs (purple) and behavior (teal) (mean \pm SD across session). The average slope (α) of a linear regression fit to each session is provided in the figure legend. The linear fit's slope significantly exceeded 0.5 for SWRs and run snippets for all sessions but not for simulated diffusion (inset; SD error bars obscured by dots). See also Figure S4.

Decoded trajectories confirm spatiotemporal momentum dynamics

Recently, Stella et al. (2019) found that the reactivated trajectories in hippocampal replay during sleep followed Brownian diffusion dynamics, akin to our diffusion dynamics model without momentum. Specifically, they identified these dynamics by a power-law relationship between time elapsed and distance traveled with an exponent of 0.5 (Rudnick and Gaspari, 2004). Applying the same analysis to place cell activity during movement, they again found a power law relationship, but this time with an exponent greater than 0.5, inconsistent with Brownian diffusion. Hence, replay events during sleep do not seem to follow dynamics resembling natural movement. This stands in contrast to our finding that awake replay follows dynamics with momentum which resemble natural movement and predict exponents >0.5 (see Supplemental information), pointing to a potential difference in the mechanism generating awake and sleep replay.

To ensure that this discrepancy did not arise from the difference in applied methodology, we replicated our analysis using the method of Stella et al. (2019). First, for each SWR we inferred the most likely trajectory by the Viterbi algorithm. This approach takes temporal continuity of latent positions into account and is thus more robust to decoding noise than the traditional approach of concatenating the most likely position across individual time bins (Figures 5A and 5B; STAR Methods). To avoid confounding our analysis with non-trajectory SWRs, we considered only SWRs that are best described by one of the trajectory models (2,366 SWRs, 82.1% of the total dataset; Figure S1B). Furthermore, we applied the Viterbi algorithm to trajectories inferred using diffusion rather than momentum dynamics, to avoid biasing the inference toward momentum dynamics. Separately, we

applied the same approach to infer trajectories from the neural data during movement, using the same run snippets we used for our model comparison analysis further above. The most likely position sequence aligned well with the marginal position estimates for SWRs (Figure 5A) and with both the rat's behavioral trajectory and the marginal position estimates for run snippets (Figure 5B; see Figure S4 for run snippet decoding accuracy).

Plotting the distance traveled within individual, inferred trajectories over elapsed time (Figure 5C), we observed a power-law relationship between these quantities with an exponent significantly exceeding 0.5 for all sessions for both the trajectories decoded from run snippets as well as from SWRs (bootstrap, both one-sided $p < 1 \times 10^{-6}$; Figure 5C, inset). This above-0.5 exponent suggests trajectory dynamics with momentum, in line with our previous model comparison. For verification, the same coefficient computed from simulated trajectories with diffusion dynamics resulted in an exponent that was not significantly larger than 0.5 (bootstrap, one-sided $p = 0.71$). This confirms that, in contrast to replay events during sleep, awake replay events appear to encode trajectory dynamics with momentum, which hints at different mechanisms underlying replay events during sleep and wakefulness.

Sub-selecting SWRs by heuristics biases the analysis of encoded trajectories

Finally, we asked if sub-selecting SWRs according to the traditional classification criterion for replay events introduced biases in analyzing how replayed trajectories perform computations supporting navigational planning. In the analyzed dataset, the animals alternated between two trial types: (1) foraging for a food reward hidden in one of the 36 food wells at random (the "goal" well) and (2) returning to collect a food reward in a known

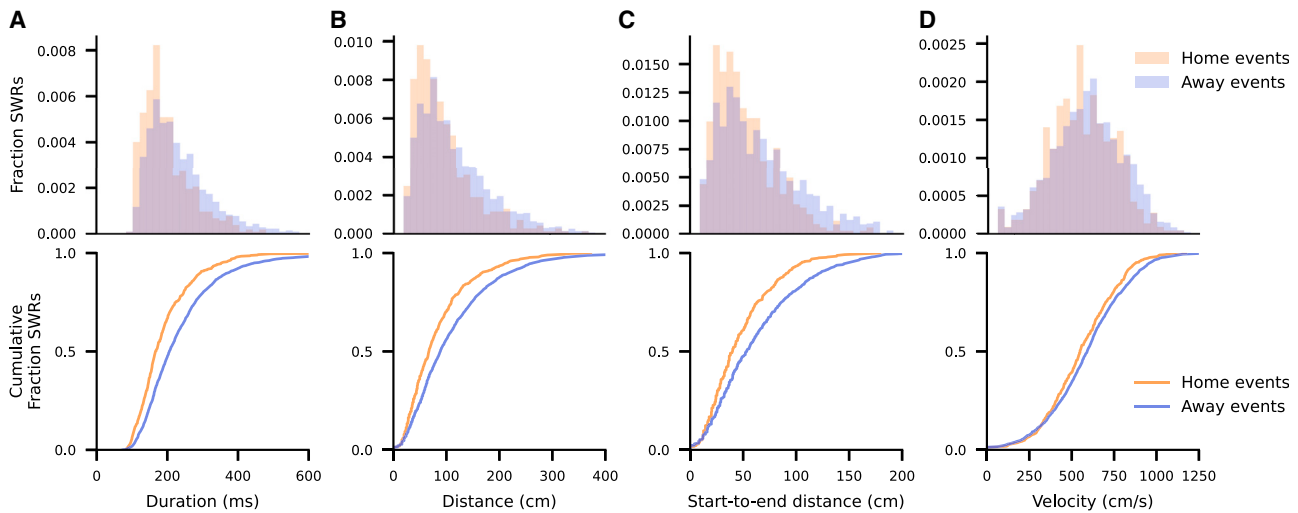


Figure 6. Decoded trajectories of home events tend to be shorter and slower than of away events

(A–D) Histogram (top row) and cumulative fraction (bottom row) of replay trajectory descriptive statistics split by home and away events depicting (A) the duration, (B) the total distance, (C) the start-to-end distance, and (D) the velocity, defined as total distance divided by duration.

See also [Figure S5](#).

“home” well that was consistent across trials within a session (see [Figure 1A](#) for well placements). This task structure allowed us to compare replay trajectories between these two types of goal-directed actions: returning to the “home” well, whose location is known, and finding the “goal” well, whose location is not. Following [Pfeiffer and Foster \(2013\)](#), we group SWRs into “home events” and “away events” on the basis of the current location of the animal, corresponding to being either at the home well or elsewhere in the environment, respectively. Although [Pfeiffer and Foster \(2013\)](#) used HSEs (STAR Methods) to identify events of interest, they switched to the use of SWRs in [Pfeiffer and Foster \(2015\)](#). For consistency, we also used SWRs, but in contrast to [Pfeiffer and Foster \(2015\)](#), who focused on the 23.7% of SWRs classified as replay events by the traditional method, we analyzed all 82.1% of SWRs that were best fit by one of our trajectory models. Furthermore, as described above, we used the Viterbi algorithm applied to the diffusion dynamics model rather than traditional decoding approaches to decode the most likely trajectories from SWRs.

First, we compared decoded trajectories between these two trial types using a set of simple descriptive statistics, namely, the duration, the total distance traveled, the start-to-end distance, and the average velocity of each of the trajectories ([Figure 6](#)). These trajectory statistics varied considerably across trajectories, with many trajectories having shorter distances than the minimum distance used in common SWR classification criteria. Furthermore, replay trajectories at the home location had a significantly shorter duration (independent t test: $t[2,366] = 7.2$, two-sided, Bonferroni-corrected $p < 1 \times 10^{-6}$), total distance (independent t test: $t[2,366] = 6.4$, two-sided, Bonferroni-corrected $p < 1 \times 10^{-6}$), and start-to-end distance (independent t test: $t[2,366] = 7.3$, two-sided, Bonferroni-corrected $p < 1 \times 10^{-6}$) than replay trajectories elsewhere in the environment and were also slightly, but significantly, slower (independent

t test: $t[2,366] = 2.9$, two-sided, Bonferroni-corrected $p = 0.013$). These differences increased for a more restrictive definition of “away events” that considered only replay events while the animal was at the goal location ([Figure S5A](#)), even after controlling for the difference in trajectory durations ([Figure S5B](#)). However, they vanished for all but duration once we only considered the subset of trajectories extracted with the traditional method for trajectory classification (as used by [Pfeiffer and Foster \(2015\)](#); [Figure S5C](#)) and thus have eluded discovery so far. More generally, irrespective of trial type, previously classified trajectories were longer in duration, distance traveled, and start-to-end distance and faster than the trajectories analyzed here ([Figure S2B](#)). Last, the differences between trial types were not present in the rodents’ movement (independent t test, two-sided, Bonferroni-corrected p values: duration, $t[2,122] = 2.59$, $p = 0.029$; total distance, $t[2,122] = 2.27$, $p = 0.070$; start-to-end distance, $t[2,122] = 0.38$, $p = 1.00$; velocity, $t[2,122] = 0.14$, $p = 1.00$), suggesting that they do not derive from the rodents’ behavior.

Second, we revisited the main finding of [Pfeiffer and Foster \(2013\)](#) that replay events are predictive of the future path taken by the animal. Specifically, they found that replayed trajectories were somewhat predictive of future but not past paths for “home events,” that is, when the animal was at the home well. They were predictive of both future and past paths for “away events,” that is, when the animal was elsewhere in the area, but more strongly so for future paths. We replicated these effects for future paths when considering only the subset of SWRs classified as replay events in [Pfeiffer and Foster \(2015\)](#), while decoding trajectories using our probabilistic model ([Figure 7B](#), green lines). Once we included all SWRs deemed as encoding trajectories, the effects remained for “away events” but vanished for “home events” ([Figure 7B](#), purple lines), and similar results were observed for the more restrictive definition of “away

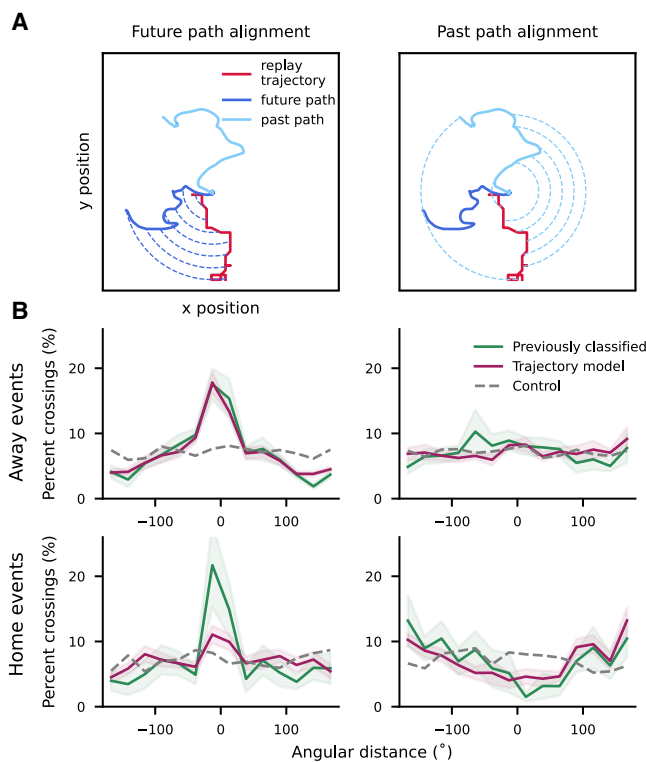


Figure 7. Away events are more biased to the future path than home events

(A) Alignment between replayed trajectories and the animal's path was quantified as the angular distance between the decoded replay trajectory and the future (left) or past path (right) at a series of concentric circles expanding outward from the current location of the animal (Pfeiffer and Foster, 2013; see STAR Methods). Using a criterion that depends on distance but not velocity supports comparing replay and behavioral trajectories that evolve at different speeds.

(B) The plots show histograms of angular alignments, computed as described in (A). A peak at zero indicates that the replayed trajectories are predictive of the animal's taken path. The different plots show these alignments for "away events" (top) and "home events" (bottom), and for future (left) and past (right) paths, when including only previously classified SWRs (green; mean \pm SD within each alignment bin) or all the SWRs we classified as encoding trajectories (purple; mean \pm SD within each alignment bin). The dashed gray lines indicate alignment from shuffled events (see STAR Methods). See also Figure S6.

events" (Figure S6). We did not identify any predictivity for past paths.

Although a detailed computational account of these findings is beyond the scope of our work, these results suggest that, preceding directed movement, SWRs simulate longer trajectories that more closely resemble the full future path of the animal. In contrast, preceding random foraging, they simulate shorter possible paths that do not necessarily predict immediate future behavior.

DISCUSSION

Applying a probabilistic method to classify SWRs, we found that a large majority of SWRs during exploration of a 2D environment

contain temporally continuous spatial information. Furthermore, the encoded trajectories evolve under dynamics with momentum, similar to real movement through the environment. We showed this by using Bayesian model comparison to infer which of a set of five state-space models, each making different explicit assumptions about the encoded position dynamics, is able to explain each SWR's spike pattern. Our result stands in contrast to previous, heuristic approaches, which identified only the most salient trajectories but left the role of the remaining SWRs unresolved (Foster and Wilson, 2006; Karlsson and Frank, 2009; Pfeiffer and Foster, 2015; Tingley and Peyrache, 2020). Observing that the replayed trajectories featured momentum further contrasts with previous work showing that trajectories replayed during sleep followed Brownian motion without momentum (Stella et al., 2019). Thus, awake replay might result from different underlying mechanisms and feature a different purpose. As discarding the majority of SWRs might have led to biases in follow-up analyses, we replicated the analyses of Pfeiffer and Foster (2013) on the full set of SWRs and identified small, but significant, differences, as well as previously underappreciated heterogeneity of the embedded trajectories. More generally, our method promises a less biased analysis of hippocampal replay that we expect to aid future research on identifying the computational role of replay in planning and learning.

Although our model comparison classifies a small fraction of SWRs as non-trajectories, even these might correspond to replay events. They might, for example, constitute replay of close to stationary trajectories (Denovellis et al., 2021; Yu et al., 2017), which are in turn best captured by our stationary (non-trajectory) model. Furthermore, some SWRs might replay trajectories in a different environment, like the rodent's home cage (Karlsson and Frank, 2009). This implies a different mapping between place cell activity and encoded location, in which case the trajectory should appear random to the model (Leutgeb et al., 2004). The low prevalence of SWRs being classified as random suggests that such replays are rare or non-existent in the analyzed data.

Observing that the trajectories underlying awake replay feature momentum leads to questions about the mechanisms that generate them. During movement, encoded trajectories are expected to feature momentum, as they mirror the animal's natural movement (Figures 4B and 5C). Absent such movement, the origin of momentum is less clear. Momentum could arise in at least two ways, both involving some notion of the animal's velocity. Representing the immediately preceding location in addition to the current one would allow the next location to depend on both, such that the distance between current and next location to relate to that between preceding and current one, effectively implementing momentum. This is, in fact, how we implemented the momentum model in our model comparison approach (see STAR Methods; Supplemental information). Alternatively, awake replays could engage consistent activity in areas representing the animal's velocity (Kropff et al., 2015) and heading direction (Taube, 2007), similar to when the animal is moving through the environment. These alternatives lead to different predictions about neural activity during replay events that can be distinguished experimentally. Furthermore, it would distinguish them from replay events during sleep, during which momentum is

not observed, and might elucidate how these two types of replay events differ in terms of function.

Although our trajectory decoding approach did not explicitly decode velocities, it is worth considering how a consistent velocity representation during replayed trajectories might aid position decoding by downstream brain regions. For example, it could augment uncertain decoding of the currently replayed position by predicting this position from past position estimates and decoded velocities. In fact, the same principle has been used for decoding position during behavior (Zhang et al., 1998). The nervous system, in contrast, might not benefit from adding velocity to its readout for follow-on computations, as it has access to all position-encoding neurons. This might yield highly precise position information that added velocity information might not improve upon.

Decoding replayed trajectories with our model revealed subtle, task-specific details in the trajectories' dynamics whose highly heterogeneous structure might have been underappreciated by past heuristic-driven analyses. In particular, we found replay events preceding random foraging to be shorter and slower than events preceding directed movement, a detail that is missed when only considering the subset of SWRs that meet the heuristic classification criteria. Additionally, we can interpret our trajectories in light of a recent theory by Mattar and Daw (2018), which suggests that replay trajectories reflect previous experience replayed to improve action selection, prioritized to balance propagating new information about reward with focusing on the most imminent locations in the environment. Thus, for the analyzed dataset, replay trajectories driven by consideration of imminent choices should predict future paths for "away events," when the future goal location is known, but not for "home events," when the future location is unknown. Although in conflict with previous analysis of Pfeiffer and Foster (2013) that focused on the most salient replayed trajectories, we confirm these predictions once we include all events we identified as trajectories (Figure 7B). However, we would furthermore expect replay trajectories at the home well to predict past paths, driven by propagation of reward, which we do not see. Although we consider an interpretation of these details beyond the scope of our work, they are examples of the kinds of characteristics of replay events that would be useful to consider in further studies assessing the computational role of replay, but that might be missed without a systematic and unbiased characterization of their spatiotemporal structure.

Treating the identification of latent position dynamics as a probabilistic inference problem has several benefits. First, it allows use of the full posterior estimate of the position across time rather than relying on point estimates of the most likely position, which, because of noisy neural data, will be noisy and thus unreliable. Second, the models capture uncertainty in the encoded spatial pattern, which in turn is reflected in the relative certainty between possible underlying dynamics per SWR. Last, it provides a consistent and unbiased approach to focusing on the aspects of the dynamics of interest, by formulating a targeted set of models that differ in their explicit assumptions about the underlying dynamics, like diffusion with or without momentum.

Our approach relied on spike-sorted place field data, which might discard a significant number of spike-like voltage traces,

whose inclusion might boost our model comparison's power. Recent work side-stepped spike sorting and decoded place cell data from raw voltage traces instead (Deng et al., 2015; Kloosterman et al., 2014), thereby improving decoding performance (Kay et al., 2020). Given the large number of recorded neurons in our data, and the already fairly conclusive model comparison results, we do not expect this alternative approach to significantly improve our analysis. However, it might benefit similar analyses applied to recordings from smaller populations, as was the case in aforementioned work.

One of the limitations of our approach, shared to our knowledge by all work that decodes location from place cell activity, is that we assume neural activity of these cells for each location to be statistically independent of each other. Noise correlations modulate information and for the considered population sizes could either boost or decrease it (Averbeck and Costa, 2017; Kohn et al., 2016). Although data sparsity makes it impossible to precisely quantify such correlations, we do not expect that taking them into account would qualitatively change our results. Additionally, we have assumed the place cell activity's rate during SWRs to scale up by a constant factor, shared by all neurons, compared with activity during movement (Figure S3A). Allowing this factor to vary across SWRs and neurons, which effectively leads to over-dispersed SWR spike counts as suggested by Chandrasekaran et al. (2018), neither improved our model fits nor changed the outcome of our model comparison (Figures S1K–S1M). Relatedly, our model comparison approach is restricted to the five dynamics models we have considered, and trajectories that match neither dynamics model would erroneously be attributed to one of them. Given the high flexibility with which we formulated the dynamics models, we do not expect this to have confounded our results. Nonetheless, an interesting avenue for future work could be to model the latent location sequence as a hierarchical Gaussian process, as in Wu et al. (2017), in which case structural assumptions are flexibly captured by the covariance matrix governing location evolution. This would allow even greater flexibility in defining possible transition structures of interest. More generally, Bayesian model comparison can be sensitive to slight model misspecification (Chandrasekaran et al., 2018), which we hope to have avoided by our exploration of different variants of both place cell activity likelihoods and model dynamics. Last, we have focused on identifying the macrodynamics of replay trajectories, which does not capture previously identified fine-scale dynamics within replay events in which the position "jumps" between successive positions (Pfeiffer and Foster, 2015).

Our work identifies the dominant function of SWRs as simulating trajectories and lays a principled foundation for future work examining the computational role of the replayed trajectories. Our finding that their spatiotemporal dynamics feature momentum informs the possible computational roles they could play. Specifically, this argues against awake replay events being generated by local Markovian dynamics based on location only and instead suggests that they emerge from network dynamics that incorporate a notion of velocity. Overall, the use of state-space models provides a principled method for analyzing the spatiotemporal structure contained within sharp-wave ripples that future work on their computational function and underlying mechanisms can build upon.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
 - Animal model
- METHOD DETAILS
 - Data acquisition and behavior
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Estimating place fields
 - SWR and HSE detection
 - Extracting population bursts from SWRs
 - Description of state-space models used to characterize SWRs
 - Poisson spike generation model
 - Over-dispersed spike generation model
 - Dynamics models
 - Relationship between dynamics models
 - Marginalizing over latent position sequences and model parameters
 - Computing fraction deviance explained
 - Best-fitting models, and fixed and random effects model comparison of dynamics models
 - Model recovery from simulated data
 - Extracting maximum likelihood trajectories
 - Traditional method for classifying SWRs
 - Analysis of place cell activity during movement
 - Diffusion coefficient analysis
 - Splitting SWRs into “home” and “away” events
 - Analysis of correlation between replay trajectories and behavior
 - Figure details

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2021.11.014>.

ACKNOWLEDGMENTS

We thank Brad Pfeiffer and David Foster for sharing their data with us. Furthermore, we would like to thank Matt Wilson, Sam Gershman, and members of the Drugowitsch lab for feedback on the work and Anna Kutschireiter, Christopher Harvey, Johannes Bill, and John Vastola for comments on an early draft of this manuscript. The work was supported by a James S. McDonnell Foundation Scholar Award for Understanding Human Cognition (J.D., grant 220020462) and a National Defense Science and Engineering Graduate Fellowship (E.L.K.).

AUTHOR CONTRIBUTIONS

E.L.K. and J.D. conceptualized and designed the research. E.L.K. performed the analysis. E.L.K. and J.D. interpreted the results. E.L.K. wrote the initial draft of the manuscript. E.L.K. and J.D. revised the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 26, 2021

Revised: October 6, 2021

Accepted: November 12, 2021

Published: December 3, 2021

REFERENCES

- Averbeck, B.B., and Costa, V.D. (2017). Motivational neural circuits underlying reinforcement learning. *Nat. Neurosci.* *20*, 505–512.
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning* (New York: Springer).
- Buzsáki, G. (2015). Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. *Hippocampus* *25*, 1073–1188.
- Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* *14*, 147–153.
- Chandrasekaran, C., Soldado-Magraner, J., Peixoto, D., Newsome, W.T., Shenoy, K.V., and Sahani, M. (2018). Brittleness in model selection analysis of single neuron firing rates. *bioRxiv*. <https://doi.org/10.1101/430710>.
- Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. *Neuron* *63*, 497–507.
- Deng, X., Liu, D.F., Kay, K., Frank, L.M., and Edén, U.T. (2015). Clusterless decoding of position from multiunit activity using a marked point process filter. *Neural Comput.* *27*, 1438–1460.
- Denovellis, E.L., Gillespie, A.K., Coulter, M.E., Sosa, M., Chung, J.E., Edén, U.T., and Frank, L.M. (2021). Hippocampal replay of experience at real-world speeds. *eLife* *10*, e64505.
- Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* *10*, 1241–1242.
- Driscoll, L.N., Pettit, N.L., Minderer, M., Chettih, S.N., and Harvey, C.D. (2017). Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* *170*, 986–999.e16.
- Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* *440*, 680–683.
- Gillespie, A.K., Astudillo Maya, D.A., Denovellis, E.L., Liu, D.F., Kastner, D.B., Coulter, M.E., Roumis, D.K., Edén, U.T., and Frank, L.M. (2021). Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *bioRxiv*. <https://doi.org/10.1101/2021.03.09.434621>.
- Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science* *336*, 1454–1458.
- Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* *27*, 12176–12189.
- Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. *Nat. Rev. Neurosci.* *19*, 744–757.
- Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* *12*, 913–918.
- Kay, K., Chung, J.E., Sosa, M., Schor, J.S., Karlsson, M.P., Larkin, M.C., Liu, D.F., and Frank, L.M. (2020). Constant sub-second cycling between representations of possible futures in the hippocampus. *Cell* *180*, 552–567.e25.
- Kloosterman, F., Layton, S.P., Chen, Z., and Wilson, M.A. (2014). Bayesian decoding using unsorted spikes in the rat hippocampus. *J. Neurophysiol.* *111*, 217–227.
- Kohn, A., Coen-Cagli, R., Kanitscheider, I., and Pouget, A. (2016). Correlations and neuronal population information. *Annu. Rev. Neurosci.* *39*, 237–256.

- Kropff, E., Carmichael, J.E., Moser, M.-B., and Moser, E.I. (2015). Speed cells in the medial entorhinal cortex. *Nature* *523*, 419–424.
- Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.B., and Moser, E.I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* *305*, 1295–1298.
- Linderman, S.W., Johnson, M.J., Wilson, M.A., and Chen, Z. (2016). A Bayesian nonparametric approach for uncovering rat hippocampal population codes during spatial navigation. *J. Neurosci. Methods* *263*, 36–47.
- Maboudi, K., Ackermann, E., de Jong, L.W., Pfeiffer, B.E., Foster, D., Diba, K., and Kemere, C. (2018). Uncovering temporal structure in hippocampal output patterns. *eLife* *7*, e34467.
- MacKay, D.J.C. (1995). *Information Theory, Inference, and Learning Algorithms* (Cambridge, UK: Cambridge University Press).
- Mattar, M.G., and Daw, N.D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nat. Neurosci.* *21*, 1609–1617.
- Minderer, M., Brown, K.D., and Harvey, C.D. (2019). The spatial structure of neural encoding in mouse posterior cortex during navigation. *Neuron* *102*, 232–248.e11.
- Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* *19*, 9497–9507.
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* *34*, 171–175.
- Penny, W.D., Stephan, K.E., Daunizeau, J., Rosa, M.J., Friston, K.J., Schofield, T.M., and Leff, A.P. (2010). Comparing families of dynamic causal models. *PLoS Comput. Biol.* *6*, e1000709.
- Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* *497*, 74–79.
- Pfeiffer, B.E., and Foster, D.J. (2015). Place cells. Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science* *349*, 180–183.
- Rudnick, J., and Gaspari, G. (2004). *Elements of the Random Walk: An Introduction for Advanced Students and Researchers* (Cambridge, UK: Cambridge University Press).
- Shin, J.D., Tang, W., and Jadhav, S.P. (2019). Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. *Neuron* *104*, 1110–1125.e7.
- Singer, A.C., Carr, M.F., Karlsson, M.P., and Frank, L.M. (2013). Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* *77*, 1163–1173.
- Stella, F., Baracska, P., O'Neill, J., and Csicsvari, J. (2019). Hippocampal reactivation of random trajectories resembling Brownian diffusion. *Neuron* *102*, 450–461.e7.
- Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., and Friston, K.J. (2009). Bayesian model selection for group studies. *Neuroimage* *46*, 1004–1017.
- Sutton, R.S., and Barto, A.G. (1998). *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press).
- Taube, J.S. (2007). The head direction signal: origins and sensory-motor integration. *Annu. Rev. Neurosci.* *30*, 181–207.
- Tingley, D., and Peyrache, A. (2020). On the methods for reactivation and replay analysis. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *375*, 20190231.
- Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* *265*, 676–679.
- Wu, A., Roy, N.A., Keeley, S., and Pillow, J.W. (2017). Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Adv. Neural Inf. Process. Syst.* *30*, 3496–3505.
- Xu, H., Baracska, P., O'Neill, J., and Csicsvari, J. (2019). Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. *Neuron* *101*, 119–132.e4.
- Yu, J.Y., Kay, K., Liu, D.F., Grossrubatscher, I., Loback, A., Sosa, M., Chung, J.E., Karlsson, M.P., Larkin, M.C., and Frank, L.M. (2017). Distinct hippocampal-cortical memory representations for experiences associated with movement versus immobility. *eLife* *6*, e27621.
- Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* *79*, 1017–1044.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms and strains		
Long Evans Rats	Pfeiffer and Foster, 2013, 2015	N/A
Software and algorithms		
Python	https://www.python.org	Python 3.8.2
NumPy	https://numpy.org	Numpy 1.18.2
SciPy	https://scipy.org	SciPy 1.4.1
Pandas	https://pandas.pydata.org	Pandas 1.0.3
PyTorch	https://pytorch.org	PyTorch 1.4.0
Custom code used to analyze the data and generate the figures	https://zenodo.org/record/5651913	https://doi.org/10.5281/zenodo.5651913

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Jan Drugowitsch (jan_drugowitsch@hms.harvard.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- The paper analyzes data from [Pfeiffer and Foster \(2013, 2015\)](#). Please contact Brad Pfeiffer and David Foster to obtain the analyzed datasets.
- The code used to analyze the data and generate the figures has been deposited at Zenodo and is publicly available as of the date of publication. The DOI is listed in the key resources table.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animal model

The dataset used in this study has been described in detail previously in [Pfeiffer and Foster \(2013, 2015\)](#). All procedures were approved by the Johns Hopkins University Animal Care and Use Committee and followed US National Institutes of Health animal use guidelines. A total of four wild-type male Long Evans rats (10-20 weeks old, 450-550 g) were used for this study. Animals were housed in a standard, non-inverted, 12-hour light cycle, handled daily, and food-restricted to 85%–90% of their free-feeding weight. Behavioral training and recording were performed in the afternoon to early evening.

METHOD DETAILS

Data acquisition and behavior

As described in [Pfeiffer and Foster \(2013, 2015\)](#), animals were trained on a foraging task in a 2 m x 2 m open field environment, in which the animal foraged for food hidden in one of 36 wells, spaced evenly in a 6 well x 6 well grid. Neural recordings were obtained via a microdrive array containing 40 gold-plated tetrodes implanted in CA1 of dorsal hippocampus. Data was obtained from each rat for two days (a total of 8 sessions across rats). Individual units were identified by [Pfeiffer and Foster \(2015\)](#) by manual clustering based on the spike waveform peak amplitudes obtained from custom software (xclust2, Matt A. Wilson), and inhibitory units were excluded on the basis of spike width and mean firing rate.

The animal's position was determined by two distinctly colored, head-mounted LEDs, recorded at 60Hz by an overhead video system, and down-sampled to 30Hz before analysis. The animal's speed was calculated as the distance between successive recorded

positions over the recording resolution rate, and the data was split in to “moving” and “not moving” periods by a threshold of 5 cm/s on the running speed.

During data collection, animals engaged in a foraging task that consisted of two alternating trial types: searching for food in a well that was in a different location across trials (the “goal” well), or a well that was consistent location across trials (“home” well). The goal well was in a random location on each trial, excluding the home well and the goal well from the previous trial. The home well was in a consistent location within each session, but changed across sessions.

QUANTIFICATION AND STATISTICAL ANALYSIS

Estimating place fields

Place fields were fit using spiking data from periods in which the rat was considered moving (see above). The animal’s position, $\mathbf{z} = [z_x, z_y]$, was binned into a 50×50 grid (4 cm x 4 cm bins), leading to 2,500 unique positions within the arena. For each cell, i , the corresponding place field $f_i(\mathbf{z})$ was defined as the maximum a-posteriori estimate assuming Poisson spiking, $\text{Pois}(x_t | f_i(\mathbf{z}_t))$, and a Gamma prior on the firing rate, $\text{Gam}(f_i(\mathbf{z}_t) | \alpha, \beta)$, such that for each cell i and spatial bin, $\mathbf{z}_t = k$,

$$f_i(\mathbf{z}_t = k) = \frac{1}{\beta + T_{k,run} \delta t} \left(\sum_{t=1}^{T_{k,run}} x_{i,k,t} + \alpha - 1 \right) \quad (\text{Equation 1})$$

where $T_{k,run}$ is the total number of time bins spent in spatial bin k , $x_{i,k,t}$ is the spike count x emitted by cell i in spatial bin k at time bin t , and δt is the time bin size (recording resolution of 30 Hz, such that, here $\delta t = 1/30$). A weak prior was used with parameter settings $\alpha = 1.01$ and $\beta = 0.01$, such that the maximum a-priori firing rate is $(\alpha - 1)/\beta = 1$ spikes/s. After estimation, place fields $f_i(\mathbf{z}_t)$ were smoothed with a Gaussian kernel (4 cm standard deviation). A cell was considered a place cell if the peak of its tuning curve, $\max_{\mathbf{z}} f_i(\mathbf{z})$, exceeded 2 spikes/second. Only cells identified as place cells were used in further analysis.

SWR and HSE detection

For our study, we used the set of SWRs identified in [Pfeiffer and Foster \(2015\)](#) based on features of the local field potential (LFP). As described previously, a representative electrode was chosen for each tetrode, the LFP for that tetrode was band-pass filtered between 150 and 250 Hz, and the absolute value Hilbert transform of the filtered signal was smoothed with a Gaussian kernel (12.5 ms standard deviation). The average of this signal across all tetrodes was then used to identify SWRs as a local peak with an amplitude greater than 3 standard deviations above the mean, with start and end boundaries as the point when the signal crossed the mean. Only SWRs longer than 50 ms and shorter than 2 s that occurred when the rat was considered “not moving” were included for further analysis.

In addition to SWRs, we also conducted our trajectory identification analysis on periods of high neuronal synchrony, termed high-synchrony events (HSEs), which are commonly used in studies of hippocampal replay ([Davidson et al., 2009](#); [Pfeiffer and Foster, 2013](#); [Xu et al., 2019](#)). Following [Pfeiffer and Foster \(2013\)](#), a histogram of total spikes per non-overlapping 1 ms time bins was calculated and smoothed using a Gaussian kernel (20 ms standard deviation). HSEs were identified as peaks in the smoothed histogram greater than 3 standard deviations above the mean, with the starting point and end point of each HSE defined as the time point in which the smoothed histogram crossed the mean. Only time periods in which the animal was considered “not moving” were included for further analysis.

Extracting population bursts from SWRs

To ensure low firing-rate periods flanking the population activity burst of interest within an SWR did not impact our analysis, we sub-selected a time period of each SWR in which the average firing rate across all neurons was above 2 spikes/s per neuron. Specifically, for each SWR, the spiking activity was binned using a 3 ms time bin (unless stated otherwise, [Figure S1A](#)), and the “population burst” was defined as the period of time from the first upward-crossing to the last downward-crossing of the firing rate threshold within the SWR ([Figure S7A](#)). Population bursts shorter than 30 ms were excluded from the analysis (73 SWRs, or 2.5% of the 2956 total SWRs were excluded). This resulted in population bursts that were shorter than the original SWRs, but had roughly the same number of spikes ([Figures S7B and S7C](#)). Throughout all further analyses we only analyzed neural data within the extracted population burst, but referred to the extracted population bursts as SWRs, for simplicity.

Description of state-space models used to characterize SWRs

All five state-space models that we use to distinguish between the spatio-temporal dynamics of the SWRs consist of two components: (i) a spike generation model, which is common across all models, and (ii) a dynamics model, which differed across models. The spike generation model, $p(\mathbf{x}_{1:T} | \mathbf{z}_{1:T})$, describes how the latent position at each discretized time bin, $\mathbf{z}_{1:T} = \mathbf{z}_1, \dots, \mathbf{z}_T$ for a sequence of length T , is assumed to generate spike trains, $\mathbf{x}_{1:T} = \mathbf{x}_1, \dots, \mathbf{x}_T$. Here, \mathbf{z}_t denotes the latent grid position at time t , and \mathbf{x}_t is a vector of spike counts across cells for the t^{th} time bin of size δt ($\delta t = 3\text{ms}$ unless stated otherwise, see [Figure S1](#)). The dynamics model, $p(\mathbf{z}_{1:T} | M)$ differs across models, M , and describes the probability of observing a specific position sequence $\mathbf{z}_{1:T}$ for each of these models. Together, they result in the joint probability across spikes and latent positions for state-space model M to be given by

$$p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T} | M) = p(\mathbf{x}_{1:T} | \mathbf{z}_{1:T}) p(\mathbf{z}_{1:T} | M). \tag{Equation 2}$$

Here, and in much of the below, the dynamics model is implicitly also conditional on the model-specific parameters, θ_M . We make this explicit whenever required for clarity.

Poisson spike generation model

In this and the next section, we consider two spike generation models. As we show in Figure S1L, the first explains the data better, such that we focus on this model in the main text. For both models, we assume the spike generation probabilities (that is, the HMM emission probabilities) $p(\mathbf{x}_{1:T} | \mathbf{z}_{1:T})$ to factor across time, that is $p(\mathbf{x}_{1:T} | \mathbf{z}_{1:T}) = \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{z}_t)$. The first model assumes that, within each time bin, spikes are drawn from a Poisson distribution, independent across cells, with rates determined by the place fields fit during movement (see Estimating place cells), that is

$$p(\mathbf{x}_t | \mathbf{z}_t) = \prod_{i=1}^N \text{Pois}(x_{t,i} | f_i(\mathbf{z}_t) \gamma_{\text{spike}} \delta t), \tag{Equation 3}$$

where $x_{t,i}$ is the spike count of cell i in time bin index t , δt is the time bin size, and γ_{spike} is the spike count scaling factor (described below). This is equivalent to assuming that the spike sequence of each cell is drawn according to a time-discretized inhomogeneous Poisson process with instantaneous spike rate $f_i(\mathbf{z}_t) \gamma_{\text{spike}}$ that varies across positions \mathbf{z}_t , and thus time.

We accounted for the increase in firing rate within SWRs as compared to movement by calculating a spike count scaling factor, γ_{spike} (as described in Figure S3A). Average firing rate was calculated for each unit as the total number of spikes emitted over the total duration within periods that the rat was either considered “moving,” $\tilde{f}_{i,\text{movement}}$, or within population bursts, $\tilde{f}_{i,\text{popburst}}$ for each cell i . A linear regression was performed over the set of $(\tilde{f}_{i,\text{movement}}, \tilde{f}_{i,\text{popburst}})$ pairs per session, and the average slope, or average population activity scaling, across all sessions, $\gamma_{\text{spike}} = 2.9$, was used as the spike count scaling factor.

Over-dispersed spike generation model

The spike count scaling factor’s value, γ_{spike} , is uncertain, as it is estimated from noisy data. Our second spike generation model makes this uncertainty explicit, by modeling γ_{spike} are a random variable. Specifically, we assume that, within each session, γ_{spike} follows a Gamma distribution $\gamma_{\text{spike}} \sim \text{Gam}(\alpha_\lambda, \beta_\lambda)$. The distribution’s parameters, α_λ and β_λ , are, separately for each session, estimated by maximum likelihood from $\tilde{f}_{i,\text{popburst}} / \tilde{f}_{i,\text{movement}}$ measures across all neurons i recorded in this session. Furthermore, for each SWR, we assume each neuron’s spike scaling factor, $\gamma_{i,\text{spike}}$ for neuron i to be drawn independently from this distribution. With these assumptions, the spike count likelihood becomes

$$p(\mathbf{x}_t | \mathbf{z}_t) = \prod_{i=1}^N \int \text{Pois}(x_{t,i} | f_i(\mathbf{z}_t) \gamma_{i,\text{spike}} \delta t) \text{Gam}(\gamma_{i,\text{spike}} | \alpha_\lambda, \beta_\lambda) d\gamma_{i,\text{spike}} = \prod_{i=1}^N \text{NB}\left(x_{t,i} \mid \alpha_\lambda, \frac{f_i(\mathbf{z}_t) \delta t}{f_i(\mathbf{z}_t) \delta t + \beta_\lambda}\right). \tag{Equation 4}$$

Thus, the likelihood becomes a product of negative binomial distributions. For neuron i , the associated negative binomial distribution has mean $\langle x_{t,i} | \mathbf{z}_t \rangle = f_i(\mathbf{z}_t) \delta t \langle \gamma_{i,\text{spike}} \rangle$, just like the above Poisson spike generation model. However, in contrast to that first model, it features a larger variance, $\text{var}(x_{t,i} | \mathbf{z}_t) = \frac{f_i(\mathbf{z}_t) \delta t + \beta_\lambda}{\beta_\lambda} \langle x_{t,i} | \mathbf{z}_t \rangle$ (note the ≥ 1 pre-factor). Thus, taking into account our uncertainty in γ_{spike} results in an over-dispersed spike generation model

This second model is more complex than the first, such that Bayesian model comparison only prefers it if this additional complexity is warranted to explain the observed data. As shown in Figure S1L, our data does not support this additional complexity, such that we only focus on the first, Poisson spike generation model, in the main text.

Dynamics models

Different dynamics models, M , are characterized by how the latent position evolves over time,

$$p(\mathbf{z}_{1:T} | M) = p(\mathbf{z}_1 | M) \prod_{t=2}^T p(\mathbf{z}_t | \mathbf{z}_{1:t-1}, M). \tag{Equation 5}$$

We define five dynamics models: 1) *diffusion* - position evolves as a random walk, 2) *diffusion with momentum* - velocity evolves as a random walk with decay (that is, an Ornstein-Uhlenbeck process), 3) *stationary* - position remains constant, 4) *stationary Gaussian* - position is at each time point drawn from a Gaussian with constant moments, 5) *random* - position is at each time point drawn from a uniform distribution. Each of these models correspond to a first order Hidden Markov Model, except for the momentum model, which corresponds to a second order Hidden Markov Model (Figure 1D, bottom row). In continuous time, the dynamics of latent position \mathbf{z}_t in the diffusion and momentum models is described by

$$\text{Diffusion: } \frac{dz_j}{dt} = \sigma_d \eta_{j,d}(t), \quad (\text{Equation 6})$$

$$\text{Momentum: } \frac{dz_j}{dt} = v_j(t), \quad \frac{dv_j}{dt} = -\lambda_m v_j(t) + \sigma_m \eta_{j,m}(t), \quad (\text{Equation 7})$$

where the processes are independent across spatial dimension $j \in \{x, y\}$, and $\eta_{\cdot, \cdot}(t)$ are Gaussian white noise processes. Here, σ_d is the diffusion coefficient of the diffusion model, and λ_m and σ_m are the decay and diffusion coefficient, respectively, of the momentum model. As both models describe continuous spatial trajectories, we refer to them collectively as *trajectory* models.

We implement these models in discrete time, where they are given by

$$\text{Diffusion: } p(\mathbf{z}_t | \mathbf{z}_{t-1}) = N(\mathbf{z}_t | \mathbf{z}_{t-1}, \sigma_d^2 \delta t \mathbf{I}), \quad (\text{Equation 8})$$

$$\text{Momentum: } p(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{z}_{t-2}) = N\left(\mathbf{z}_t | (1 + e^{-\lambda_m \delta t}) \mathbf{z}_{t-1} - e^{-\lambda_m \delta t} \mathbf{z}_{t-2}, \frac{\sigma_m^2 \delta t^2}{2\lambda_m} (1 - e^{-2\lambda_m \delta t}) \mathbf{I}\right), \quad (\text{Equation 9})$$

$$p(\mathbf{z}_2 | \mathbf{z}_1) = N(\mathbf{z}_2 | \mathbf{z}_1, \sigma_{m0}^2 \delta t \mathbf{I}), \quad (\text{Equation 10})$$

In the above, all normal distributions $N(\cdot)$ are bivariate normal distributions over both spatial dimensions with diagonal covariance matrices, and discretized and appropriately normalized over the spatial grid. The detailed derivation for the momentum model is provided in [Supplemental information](#). It results in a second-order Markov model, where the first step, $p(\mathbf{z}_2 | \mathbf{z}_1)$ is modeled separately as a diffusion with variance $\sigma_{m0}^2 \delta t$, which effectively implements a prior over the initial velocity, $v_j(0)$. The diffusion model describes trajectories with a diffusion coefficient of $\alpha = 0.5$ (see below and main text). Over the short time-scales of SWRs, the momentum model describes super-diffusive trajectories with $\alpha > 0.5$ (see [Supplemental information](#) for derivation).

The non-trajectory models are only specified in discrete time, and are given by

$$\text{Stationary: } p(\mathbf{z}_t | \mathbf{z}_{1:t-1}) = p(\mathbf{z}_t) = \delta_{\mathbf{z}_t, \mathbf{z}_s}, \quad (\text{Equation 11})$$

$$\text{Gaussian: } p(\mathbf{z}_t | \mathbf{z}_{1:t-1}) = p(\mathbf{z}_t) = N(\mathbf{z}_t | \boldsymbol{\mu}, \sigma_g^2 \mathbf{I}), \quad (\text{Equation 12})$$

$$\text{Random: } p(\mathbf{z}_t | \mathbf{z}_{1:t-1}) = p(\mathbf{z}_t) = U(\mathbf{z}_t), \quad (\text{Equation 13})$$

for all $t = 1, \dots, T$, and where $\delta_{\mathbf{z}_t, \mathbf{z}_s}$ is the Kronecker delta function that is one if $\mathbf{z}_t = \mathbf{z}_s$, and zero otherwise. The normal distribution $N(\mathbf{z}_t | \boldsymbol{\mu}, \sigma_g^2 \mathbf{I})$ is, as before, a bivariate normal distribution over both spatial dimensions with a diagonal covariance matrix, and $U(\mathbf{z}_t) = 1/K$ is the uniform distribution across all $K = 2,500$ spatial bins. The non-trajectory models are specified by parameters \mathbf{z}_s for the stationary location, and $\boldsymbol{\mu}$ and σ_g^2 for the Gaussian mean and variance, respectively.

The model parameters σ_d , σ_m , λ_m , σ_g , $\boldsymbol{\mu}$, and \mathbf{z}_s were fit over a grid and marginalized out for model comparison (see below section for details). σ_{m0} in the momentum model, which is the standard deviation of the diffusion process on the first time step, $p(\mathbf{z}_2 | \mathbf{z}_1)$, was set a single pre-determined value of 10 m/s, representing a wide prior on the initial velocity of the trajectory. For all models a uniform prior was used for the initial position, $p(\mathbf{z}_1 | M) = U(\mathbf{z}_1) = 1/K$, where $K = 2,500$ is the number of spatial bins.

Relationship between dynamics models

Both trajectory and non-trajectory model pairs generate very similar dynamics predictions in certain parameter limits that we discuss here. If two dynamics models generate similar predictions, Bayesian model comparison will prefer the simpler of the two ([MacKay, 1995](#)), which is usually the one with the lower number of parameters.

For trajectory models, the momentum model predicts similar dynamics as the diffusion model once λ_m becomes large. This can be seen when inspecting [Equation 9](#), where $e^{-\lambda_m \delta t} \rightarrow 0$ with $\lambda_m \rightarrow \infty$. In this case, its mean becomes \mathbf{z}_{t-1} , and its variance $\sigma_m^2 \delta t^2 / (2\lambda_m)$, matching the structure of the diffusion model's dynamics, [Equation 8](#). These relationships are illustrated in [Figure S7E](#). If both models predict the latent position dynamics equally well, Bayesian model comparison will prefer the less complex diffusion model.

For non-trajectory models, the Gaussian model, [Equation 12](#), can mimic the dynamics of both the stationary and the random model, [Equations 11](#) and [13](#), in different parameter limits. In the limit of $\sigma_g^2 \rightarrow 0$, it approaches the stationary model with $\boldsymbol{\mu} = \mathbf{z}_s$. In the limit of $\sigma_g^2 \rightarrow \infty$, it approaches the random model, and $\boldsymbol{\mu}$ becomes irrelevant. These relationships are illustrated in [Figure S7D](#). If the Gaussian and stationary model, or the Gaussian and random model fit the data equally well, Bayesian model comparison will prefer the less complex stationary or random model, respectively.

Marginalizing over latent position sequences and model parameters

For each SWR, we calculated the likelihood of observing the recorded sequence of spikes under each dynamics model, marginalizing over the sequence of latent positions,

$$\rho(\mathbf{x}_{1:T}|M, \theta_M) = \sum_{\mathbf{z}_{1:T}} \rho(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}|M, \theta_M) = \sum_{\mathbf{z}_{1:T}} \rho(\mathbf{x}_1|\mathbf{z}_1) \rho(\mathbf{z}_1|M, \theta_M) \prod_{t=2}^T \rho(\mathbf{x}_t|\mathbf{z}_t) \rho(\mathbf{z}_t|\mathbf{z}_{1:t-1}, M, \theta_M), \quad (\text{Equation 14})$$

where θ_M are the parameters of model M . This provides a relative comparison of how likely each model is to have generated the spiking data per SWR, for a fixed set of model parameters θ_M . For both the diffusion and momentum models, we used the forward pass of the forward-backward algorithm to calculate the data likelihood (Bishop, 2006) while exploiting their dynamics model structure to keep the computational complexity manageable. Detailed expressions for each model and their associated time and space complexities are provided in the Supplemental information.

To compare models with different numbers of parameters, θ_M , we calculated the data likelihood $\rho(\mathbf{x}_{1:T}|M, \theta_M)$ over a parameter grid and marginalized over these parameters,

$$\rho(\mathbf{x}_{1:T}|M) = \sum_{\theta_M} \rho(\mathbf{x}_{1:T}|M, \theta_M) \rho(\theta_M|M). \quad (\text{Equation 15})$$

For μ and \mathbf{z}_s we used a uniform grid across all spatial locations. For the remaining parameters, σ_d , σ_m , λ_m , and σ_g , we determined the prior in two steps. First, we evaluated the data likelihood, $\rho(\mathbf{x}_{1:T}|M, \theta_M)$ over a uniform grid in log-space: 30 bins from 10 to $630 \frac{\text{cm}}{\sqrt{s}}$ for σ_d , 30 bins 1 to 200 cm for σ_g , 30 bins from 40 to $400 \frac{\text{m}}{s^{3/2}}$ for σ_m , and the vector $[1, 25, 50, 75, 100, 200, 300, 400, 500, 800] \frac{1}{s}$ for λ_m for SWRs; 30 bins from 1 to $100 \frac{\text{cm}}{\sqrt{s}}$ for σ_d , 30 bins 1 to 200 cm for σ_g , 25 bins from 1 to $250 \frac{\text{m}}{s^{3/2}}$ for σ_m , and the vector $[1, 10, 20, 40, 80, 120, 200, 400, 800, 1200, 2000, 4000] \frac{1}{s}$ for λ_m for run snippets. We chose these grids to ensure that most mass of the parameter likelihood $\rho(\mathbf{x}_{1:T}|M, \theta_M)$ lies within this grid. Second, we found for each session and relevant model M the maximum likelihood parameter fits over the parameter grid, $\hat{\theta}_{M,ML} = \text{argmax}_{\theta_M} \rho(\mathbf{x}_{1:T}|M, \theta_M)$, for each SWR, and used the distribution of $\hat{\theta}_{M,ML}$'s across SWRs to fit an appropriate prior, $\rho(\theta_M|M)$. Specifically, for the standard deviation parameters in the diffusion and Gaussian models, σ_g and σ_d , we used an inverse Gamma distribution as the prior. For the Gaussian model, the mean μ was marginalized out before performing these fits. For the standard deviation and decay parameter in the momentum model, σ_m and λ_m , we fit a multivariate log-normal distribution as the prior. As mentioned before, we assumed a uniform prior over \mathbf{z}_s for the stationary model, such that no prior fitting was required.

Computing fraction deviance explained

For each SWR, the model fit quality was quantified (Figure 3D) as the fraction of deviance explained by the best fit model compared to a null model, $1 - \frac{D_{\text{model}}}{D_{\text{null}}}$, where deviance D . is computed as:

$$D_{\text{model}} = 2(\log \rho(\mathbf{x}_{1:T}|M_{\text{sat}}) - \log \rho(\mathbf{x}_{1:T}|M)), \quad (\text{Equation 16})$$

$$D_{\text{null}} = 2(\log \rho(\mathbf{x}_{1:T}|M_{\text{sat}}) - \log \rho(\mathbf{x}_{1:T}|M_{\text{null}})), \quad (\text{Equation 17})$$

Both the null model, M_{null} , and saturated model M_{sat} , assume spikes to be generated by draws from a Poisson distribution, independent across cells and time. They differ, however, in the assumed spike rates: the null model assumes this rate to equal the mean firing rate within population bursts across all neurons within all SWRs in a session at each time bin, while the saturated model assumes this rate to correspond to the observed spike count in each time bin. $\rho(\mathbf{x}_{1:T}|M)$ was computed for each model M as described above. We computed deviance explained on one hand for the best-fit model for each SWR, and on the other hand for the random model, as comparison.

Best-fitting models, and fixed and random effects model comparison of dynamics models

We inferred the distribution of dynamics models underlying the SWRs in a session by random effects analysis (Stephan et al., 2009). Random effects analysis assumes that the dynamics underlying individual SWRs is drawn from a fixed distribution $p(M)$ over the five dynamics models for each session, and recovers this distribution, $\rho(M|\text{all SWRs})$ from the observed SWRs. It does so by using the likelihood distribution over models $\rho(\mathbf{x}_{1:T_n}^{(n)}|M)$ within each SWR n , and thus takes the uncertainty across models within each SWR into account.

Specifically, random effects model comparison assumes a prior over models $\rho(M|\alpha) = r_M$, where r_M is the probability of an SWR being generated according to dynamics model M . Across models, the r_M 's form the vector \mathbf{r} which has a Dirichlet prior $\text{Dir}(\mathbf{r}|\alpha)$ with concentration parameters α . We set α to get a weak uniform distribution across models, equivalent to about 10% of data points (each element of α is set to 10 for analysis of real data, and 2 for simulated data). Based on this generative model, and given the set of all spiking data within a session, $X = \{\mathbf{x}_{1:T_n}^{(n)} | n = 1, \dots, N\}$ where $\mathbf{x}_{1:T_n}^{(n)}$ is the spiking activity within SWR n and N is the total number of SWRs in the session, we compute the posterior $\rho(M|X)$ using Gibbs sampling, as described in Penny et al. (2010). To compare the two spike emission models described above, we use factorial random effects model comparison (Stephan et al., 2009) to infer the posterior

over dynamics/emission models by summing each Dirichlet parameter sample across models with the same dynamics/emission model, respectively (Figures S1K and S1L).

For comparison, we also calculate the inferred distribution of models using the more standard fixed effects analysis, which assumes each SWR is drawn from the same model. Here the likelihood of the spiking data across SWRs within a session, X , is calculated as $p(X|M) = \prod_{n=1}^N p(\mathbf{x}_{1:T_n}^{(n)}|M)$, which we then invert using Bayes rule assuming a uniform prior on M to obtain the posterior distribution over models, $p(M|X)$.

For per-SWR analyses, each SWR was “classified” as the model M that had the highest (marginalized) log-likelihood, $\log p(\mathbf{x}_{1:T}|M)$. The fraction trajectory (non-trajectory) models is the fraction SWRs best described by the diffusion or momentum (stationary, Gaussian, or random) models.

Model recovery from simulated data

For each model M , we simulated 100 position sequences in a 2D continuous space (2 m x 2 m) with a time bin size of $\delta t = 1$ ms, according to the dynamics of the respective model. For models with parameters, we drew these parameters for each trajectory from the prior, $p(\theta_M|M)$, fitted as described above. For all models, we generated neural activity from the simulated trajectory by binning the trajectory into the same 50×50 grid (with 4 cm x 4 cm bins) used for estimating place fields, and, for each place cell, drawing their activity within each time bin from a Poisson distribution with a rate given by the cell’s place field for the specific position in that time bin (Equation 3) with γ_{spike} set to 2.9 (Figure S3A). For each set of trajectories generated for a fixed model M , we applied the same Bayesian model comparison and random effects analysis to the simulated neural data, X_{sim} , as we describe above for the real neural data, resulting in an inferred distribution across models $p(M|X_{\text{sim}})$. Even though we simulated data in time steps of 1 ms, for model recovery we used time bins of $\delta t = 3$ ms, as for the real neural data. F-score (computed as the harmonic mean of the precision, $TP/(TP+FP)$, and sensitivity, $TP/(TP+FN)$, where TP = true positive rate, FP = false positive rate, and FN = false negative rate) was used to quantify our model recovery results. To quantify the reliability of distinguishing trajectory versus non-trajectory dynamics, we assumed true/false corresponded to trajectory/non-trajectory, respectively, and to quantify the reliability of distinguishing momentum from diffusion dynamics we assumed that true/false corresponded to momentum/diffusion, respectively.

Extracting maximum likelihood trajectories

Maximum likelihood trajectories were extracted for all SWRs classified as a trajectory model using the Viterbi algorithm (Bishop, 2006). The Viterbi algorithm finds the sequence of positions that gives the highest data likelihood,

$$\hat{\mathbf{z}}_{1:T,ML} = \underset{\mathbf{z}_{1:T}}{\operatorname{argmax}} \left[p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}|M) \right] = \underset{\mathbf{z}_{1:T}}{\operatorname{argmax}} \left[p(\mathbf{x}_{1:T}|\mathbf{z}_{1:T})p(\mathbf{z}_{1:T}|M) \right], \quad (\text{Equation 18})$$

using the max-sum algorithm applied to our diffusion model, $p(\mathbf{z}_{1:T}|M = \text{diffusion}, \sigma_d = 74 \text{ cm})$ and spike generation model $p(\mathbf{x}_{1:T}|\mathbf{z}_{1:T})$, where 74 cm is the mode of the distribution of maximum likelihood parameter fits for σ_d across SWRs. We chose the diffusion model to take into account temporal continuity, but not bias the trajectories to contain momentum dynamics. The Viterbi algorithm is preferred to simply taking the maximum posterior estimate at each time bin, $\hat{\mathbf{z}}_{t,ML} = \underset{\mathbf{z}_t}{\operatorname{argmax}} p(\mathbf{z}_t|\mathbf{x}_{1:T})$, as the sequence of individually most likely positions does not imply that their sequence is the most likely.

Traditional method for classifying SWRs

To compare our analysis to the traditional method for replay classification, we used the set of SWRs classified as containing a trajectory in Pfeiffer and Foster (2015). For visualization purposes, we implemented the traditional method used for replay classification as described in Pfeiffer and Foster (2015) (Figure 2C). Specifically, for each SWR we binned spiking activity using a sliding window of 20 ms bins, advanced in increments of 5 ms. Within each time bin we found the position estimate as the posterior mean associated with the spike likelihood, $\hat{\mathbf{z}}_t = \sum_{\mathbf{z}} \mathbf{z} p(\mathbf{z}|\mathbf{x}_t)$ where $p(\mathbf{z}|\mathbf{x}_t) \propto p(\mathbf{x}_t|\mathbf{z})$ and \mathbf{x}_t here denotes the sliding window-smoothed spike count. The trajectories plotted in Figure 2C are extracted as the longest sequence of most likely position estimates for which position estimates across consecutive time bins are at most 50 cm apart. Pfeiffer and Foster (2015) classified SWRs as encoding trajectories if consecutive decoded positions were less than 50 cm apart, and had a start-to-end distance of at least 80 cm.

Analysis of place cell activity during movement

To evaluate how the spatio-temporal dynamics of neural activity during SWRs compare to the spatio-temporal activity of place cells during movement, we applied our Bayesian model comparison analysis described above to snippets of neural activity during movement (Figure 4). Run snippets were selected to approximately match the distribution of total distance traveled within the trajectories decoded from SWRs (see Extracting maximum likelihood trajectories and Figure S3). Specifically, we first determined a velocity scaling factor, γ_v that represented, on average, how much faster trajectories within SWRs evolved in comparison to real movement of the animal. For each session, we calculated the velocity of each replay trajectory (total distance/duration) and each run period (a continuous period of time within a session in which the animal was considered “moving” for at least 2 s). We calculated the scaling

factor between the mean replay trajectory velocity and mean run period velocity for each session, and defined the velocity scaling factor γ_v , as the average of the scaling factors across all sessions. This procedure found that the velocity of replay trajectories was on average 19.7 times higher than real movement.

Next, we selected the run snippets. For each session, the duration distribution of run snippets was determined by up-scaling the duration of SWRs by the velocity scaling factor. We selected run snippets by randomly sampling a run period and then randomly sampling a start time within that run period, such that the entire duration of the run snippet fell within the run period. In order to match the distribution of sequence lengths T between SWRs and run snippets, we also up-scaled the time bin size used for run snippet analysis by the velocity scaling factor, to $\delta t = 60$ ms (Figure S3). We then applied the model comparison analysis to the neural data within the set of run snippets exactly as outlined above, except that in Equation 3 we set γ_{spike} to reflect the fact that neural activity came from the same “movement” periods that were used to estimate place fields.

Diffusion coefficient analysis

Following Stella et al. (2019), we asked if the relationship between distance and time for replay trajectories could be described by a power law relationship of the form

$$\sqrt{\langle d(t)^2 \rangle} \propto t^\alpha, \quad (\text{Equation 19})$$

where d is the distance between two points in time within a replay trajectory, and t is the time elapsed between those two time points. An α -value of 0.5 corresponds to Brownian motion, or a random walk, while deviations from 0.5 are inconsistent with Brownian motion. We evaluate if replay trajectories evolve according to Brownian motion by plotting the time-distance relationship in log-space, and use linear regression applied to this log-log plot to find the coefficient α , which corresponds to the slope of the log-log time-distance relationship. To avoid confounds, we restricted this analysis to SWRs best-fit by a trajectory model (either the diffusion or momentum model). Replay trajectories were decoded from neural activity as described in [Extracting maximum likelihood trajectories](#). For each replay trajectory within a session we found the squared distance d_j^2 between all pairs of decoded positions separated by all multiples of the time bin used for trajectory decoding $\Delta t / \delta t$, where Δt is the time elapsed between decoded positions and δt is the decoding time bin. This resulted in a set of distance-time pairs, $(d_j^2, \Delta t_j / \delta t)$, for each decoded position pair j across all trajectories within a session. We then found the mean distance, \bar{d}^2 for each multiple of the decoding time bin, fit a linear regression model per session to $\left(\log \sqrt{\bar{d}^2}, \log \left(\frac{\Delta t}{\delta t} \right) \right)$, which gave the estimate of coefficient α for that session as the slope of the regression.

To assess if the coefficient α was significantly larger than 0.5, we generated 1000 bootstrap samples with replacement of sets of trajectories within a session and applied the same procedure to each session. We then computed the fraction of the bootstrap samples for which α was less than 0.5, which results in the p values reported in the main text and Figure 5. We applied this same analysis across sessions to trajectories decoded from neural activity during run snippets (the same run snippets selected for the model comparison analysis on movement data), as well as to the set of 100 trajectories simulated under diffusion dynamics (as described in [Model recovery from simulated data](#)). We chose to visualize the distance-time relationship in terms of number of time bins, rather than time elapsed so that SWRs and run snippets could be visualized at the same scale.

Splitting SWRs into “home” and “away” events

SWRs were split into “home” and “away” events according to the current location of the animal at SWR onset: a “home” event if the animal is less than 5 cm from the home well, and an “away” event if the animal is elsewhere in the environment. In Figures S5A, S5B, and S6 “away” event is redefined as if the animal is less than 5 cm from the goal well, rather than anywhere in the environment other than within 5 cm of the home well.

Analysis of correlation between replay trajectories and behavior

To quantify the correlation between replay trajectories and behavior, we implemented the method described in Pfeiffer and Foster (2013). Replay trajectories were extracted from SWRs as described in [Extracting maximum likelihood trajectories](#). The future/past path of the animal was defined as the immediate future/past path of the animal for either 10 s after/before the SWR or until a direct distance of 75 cm from the current location of the animal during the SWR was reached, whichever threshold gives the greater direct distance. To calculate the correlation between the replay trajectory and the behavioral path, the angular distance between the replay trajectory and behavioral trajectory was found at a series of progressively larger radii from the current location of the animal during the SWR. Specifically, the crossing points between the behavioral trajectory and the replay trajectory with a circle centered on current location of the animal was found for each radius, and the angular distance between these two crossing points was calculated (as visualized in Figure 7A). The minimum radius was 5 cm, increased in increments of 3 cm until either 75 cm or the end of the replay trajectory was reached. For each session, a histogram of angular displacements was calculated using all crossings across SWRs within the session, with SWRs split into “home” and “away” events as described in [Splitting SWRs into “home” and “away” events](#). This analysis was applied separately to either all SWRs classified as a trajectory model or all SWRs previously classified by the traditional method. Chance correlation was found by 2000 shuffles in which a behavioral path and SWR was selected at random from any session, the behavioral

path was shifted to the current location of the animal during the SWR, and the angular distance was calculated using the same method described above.

Figure details

Figure 1. Panel **B** shading visualizes “moving” and “not moving” periods as described in [Data acquisition and behavior](#), and the example place fields in panel **C** were computed using the place field data from “moving” periods as described in [Estimating place fields](#). The spike generation model in panel **C** and the dynamics models in panel **D** are described in the [Spike generation model](#) and [Dynamics models](#) sections. The graphical models in **D** describe the statistical relationship between variables in each dynamics model, whose details are given in [Dynamics models](#).

Figure 2. The heatmaps in panel **A** were computed from the marginal decoded position per time bin $p(\mathbf{z}_t|\mathbf{x}_{1:T}, M)$ for each dynamics model described in [Description of state-space models](#). For the diffusion and momentum models, the marginal decoded positions were found by the forward-backward algorithm ([Bishop, 2006](#)). The other, non-trajectory, models had a simpler form that supported a simpler computation of the position estimates. We normalized these position estimates separately for each time bin, $\tilde{p}(\mathbf{z}_t|\mathbf{x}_{1:T}, M) = p(\mathbf{z}_t|\mathbf{x}_{1:T}, M) / \sum_{\mathbf{z}_t} p(\mathbf{z}_t|\mathbf{x}_{1:T}, M)$, summed them over time $\tilde{p}(\mathbf{z}_{1:T}|\mathbf{x}_{1:T}, M) = \sum_{t=1}^T \tilde{p}(\mathbf{z}_t|\mathbf{x}_{1:T}, M)$, and plotted this sum. While the sum is no longer a real probability distribution, it is nonetheless useful for visualization. In panel **B**, the distribution over models per SWR was computed as described in [Marginalizing over latent position sequences and model parameters](#). In panel **C**, the heatmaps were computed by obtaining the decoded position per time bin $p(\mathbf{z}_t|\mathbf{x}_t)$ as calculated by the traditional method described in [Traditional method for classifying SWRs](#), then normalized, summed, and plotted as described for panel **A**. The trajectory plotted is as described in [Traditional method for classifying SWRs](#).

Figure 3. The distribution over models in panel **A** was inferred by random effects analysis as explained in [Best-fitting models, and fixed and random effects model comparison of dynamics models](#), the comparison to the traditional method as described in [Traditional method for classifying SWRs](#), the simulated data model recovery as described in [Model recovery from simulated data](#), and the deviance explained as described in [Computing fraction deviance explained](#).

Figure 4. The run snippets visualized in panel **A** were selected as described in [Analysis of place cell activity during movement](#) for example session 1 (rat 1, day 1). The distribution over models was computed by random effects analysis as described in [Best-fitting models, and fixed and random effects model comparison of dynamics models](#). The example run snippets decoded positions are visualized following the same procedure outlined for [Figure 2A](#) (above), and the distribution over models per-run snippet was computed as described in [Marginalizing over latent position sequences and model parameters](#).

Figure 5. The heatmaps in panels **A** and **B** were computed as described for [Figure 2A](#) (above), and the decoded trajectories were obtained as described in [Extracting maximum likelihood trajectories](#). For panel **C**, the time-distance relationship, estimated diffusion coefficient in panel, and bootstrap significance testing (inset) are described in [Diffusion coefficient analysis](#).

Figure 6. The data presented here agglomerates SWRs best fit by a trajectory model, as described in [Best-fitting models, and fixed and random effects model comparison of dynamics models](#), across all sessions, and splits SWRs into “home” and “away” events as described in [Splitting SWRs into “home” and “away” events](#).

Figure 7. “Trajectory model” SWRs are selected as described in [Best-fitting models, and fixed and random effects model comparison of dynamics models](#), “previously classified” SWRs are selected as described in [Traditional method for classifying SWRs](#), and SWRs are split into “home” and “away” events as described in [Splitting SWRs into “home” and “away” events](#). The angular distance was quantified as described in [Analysis of correlation between replay trajectories and behavior](#).

Neuron, Volume 110

Supplemental information

**A large majority of awake hippocampal sharp-wave
ripples feature spatial trajectories with momentum**

Emma L. Krause and Jan Drugowitsch

Supplementary Notes

Momentum dynamics model formulation as a second-order Markov model

As described in Methods, the momentum model assumes that the velocity follows an Ornstein-Uhlenbeck (OU) that is independent in each spatial dimension $j \in \{x, y\}$, and given by

$$\frac{dv_j}{dt} = -\lambda_m v_j(t) + \sigma_m \eta_{j,m}(t), \quad (\text{S1})$$

where λ_m determines the decay, σ_m is the diffusion coefficient, and $\eta_{j,m}(t)$ is a Gaussian white noise process. The latent position is fully determined by this velocity, that is $dz_j/dt = v_j(t)$.

We turn this dynamics model into a discrete-time second-order Markov model in two steps. We will do so for each spatial dimension separately, and will drop the \cdot_j subscript to keep the notation uncluttered. First, note that the linear-Gaussian nature of the OU process implies that $v(t)$ remains Gaussian at any point in time. Furthermore, if $v(t)$ is known at time t , its moments evolve for $\delta t \geq 0$ according to

$$\langle v(t + \delta t) \rangle = v(t)e^{-\lambda_m \delta t}, \quad \text{var}(v(t + \delta t)) = \frac{\sigma_m^2}{2\lambda_m} (1 - e^{-2\lambda_m \delta t}). \quad (\text{S2})$$

As a consequence, we can write $v(t + \delta t)$ as

$$v(t + \delta t) = v(t)e^{-\lambda_m \delta t} + \varepsilon_{\delta t}, \quad \text{with } \varepsilon_{\delta t} \sim \mathcal{N}\left(0, \frac{\sigma_m^2}{2\lambda_m} (1 - e^{-2\lambda_m \delta t})\right). \quad (\text{S3})$$

Second, we approximate dz/dt by finite differences, resulting in

$$\frac{z(t) - z(t - \delta t)}{\delta t} \approx v(t). \quad (\text{S4})$$

Substituting this approximation into Eq. (S3) results in

$$\frac{z(t + \delta t) - z(t)}{\delta t} = \frac{z(t) - z(t - \delta t)}{\delta t} e^{-\lambda_m \delta t} + \varepsilon_{\delta t}. \quad (\text{S5})$$

Solving for $z(t + \delta t)$ yields

$$z(t + \delta t) = (1 + e^{-\lambda_m \delta t}) z(t) - e^{-\lambda_m \delta t} z(t - \delta t) + \delta t \varepsilon_{\delta t}, \quad (\text{S6})$$

such that $z(t + \delta t)$ is distributed as

$$p(z(t + \delta t) | z(t), z(t - \delta t)) = \mathcal{N}\left(z(t + \delta t) | (1 + e^{-\lambda_m \delta t}) z(t) - e^{-\lambda_m \delta t} z(t - \delta t), \frac{\sigma_m^2 \delta t^2}{2\lambda_m} (1 - e^{-2\lambda_m \delta t})\right). \quad (\text{S7})$$

If we let $\mathbf{z}(t) = (z_x(t), z_y(t))^T$, and consider that $z_x(t)$ and $z_y(t)$ evolve independently, their joint evolution can be written as the second-order Markov chain

$$p(\mathbf{z}(t) | \mathbf{z}(t - \delta t), \mathbf{z}(t - 2\delta t)) = \mathcal{N}\left(\mathbf{z}(t) | (1 + e^{-\lambda_m \delta t}) \mathbf{z}(t - \delta t) - e^{-\lambda_m \delta t} \mathbf{z}(t - 2\delta t), \frac{\sigma_m^2 \delta t^2}{2\lambda_m} (1 - e^{-2\lambda_m \delta t}) \mathbf{I}\right). \quad (\text{S8})$$

In Methods, we provide this equation with time-discretized indices on the \mathbf{z} 's.

Due to the second-order Markov chain-nature of this process, we need to handle the first time-discretized transition $p(\mathbf{z}(\delta t) | \mathbf{z}(0))$ separately, as it cannot depend on $\mathbf{z}(-\delta t)$. To do so, we again rely on the finite-difference approximation, such that $\mathbf{v}(\delta t)\delta t = \mathbf{z}(\delta t) - \mathbf{z}(0)$. If we now assume a prior $\mathbf{v}(\delta t) \sim \mathcal{N}\left(\mathbf{0}, \sigma_{m0}^2 / \sqrt{\delta t} \mathbf{I}\right)$, the first step becomes

$$p(\mathbf{z}(\delta t) | \mathbf{z}(0)) = \mathcal{N}\left(\mathbf{z}(\delta t) | \mathbf{z}(0), \sigma_{m0}^2 \delta t \mathbf{I}\right). \quad (\text{S9})$$

Diffusion scaling of the momentum dynamics model

The diffusion scaling factor ξ is defined by how the mean square displacement $r(t)^2$ scales with time. This displacement is defined as the distance from the origin, $r(t) = \sqrt{(z_x(t) - z_x(0))^2 + (z_y(t) - z_y(0))^2}$, such that $r(t)^2 = (z_x(t) - z_x(0))^2 + (z_y(t) - z_y(0))^2$. For a scaling factor ξ , this mean square displacement scales as

$$\langle r(t)^2 \rangle \sim t^\xi. \quad (\text{S10})$$

In the main text and in Stella et al. (2019), we report scaling of the root mean square displacement $\sqrt{\langle r(t)^2 \rangle} \sim t^\alpha$, that relates to the diffusion scaling factor by $\xi = 2\alpha$. For standard diffusion dynamics, this diffusion scaling factor is $\xi = 1$ or $\alpha = 1/2$. Here, we consider this scaling factor for our momentum dynamics model.

To find this scaling factor, note that the momentum model features independent, and equivalent, dynamics along the x - and y -directions. Therefore, $\langle r(t)^2 \rangle = \langle (z_x(t) - z_x(0))^2 \rangle + \langle (z_y(t) - z_y(0))^2 \rangle = 2 \langle (z_x(t) - z_x(0))^2 \rangle$. Thus, it is sufficient to find this scaling for one spatial dimension. For this reason, we will drop the dimension-identifying subscripts in what follows.

The momentum dynamic's model is defined by Eq. (S3), $dz/dt = v(t)$, and $v(0) \sim \mathcal{N}(0, \sigma_{m0}^2)$. To find $\langle (z(t) - z(0))^2 \rangle$, let us first consider the moments of $v(t)$, which can be shown to follow

$$\langle v(t) \rangle = \langle v(0) \rangle e^{-\lambda_m t} = 0, \quad (\text{S11})$$

$$\text{cov}(v(s), w(t)) = \sigma_{m0}^2 e^{-\lambda_m(t+s)} + \frac{\sigma_m^2}{2\lambda_m} \left(e^{-\lambda_m|t-s|} - e^{-\lambda_m(t+s)} \right), \quad (\text{S12})$$

$$\text{var}(v(t)) = \sigma_{m0}^2 e^{-2\lambda_m t} + \frac{\sigma_m^2}{2\lambda_m} (1 - e^{-2\lambda_m t}), \quad (\text{S13})$$

where $v(s)$ and $w(t)$ are different realizations of the velocity process, with different associated $v(0)$ and $w(0)$ drawn from the prior. Except for the terms introduced by the prior, these moments reflect those of a standard Ornstein-Uhlenbeck process.

In the momentum dynamics model, the location $z(t)$ is simply the sum of the momentary velocities, $z(t) = z(0) + \int_0^t v(s) ds$, and so a linear function of them. As these velocities are Gaussian, so is the location. Therefore, it is sufficient to find its first- and second-order moments to fully describe the location evolution. As the mean velocity is zero, the location's mean is simply given by

$$\langle z_j(t) \rangle = z_j(0) + \langle v_j(0) \rangle \int_0^t e^{-\lambda_m s} ds = z_j(0). \quad (\text{S14})$$

To find $\langle (z(t) - z(0))^2 \rangle$, we first find $\langle z(t)^2 \rangle$, which is given by

$$\begin{aligned} \langle z(t)^2 \rangle &= z(0)^2 + 2z(0) \left\langle \int_0^t v(s) ds \right\rangle + \int_0^t \int_0^t \text{cov}(v(s), w(u)) du ds \\ &= z(0)^2 + \int_0^t \int_0^t \sigma_{m0}^2 e^{-\lambda_m(s+u)} du ds \\ &\quad + \frac{\sigma_m^2}{2\lambda_m} \int_0^t \left[\int_0^s \left(e^{-\lambda(s-u)} - e^{-\lambda_m(s+u)} \right) du + \int_s^t \left(e^{-\lambda(u-s)} - e^{-\lambda_m(s+u)} \right) du \right] ds \\ &= z(0)^2 + \frac{\sigma_{m0}^2}{\lambda_m^2} (1 - e^{-\lambda_m t})^2 + \frac{\sigma_m^2}{2\lambda_m^3} (-3 + 4e^{-\lambda_m t} - e^{-2\lambda_m t} + 2\lambda_m t). \end{aligned} \quad (\text{S15})$$

Here we have split the $|s - u|$ in the OU covariance into two separate integrals. The final expression for $\langle (z(t) - z(0))^2 \rangle$ thus reads

$$\langle (z(t) - z(0))^2 \rangle = \frac{\sigma_{v0}^2}{\lambda_m^2} (1 - e^{-\lambda_m t})^2 + \frac{\sigma_m^2}{\lambda_m^2} t + \frac{\sigma_m^2}{2\lambda_m^3} (-3 + 4e^{-\lambda_m t} - e^{-2\lambda_m t}). \quad (\text{S16})$$

To gain insight into its time scaling, let us consider the limit in which $\lambda_m \ll t$, that is, where the velocity evolves very slowly when compared to t . In this limit, the mean squared displacement becomes

$$\lim_{\lambda_m \rightarrow 0} \langle (z(t) - z(0))^2 \rangle = \sigma_{m0}^2 t^2 + \frac{\sigma_m^2}{3} t^3. \quad (\text{S17})$$

Therefore, in this limit, the momentum model features a super-diffusive diffusion scaling factor of $\xi = 3$ or $\alpha = 3/2$. If we instead look at long times t , we find

$$\lim_{t \rightarrow \infty} \langle (z(t) - z(0))^2 \rangle = \frac{\sigma_m^2}{\lambda_m^2} t, \quad (\text{S18})$$

such that, in this limit, the diffusion scaling factor approaches that of a standard diffusion, $\xi = 1$. This is to be expected as, in this limit, positive and negative velocity deflections cancel. Note that all of the above only applies for unbounded trajectories. Thus, for our use case it only applies to trajectories that are sufficiently far away from the boundary. For those trajectories, we expect to see super-diffusivity as long as λ_m is sufficiently small when compared to t .

Computational complexity of evaluating the model evidence

Here we derive the computational complexity associated with computing the model evidence $p(\mathbf{x}_{1:T}|M, \theta_M)$ for different models M , and where $\mathbf{x}_{1:T}$ is the sequence of neural population activity patterns in time bins $t = 1, \dots, T$ across N neurons, and where θ_M is the parameter vector of model M . In general, the model evidence is given by

$$p(\mathbf{x}_{1:T}|M, \theta_M) = \sum_{\mathbf{z}_{1:T}} p(\mathbf{x}_1|\mathbf{z}_1) p(\mathbf{z}_1|M, \theta_M) \prod_{t=2}^T p(\mathbf{x}_t|\mathbf{z}_t) p(\mathbf{z}_t|\mathbf{z}_{1:t-1}, M, \theta_M). \quad (\text{S19})$$

This expression simplifies in different ways for different models. We will first discuss the computational complexity for the likelihood $p(\mathbf{x}_t|\mathbf{z}_t)$, which is shared across all models. Second, we will discuss the overall computational complexity, first for non-trajectory models, and then for trajectory models.

Likelihood

The likelihood $p(\mathbf{x}_t|\mathbf{z}_t)$ is computed for a population of N neurons over a spatial grid with K bins. We need to store neural activity of N neurons for each time bin, coming at space complexity $\mathcal{O}(NT)$. When evaluating the likelihood, we need to do so for each neuron, each \mathbf{z}_t , and each time bin, such that the time complexity is $\mathcal{O}(NKT)$.

Non-trajectory models

Non-trajectory models feature model dynamics that are independent across time-steps, that is $p(\mathbf{z}_t|\mathbf{z}_{1:t-1}, M, \theta_M) = p(\mathbf{z}_t|M, \theta_M)$. Thus, the model evidence simplifies to

$$p(\mathbf{x}_{1:T}|M, \theta_M) = \prod_{t=1}^T \sum_{\mathbf{z}_t} p(\mathbf{x}_t|\mathbf{z}_t) p(\mathbf{z}_t|M, \theta_M), \quad (\text{S20})$$

where $p(\mathbf{z}_t|M, \theta_M)$ differs across the various dynamics models.

Stationary dynamics model. For the stationary dynamics model, $M = \text{stationary}$, $p(\mathbf{z}_t|M, \theta_M) = \delta_{\mathbf{z}_t, \mathbf{z}_s}$, where $\theta_M = \{\mathbf{z}_s\}$ is the only model parameter, which has a uniform prior, $p(\mathbf{z}_s) = 1/K$. For this model, the model evidence simplifies to

$$p(\mathbf{x}_{1:T}|M, \theta_M) = \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{z}_t = \mathbf{z}_s), \quad (\text{S21})$$

which has a space and time complexity of $\mathcal{O}(NT)$. The time complexity is lower than that of the likelihood computation, as the likelihood in the above expression is evaluated for a single $\mathbf{z}_t = \mathbf{z}_s$, rather than all of them. Once we marginalize over the model parameters, the model evidence becomes

$$p(\mathbf{x}_{1:T}|M) = \frac{1}{K} \sum_{\mathbf{z}_s} \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{z}_t = \mathbf{z}_s), \quad (\text{S22})$$

with a space complexity of $\mathcal{O}(NT)$ and time complexity of $\mathcal{O}(NKT)$.

Gaussian dynamics model. For the Gaussian dynamics model, $M = \text{gaussian}$, $p(\mathbf{z}_t|M, \theta_M) = \mathcal{N}(\mathbf{z}_t|\boldsymbol{\mu}, \sigma_g^2 \mathbf{I})$, with model parameters $\theta_M = \{\boldsymbol{\mu}, \sigma_g^2\}$, and uniform $\boldsymbol{\mu}$ -prior, $p(\boldsymbol{\mu}) = 1/K$. Thus, the model evidence is given by

$$p(\mathbf{x}_{1:T}|M, \theta_M) = \prod_{t=1}^T \sum_{\mathbf{z}_t} p(\mathbf{x}_t|\mathbf{z}_t) \mathcal{N}(\mathbf{z}_t|\boldsymbol{\mu}, \sigma_g^2 \mathbf{I}), \quad (\text{S23})$$

which has a space and time complexity of $\mathcal{O}(NKT)$. Marginalizing over $\boldsymbol{\mu}$ yields

$$p(\mathbf{x}_{1:T}|M, \sigma_g^2) = \frac{1}{K} \sum_{\boldsymbol{\mu}} \prod_{t=1}^T \sum_{\mathbf{z}_t} p(\mathbf{x}_t|\mathbf{z}_t) \mathcal{N}(\mathbf{z}_t|\boldsymbol{\mu}, \sigma_g^2 \mathbf{I}), \quad (\text{S24})$$

with a space complexity of $\mathcal{O}(NKT)$ and a time complexity of $\mathcal{O}(NK^2T)$. Marginalization over σ_g^2 is performed over a fixed grid of parameter values, and thus only adds an additional constant scaling to the complexity.

Random dynamics model. The random dynamics model, $M = \text{random}$ has likelihood $p(\mathbf{z}_t|M) = 1/K$ and no parameters. Thus, it's associated model evidence is

$$p(\mathbf{x}_{1:T}|M) = \frac{1}{K^T} \prod_{t=1}^T \sum_{\mathbf{z}_t} p(\mathbf{x}_t|\mathbf{z}_t), \quad (\text{S25})$$

with space and time complexities $\mathcal{O}(NKT)$.

Trajectory models

We evaluate the model evidence of trajectory models using the forward pass of the forward-backward algorithm for Hidden Markov Models (HMMs). Assuming we know $p(\mathbf{z}_{t-1}|\mathbf{x}_{1:t-1})$, and implicitly conditioning on M and θ_M , the forward pass operates in two steps. In the *prediction* step, we find $p(\mathbf{z}_t|\mathbf{x}_{1:t-1})$ by marginalization,

$$p(\mathbf{z}_t|\mathbf{x}_{1:t-1}) = \sum_{\mathbf{z}_{t-1}} p(\mathbf{z}_t|\mathbf{z}_{t-1}) p(\mathbf{z}_{t-1}|\mathbf{x}_{1:t-1}). \quad (\text{S26})$$

The *update* step includes the latest observation \mathbf{x}_t by Bayes' rule,

$$p(\mathbf{z}_t|\mathbf{x}_{1:t}) = \frac{p(\mathbf{x}_t|\mathbf{z}_t) p(\mathbf{z}_t|\mathbf{x}_{1:t-1})}{p(\mathbf{x}_t|\mathbf{x}_{1:t-1})} \quad (\text{S27})$$

For this step we first compute the numerator for all \mathbf{z}_t , and then find the normalizing denominator by $p(\mathbf{x}_t|\mathbf{x}_{1:t-1}) = \sum_{\mathbf{z}_t} p(\mathbf{x}_t|\mathbf{z}_t) p(\mathbf{z}_t|\mathbf{x}_{1:t-1})$. These two steps are alternated to compute $p(\mathbf{z}_1|\mathbf{x}_1)$, then $p(\mathbf{z}_2|\mathbf{x}_{1:2})$, and so on, until $p(\mathbf{z}_T|\mathbf{x}_{1:T})$. This update step's normalizer is in turn used across time to compute the model evidence by

$$p(\mathbf{x}_{1:T}) = p(\mathbf{x}_1) \prod_{t=2}^T p(\mathbf{x}_t|\mathbf{x}_{1:t-1}). \quad (\text{S28})$$

For both trajectory models, model parameters are marginalized over a fixed parameter grid, such that this marginalization only adds a constant scaling to the respective model complexities.

Diffusion dynamics model. The diffusion dynamics model, $M = \text{diffusion}$, uses a Gaussian transition structure, $p(\mathbf{z}_t|\mathbf{z}_{t-1}) \propto \mathcal{N}(\mathbf{z}_t|\mathbf{z}_{t-1}, \sigma_d^2 \mathbf{I})$, appropriately normalized on the discretized \mathbf{z}_t grid. Naively, the prediction step could be computed by

$$p(\mathbf{z}_T|\mathbf{x}_{1:t-1}) = \sum_{\mathbf{z}_{t-1}} \tilde{\mathcal{N}}(\mathbf{z}_t|\mathbf{z}_{t-1}, \sigma_d^2 \mathbf{I}) p(\mathbf{z}_t|\mathbf{x}_{1:t-1}), \quad (\text{S29})$$

where $\tilde{\mathcal{N}}(\cdot)$ denotes a Gaussian distribution normalized on the grid, and with space complexity $\mathcal{O}(K^2)$ (assuming a pre-computed $p(\mathbf{z}_t|\mathbf{z}_{t-1})$) and time complexity $\mathcal{O}(K^2T)$ across all time steps. The update step's time and space complexity instead is $\mathcal{O}(NKT)$ across all time steps. Thus, the overall time and space complexity of such a naive implementation is $\mathcal{O}(NK^2T)$. We can further reduce this complexity by factorizing the transition structure into $p(\mathbf{z}_t|\mathbf{z}_{t-1}) \propto \mathcal{N}(z_{x,t}|z_{x,t-1}, \sigma_d^2) \mathcal{N}(z_{y,t}|z_{y,t-1}, \sigma_d^2)$. Then, the prediction step becomes

$$p(z_{x,t}, z_{y,t}|\mathbf{x}_{1:t-1}) = \sum_{z_{y,t-1}} \tilde{\mathcal{N}}(z_{y,t}|z_{y,t-1}, \sigma_d^2) \sum_{z_{x,t-1}} \tilde{\mathcal{N}}(z_{x,t}|z_{x,t-1}, \sigma_d^2) p(z_{x,t-1}, z_{y,t-1}|\mathbf{x}_{1:t-1}), \quad (\text{S30})$$

with space complexity $\mathcal{O}(K)$ (again assuming pre-computation of $p(z_{j,t}|z_{j,t-1})$ for $j \in \{x, y\}$), and time complexity $\mathcal{O}(KT)$ across all time steps. Overall, this reduces the space and time complexity of computing the model evidence to $\mathcal{O}(NKT)$.

Momentum dynamics model. The momentum dynamics model, $M = \text{momentum}$ yields a second-order HMM with a transition structure $p(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{z}_{t-2})$ whose exact expression is given further above and in the main text. For such second-order HMMs, we propagate the joint posterior $p(\mathbf{z}_t, \mathbf{z}_{t-1}|\mathbf{x}_{1:t})$ through time, as follows. The prediction step changes to

$$p(\mathbf{z}_t, \mathbf{z}_{t-1}|\mathbf{x}_{1:t-1}) = \sum_{\mathbf{z}_{t-2}} p(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{z}_{t-2}) p(\mathbf{z}_{t-1}, \mathbf{z}_{t-2}|\mathbf{x}_{1:t-1}). \quad (\text{S31})$$

As the posterior itself requires space $\mathcal{O}(K^2)$, a naive implementation of this prediction step has space complexity $\mathcal{O}(K^3)$ (again assuming pre-computation of $p(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{z}_{t-2})$) and time complexity $\mathcal{O}(K^3T)$ across all time steps. The update step changes to

$$p(\mathbf{z}_t, \mathbf{z}_{t-1}|\mathbf{x}_{1:t}) = \frac{p(\mathbf{x}_t|\mathbf{z}_t)p(\mathbf{z}_t, \mathbf{z}_{t-1}|\mathbf{x}_{1:t-1})}{p(\mathbf{x}_t|\mathbf{x}_{1:t-1})}, \quad (\text{S32})$$

with $\mathcal{O}(NK^2T)$ space and time complexity. The model evidence is computed as before, from the product of update step normalizers across times. Overall, this results in a space and time complexity of $\mathcal{O}(NK^3T)$. This can be reduced by factorizing the transition dynamics into $p(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{z}_{t-2}) = p(z_{x,t}|z_{x,t-1}, z_{x,t-2})p(z_{y,t}|z_{y,t-1}, z_{y,t-2})$, simplifying the prediction step to

$$p(\mathbf{z}_t, \mathbf{z}_{t-1}|\mathbf{x}_{1:t-1}) = \sum_{z_{y,t-2}} p(z_{y,t}|z_{y,t-1}, z_{y,t-2}) \sum_{z_{x,t-2}} p(z_{x,t}|z_{x,t-1}, z_{x,t-2}) p(\mathbf{z}_{t-1}, \mathbf{z}_{t-2}|\mathbf{x}_{1:t-1}), \quad (\text{S33})$$

with reduced space complexity $\mathcal{O}(K^2)$ and time complexity $\mathcal{O}(K^2T)$. Overall, this reduces the computational complexity of computing the model evidence to $\mathcal{O}(NK^2T)$ in both space and time.

Supplementary Figures

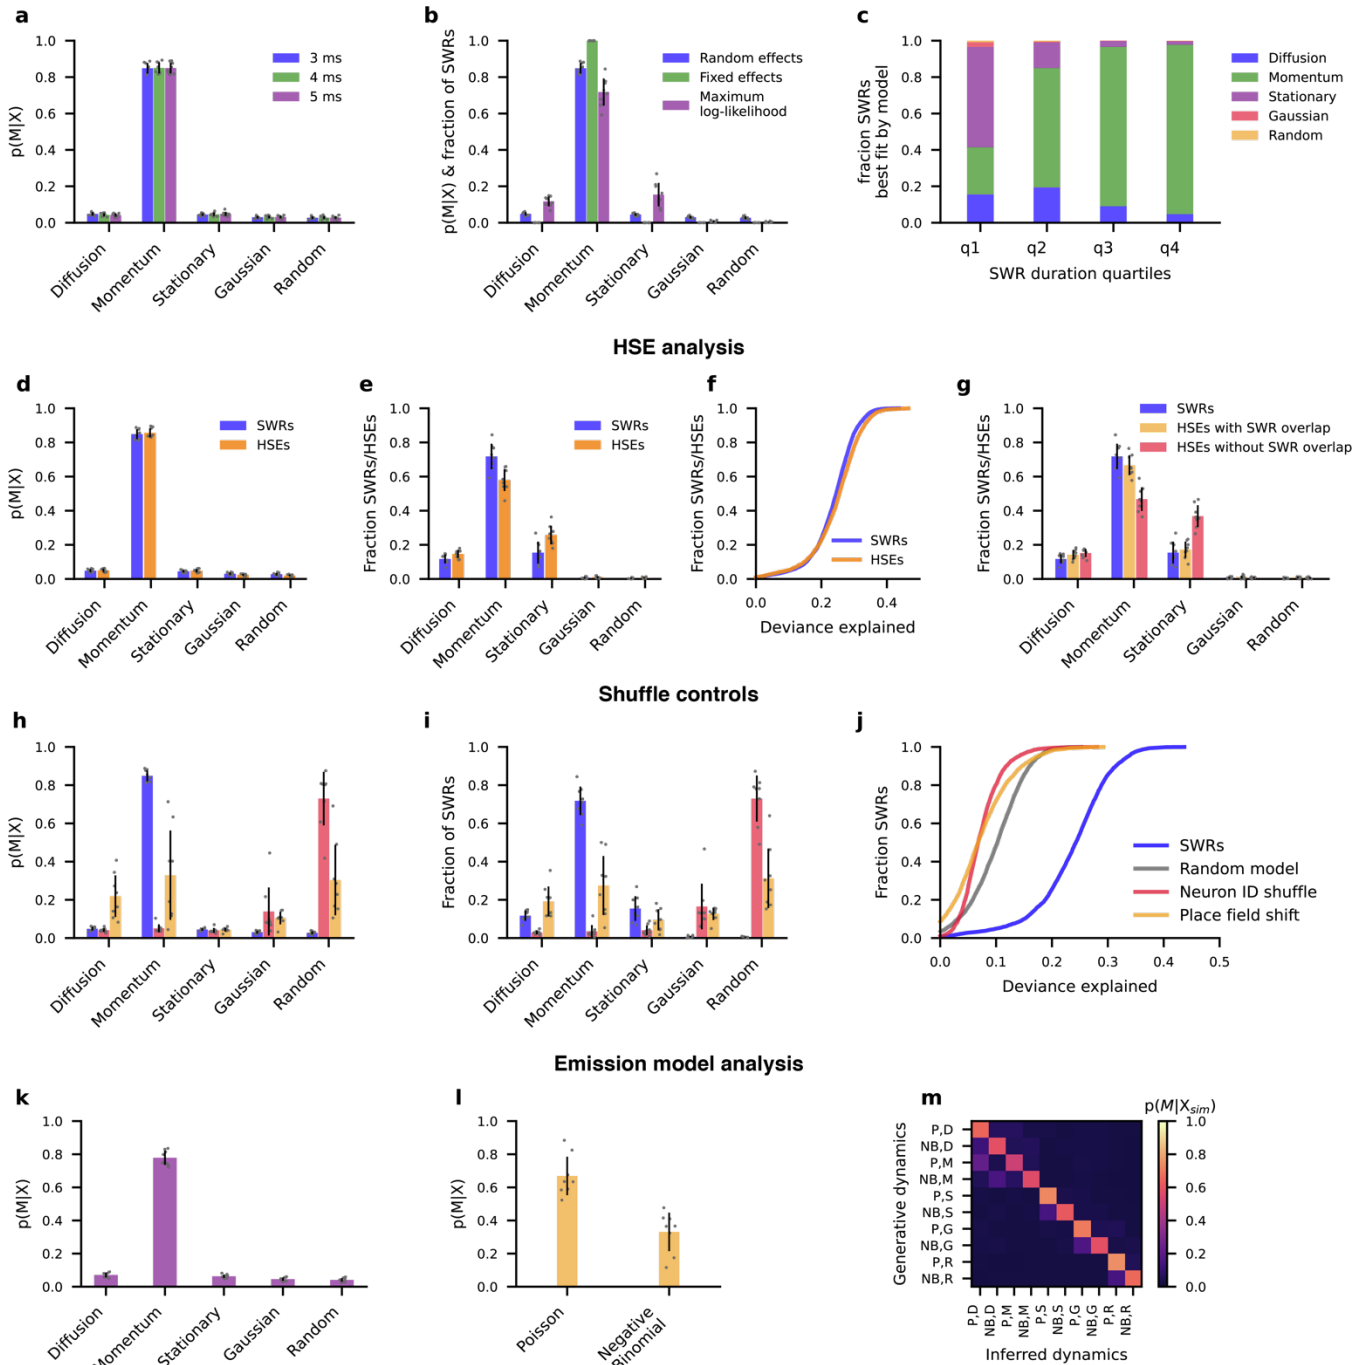


Figure S1. Analysis of model variants for our Bayesian model comparison method. Related to Figure 3.

a. Bayesian model comparison gives consistent results across different time bin sizes. We replicated our Bayesian model comparison analysis (Fig. 3a) for different time bin sizes (colors). The results for 3ms bins are described in the main text. Here we plot the inferred distributions of dynamics models underlying the generation of SWR spikes, using random effects model comparison (mean \pm SD across sessions; gray

dots = individual sessions). All time bin sizes gave very similar results (exceedance probability of momentum model ≈ 1 for all sessions and all time bin sizes).

b. Bayesian model comparison takes the relative likelihood distribution into account. Random effects analysis takes relative likelihood distribution into account and is less sensitive to outliers than standard fixed effects Bayesian model comparison. For comparison to our random effects Bayesian model comparison presented in the main text (blue), we also compute the inferred distribution of models given the data from all SWRs, $p(M|X)$, by fixed effects Bayesian model comparison (green), as well as the fraction of SWRs with maximum likelihood for each model (purple), mean \pm SD across sessions; gray dots = individual sessions. Fixed effects Bayesian model comparison compares models by summing across the log-likelihoods for all SWRs within a session, before inverting them according to Bayes rule to find the shown posterior over models. This makes this approach sensitive to outlier SWRs for which one model fits the data much better than the other models (Stephan et al., 2009). As a result, it assigns probability 1 to the momentum model for all sessions. We furthermore show for each dynamics model the fraction of SWRs within each session that yield that maximum likelihood (purple). While this fraction is insensitive to the relative difference of likelihoods across models within each SWR, it is a useful measure to decide for individual SWRs the most likely model (Fig. 2b), which we use in the main text for selecting SWRs best-fit by a trajectory model. The random effects approach, in contrast, only provides an aggregate measure, rather than a per-SWR measure.

c. Our Bayesian model comparison approach favors simpler models over more complex ones when both explain the spiking data equally well (MacKay, 1995). To show this, we here plot the distribution of best-fit models, grouping SWRs into quartiles by duration. The fraction of SWRs in each quartile best described by each dynamics model is indicated with a stacked bar plot. This illustrates how Bayesian model comparison prefers simpler models for sparser data: for short location sequences that only span a few time bins, the limited number of spikes might not support telling apart pure diffusion dynamics from dynamics with momentum, even if the true underlying dynamics feature momentum. In such circumstances, the associated SWR will be classified as momentum-free diffusion rather than diffusion with momentum. This can be seen by the larger fraction of the momentum model explaining SWRs better than the diffusion model as SWR lengths increase. For even shorter sequences with even less spikes, location sequences with momentum might equally well be explained by a random, or even a stationary model, such that our comparison will prefer non-trajectory dynamics (see first quartile). This illustrates how our method handles the irreducible uncertainty inherent to the limited data available in a principled way.

d-g. Bayesian model comparison applied to high-synchrony events (HSEs). Bayesian model comparison gives similar results when applied to HSEs as when applied to SWRs for random effects analysis (d), maximum log-likelihood (e), and fraction deviance explained (f), though a slightly smaller/larger fraction of HSEs are better described by the momentum/stationary models (exceedance probability of momentum model ≈ 1 for all sessions for both SWRs and HSEs). f. Splitting HSEs into those that overlap with SWRs ($54.3 \pm 10.4\%$, mean \pm SD across sessions) and those that do not reveals that the higher fraction of HSEs best fit by the stationary model is driven by HSE's that do not overlap with SWRs, for which a larger fraction of HSE's are best fit by the stationary model.

h-j. Bayesian model comparison applied to shuffled data. We shuffled data by either randomly shuffling the neural identity separately for each SWR (red), or by randomly shifting place fields (assuming connected top/bottom and left/right arena edges), separately for each neuron, but consistently across all SWRs (yellow). Both perturbations resulted in the momentum model to cease being the dominant model in both the random effects model comparison (h), as well as the distribution of maximum-likelihood models (i). Exceedance probabilities (mean \pm SD across sessions) for neuron ID shuffle = random: 0.93 ± 0.2 , Gaussian: 0.07 ± 0.2 , and for place field shift = random: 0.29 ± 0.42 , diffusion: 0.22 ± 0.37 , momentum: 0.5 ± 0.5 . Shuffling neuron identities resulted in a stronger drop of spatiotemporal structure and made the random model dominant. Both perturbations resulted in a poorer fit to spiking data, as illustrated by the drop

in deviance explained of the best fit model to below that of the random model applied to non-shuffled SWRs (j).

k-m. Analysis of spike emission model. Factorial random effects model comparison (Stephan et al., 2009) reveals that neural data during SWRs is better fit by assuming place fields are scaled by a constant factor across neurons (Poisson, or “P”, emission model variant), rather than assuming this scaling factor to vary independently across neurons and SWRs (negative binomial, or “NB”, emission model variant; see STAR Methods). Here we plot the inferred random effects model comparison distributions collapsed over emission models (k; exceedance probability of momentum model ≈ 1 for all sessions) or dynamics models (l; exceedance probability of Poisson model ≈ 1 for all sessions); mean \pm SD across sessions and gray dots = individual sessions. **m.** Simulating data according to each dynamics and emission model combination (similar to main text Fig. 3c), we find reliable model recovery. For the limited data of simulated SWRs, data generated with the more complex “NB” emission model variant is more likely mistaken for data generated with the simpler “P” emission model variant, illustrating the implicit model complexity penalty of Bayesian model comparison.

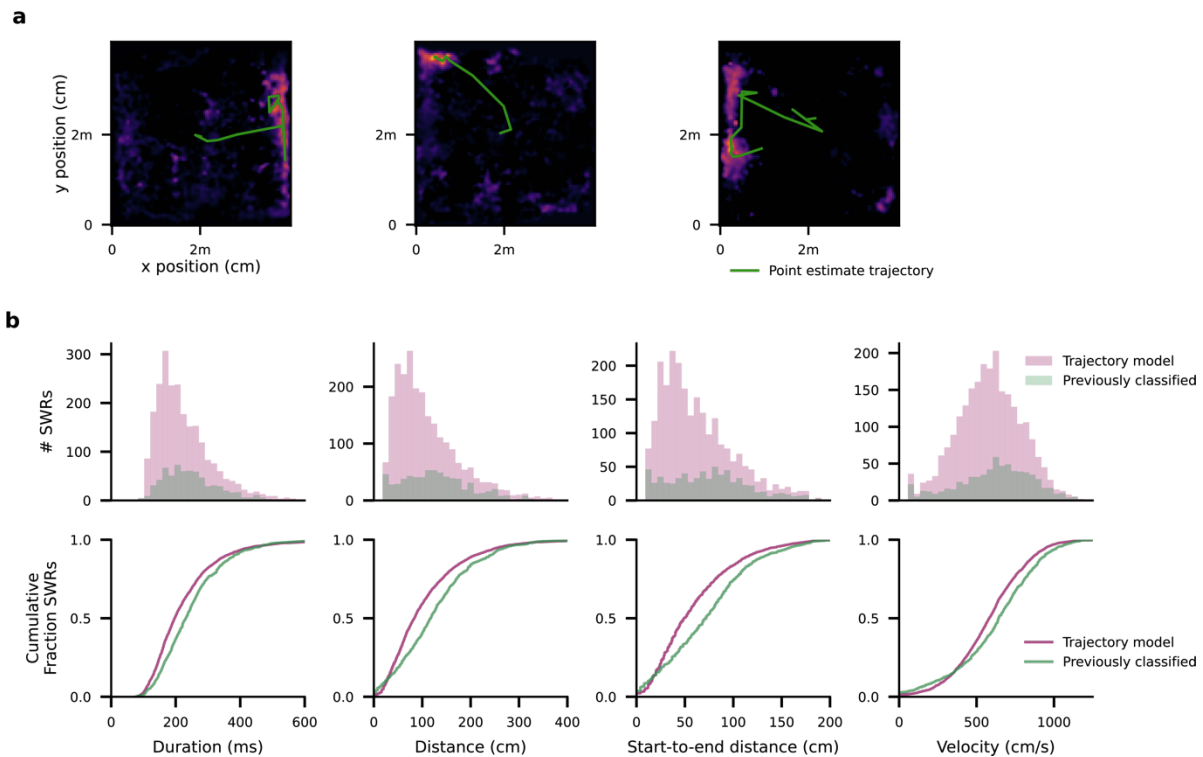


Figure S2. Comparing previously classified SWRs to those classified as trajectories by our approach. Related to Figure 3.

a. Only a small fraction of previously classified SWRs are best described by non-trajectory model. Specifically, our model comparison classified 8.3% of SWRs previously classified as trajectories as non-trajectories (out of 2956 SWRs, 48 as stationary, 2 as stationary Gaussian, and 8 did not meet the population burst threshold criteria (see Methods and Fig. S7)). Three examples of such SWRs are shown here, all classified as stationary. The green traces show the maximum a-posteriori trajectories decoded by the traditional method. All shown trajectories fit the criteria of maximum distance between consecutive points and minimum total distance used by Pfeiffer & Foster (2015). However, these criteria only seem to be satisfied due to noise in the decoded positions, as can be seen from the heatmaps, which visualizes these decoded positions by their posteriors, summed over time (see Fig. 2a).

b. Histogram (top row) and cumulative fraction (bottom row) of replay trajectory descriptive statistics depicting the duration, total distance, start-to-end distance, and velocity (defined as total distance divided by duration) for either the set of all SWRs classified as a trajectory by our model comparison (purple), or the set of previously classified SWRs (green). Previously classified SWRs featured trajectories that were longer in duration (independent t-test; $t(3058)=4.30$ two-sided, Bonferroni-corrected $p=5.34 \times 10^{-5}$), distance (independent t-test; $t(3058)=6.08$ two-sided, Bonferroni-corrected $p < 10^{-6}$), and start-to-end distance (independent t-test; $t(3058)=7.58$ two-sided, Bonferroni-corrected $p < 10^{-6}$), and were faster than the SWRs we identified as encoding trajectories (independent t-test; $t(3058)=4.12$ two-sided, Bonferroni-corrected $p=1.17 \times 10^{-4}$). Previously classified trajectories were those used in Pfeiffer and Foster (2015), who selected trajectories by applying heuristic criteria, such as minimum start to end distance, to the sequence of posterior mean decoded locations. Here, we instead decoded trajectories using the more robust Viterbi algorithm, which causes some of the SWRs previously classified as trajectories to appear to violate these heuristic criteria (e.g., shorter start-to-end distance than allowed).

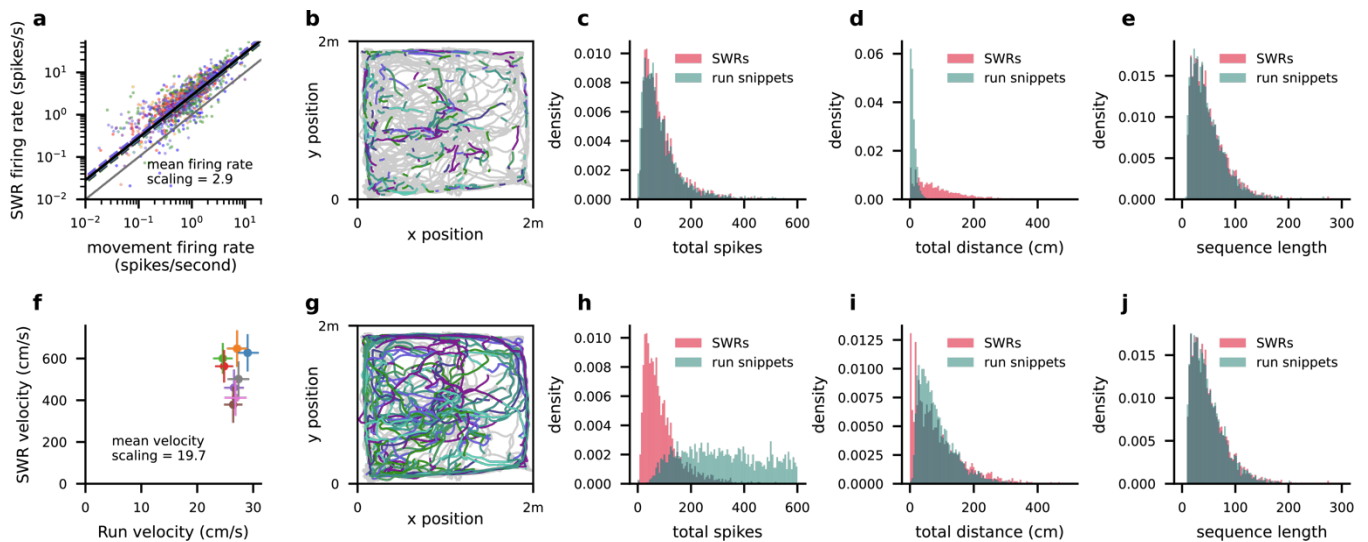


Figure S3. Velocity scaling in SWRs and run snippet selection. Related to Figure 4.

For a fair comparison between SWRs and neural activity during movement, we selected run snippets (randomly selected snippets of neural data from periods in which the rat was moving $> 5\text{cm/s}$) for our model comparison analysis that matched both the distance distribution we observed across SWRs and distribution of sequence lengths (bottom row). If we instead matched the distribution of total number spikes between SWRs and run snippets, the run snippets traverse much shorter distances than the decoded replay trajectories (top row).

a. Log-log plot of the average firing rate during movement, vs. during population bursts. Average firing rate is calculated for each unit separately (points on scatter plot, colored according to session) as the total number of spikes emitted over the total duration within periods that the rat was moving $> 5\text{cm/s}$ for each unit (movement), or within population bursts (Fig. S7). The slope of a linear regression, fitted separately to the data of each session (dashed lines, one per session), appears as an offset in the log-log plot. The solid gray line indicates the expected slope if the firing rates were equal across sessions. The average slope, or average population activity scaling across all sessions, was 2.9 (solid black line). This spike count scaling factor was used to scale the place cell firing rates that were estimated from data while the rats were moving (Fig. 1), for use in the spike generation model of SWR activity.

b. Example run snippets selected for session 1 (run snippets in color, full behavioral trajectory throughout the session in gray). Run snippet durations here were chosen by up-scaling the duration distribution of SWR population bursts by the spike count scaling factor of 2.9 determined in **a**.

c. Distribution of total number of spikes within all SWRs and all run snippets across all sessions.

d. Distribution of total trajectory distance within all decoded SWRs and all run snippets across all sessions.

e. Distribution of sequence lengths (i.e., number of time bins per sequence) within all decoded SWRs and run snippets across all sessions. In order to provide the same sequence lengths to the Bayesian model comparison for run snippets as for SWRs despite the longer durations of the run snippets, we also scaled the time bin size by the spike count scaling factor (3ms for SWRs, 9ms for run snippets).

f. To match the distance distributions rather than the distribution of spike counts, we determined the velocity scaling factor between movement and decoded SWR trajectories. For each session we compared the average velocity of the decoded replay trajectories to the average velocity across all run periods during movement (mean \pm SD for each session). SWR velocities were on average 19.7 times higher than behavioral run velocities.

g-j. Same as **b-e**, but for run snippets in which the duration distribution was calculated by scaling the population bursts duration distribution by the velocity scaling factor, rather than the spike count scaling factor. The time bin size used in **j** was also calculated by scaling the SWR time window by the velocity scaling factor (3ms for SWRs, 60ms for run snippets).

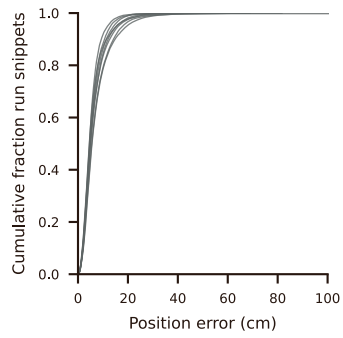


Figure S4. Decoding accuracy. Related to Figure 5.

Cumulative histograms for each session (lines = sessions) of the position error (absolute distance) between animals' true trajectories and those decoded using the Viterbi algorithm applied to the diffusion model, where the absolute distance is calculated at each time point within the trajectory.

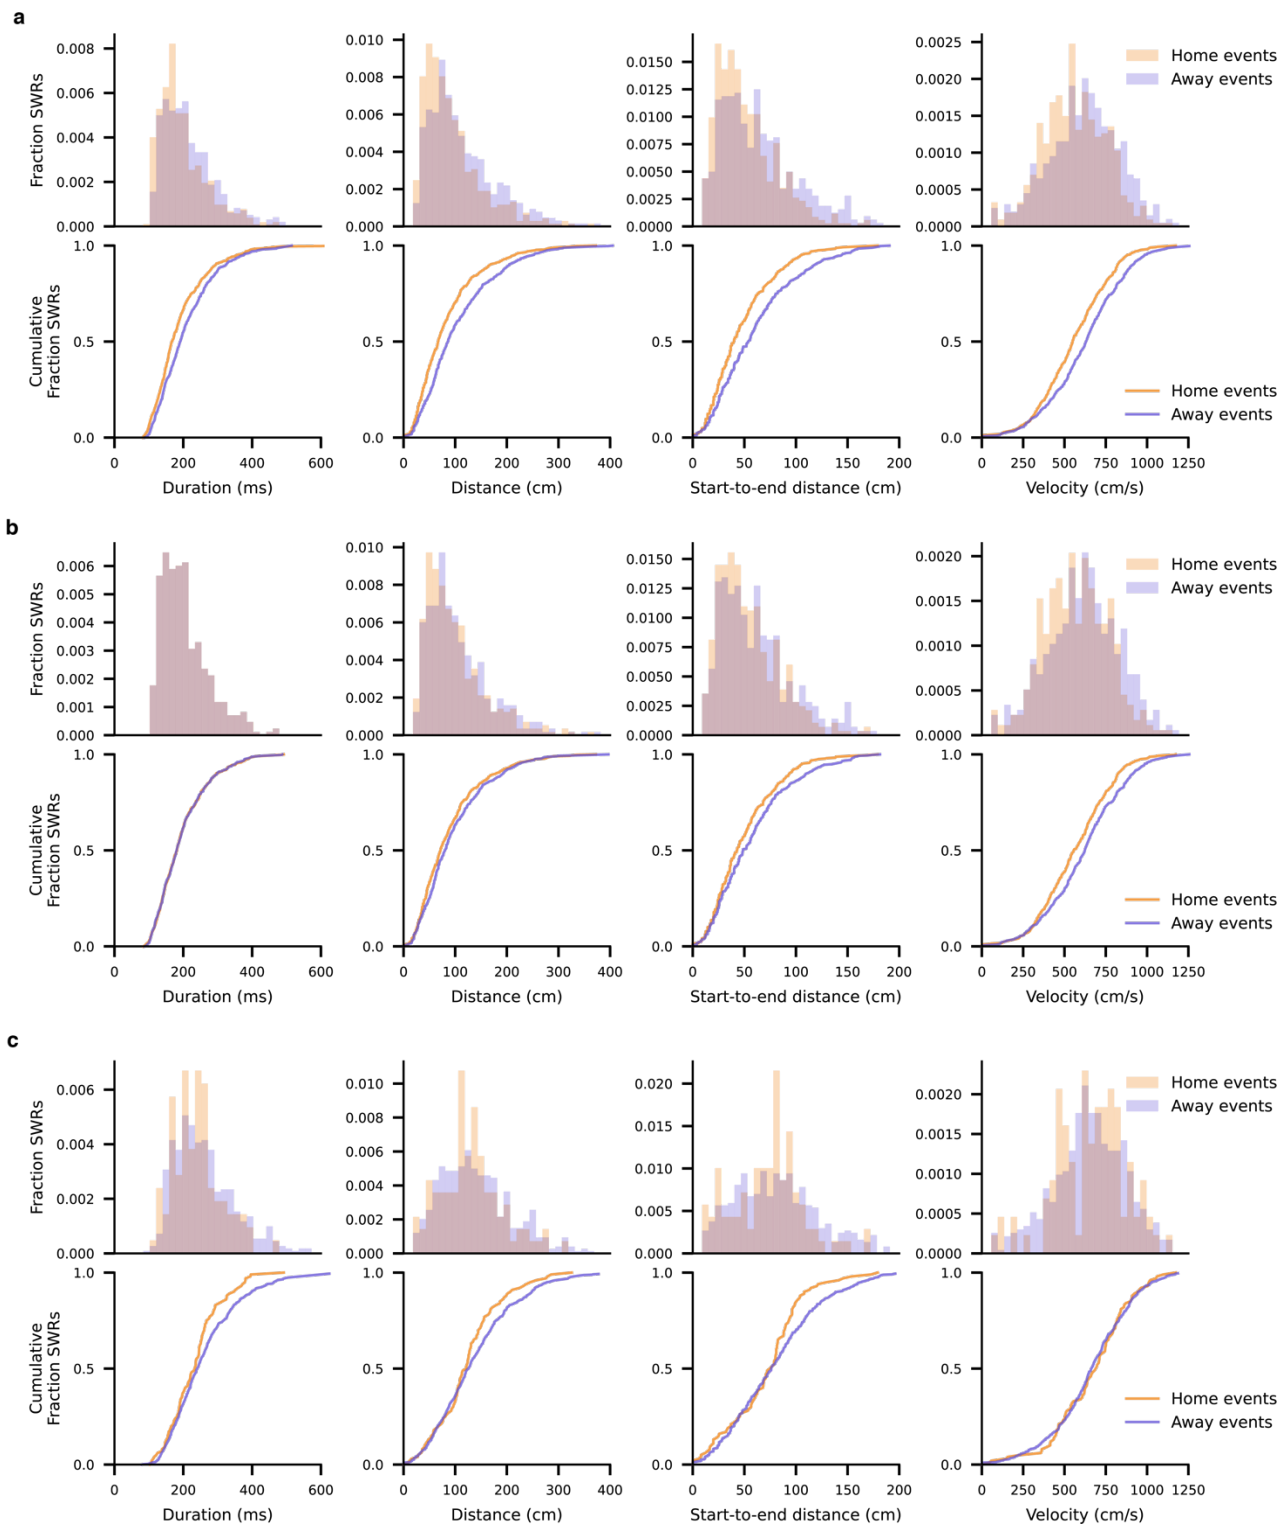


Figure S5. Differences in descriptive trajectory statistics are qualitatively preserved for more restrictive definition of “away” events, even when matching the distribution of trajectory durations, but disappears for previously classified trajectories. Related to Figure 6.

a. We here define “away” events as events initiated while the animal is at the goal location, rather than while the animal is anywhere other than the home location (used in the main text). Compared to Fig. 6 in the main text, we see the same qualitative home/away differences for all descriptive statistics duration, distance, start-to-end distance, and velocity, but with larger effect sizes. Statistical significance by independent t-test, two-

sided, Bonferroni-corrected p-values: duration $t(960)=2.9$, $p=0.011$, total distance $t(960)=4.9$, $p=4.2 \times 10^{-6}$, start-to-end distance $t(960)=5.8$, $p < 1 \times 10^{-6}$, velocity $t(960)=4.6$, $p=1.7 \times 10^{-5}$).

b. These differences persist for this definition of “away” events when matching the duration distributions by subsampling per bin from the trial type (home or away) with more trajectories. In this case, the difference in total distance becomes non-significant, but remains significant for start-to-end distance and velocities.

Statistical significance by independent t-test, two-sided, Bonferroni-corrected p-values: duration $t(820)=0.012$, $p=1.00$, total distance $t(820)=1.69$, $p=0.27$, start-to-end distance $t(820)=3.08$, $p=6.4 \times 10^{-3}$, velocity $t(820)=3.34$, $p=2.6 \times 10^{-3}$).

c. Restricting our analysis to only previously classified SWRs, none of the trajectory statistics differ significantly between home and away events except for duration, which is significant but with a small effect size. Statistical significance by independent t-test, two-sided, Bonferroni-corrected p-values: duration $t(642)=2.47$, $p=0.042$, total distance $t(642)=1.62$, $p=0.32$, start-to-end distance $t(642)=1.89$, $p=0.18$, velocity $t(642)=0.11$, $p=1.00$).

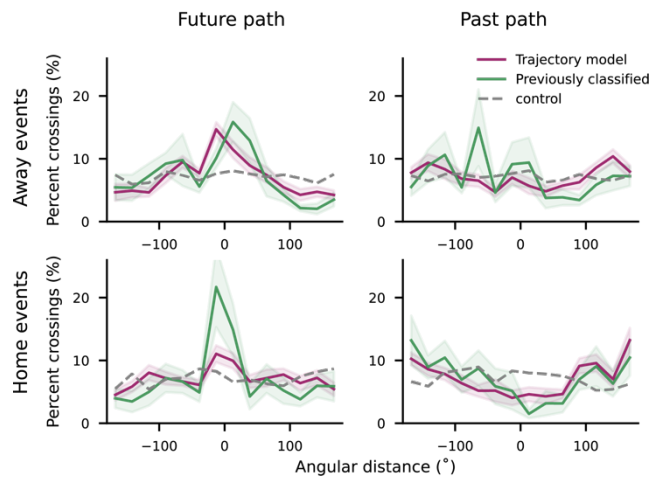


Figure S6. Predictive analysis remains qualitatively similar for more restrictive definition of “away” events. Related to Figure 7.

As in Figure S7a, we here define “away” events as events initiated while the animal is at the goal location, rather than anywhere other than the home location. We see the same results qualitatively for all conditions as shown in Figure 7.

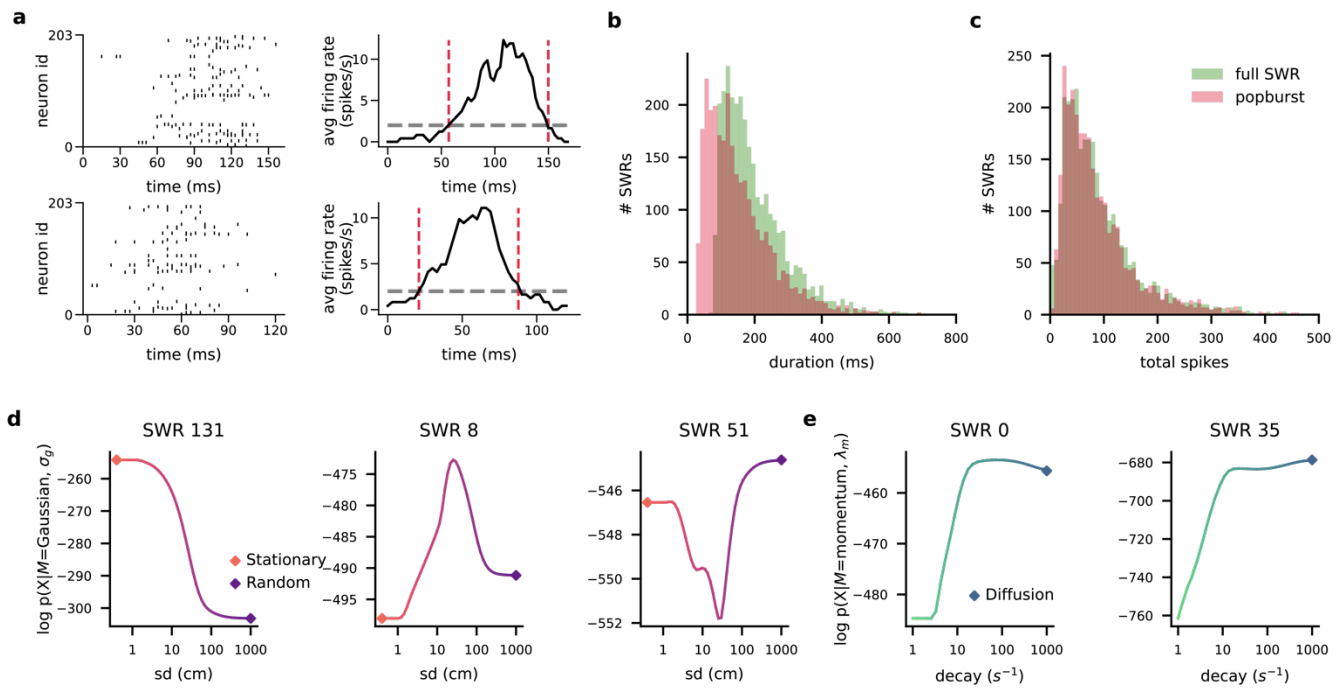


Figure S7. Population burst selection and relation between dynamics models. Related to STAR Methods.

To avoid confounding our location decoding analysis by periods of low neural activity which are uninformative about location, we restricted this analysis to periods (or “population bursts”) in which above a minimal number of spikes was observed.

a. Two example SWRs depicting the process for selecting the population burst within an SWR. Left plots show the spike raster within the LFP-identified SWR. The black trace in the right plots show the average firing rate within the SWR across time, calculated using a 12 ms moving average. Gray dashed line indicates the threshold for the population burst start and end, and the red vertical dashed lines indicate for that ripple the identified start and stop times of the population burst.

b. Histogram of the duration of full SWRs compared to a histogram of the duration of the selected population bursts.

c. Histogram of the total spikes within full SWRs compared to a histogram of the total spikes within the selected population bursts. Despite the decrease in duration from selecting population bursts, almost all spikes are preserved.

d. The Gaussian dynamics model has model evidence equivalent to the stationary/random dynamics models (orange/blue diamonds) once its standard deviation parameter becomes small/large. The three examples depict SWRs that are best described by the stationary model (left), Gaussian model (middle), and random model (right).

e. The momentum dynamics model has model evidence equivalent to the diffusion model (cyan diamond) once the decay parameter becomes large. The two examples depict SWRs that are best described by the momentum model (left), and diffusion model (right). To keep model evidences comparable across models, the model evidences shown in (d) & (e) marginalize over the latent state sequences, $\mathbf{z}_{1:T}$, but not the models’ parameters, θ . Marginalizing over parameters would lower the model evidence for more complex models (e.g., momentum when compared to diffusion, $p(X|M = \text{momentum}) \leq \max_{\theta} p(X|M = \text{momentum}, \theta)$), such that Bayesian model comparison favors simpler models as long as they can explain the data equally well. For example, for SWR 35, marginalizing over the decay parameter results in a model evidence for the momentum model that is lower than that for the diffusion model.