

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

3 Emma L. Krause¹ and Jan Drugowitsch^{1,†}

4 ¹ Department of Neurobiology, Harvard Medical School, Boston, MA, USA

5 [†] Corresponding author: jan_drugowitsch@hms.harvard.edu

6 Abstract

7 During periods of rest, hippocampal place cells feature bursts of activity called sharp-wave ripples (SWRs).
8 Heuristic approaches to their analysis have revealed that a small fraction of SWRs appear to “simulate”
9 trajectories through the environment—called awake hippocampal replay—while the functional role of a
10 majority of these SWRs remains unclear. Applying a novel probabilistic approach to characterize the spatio-
11 temporal dynamics embedded in SWRs, we instead show that almost all SWRs of foraging rodents simulate
12 such trajectories through the environment. Furthermore, these trajectories feature momentum, that is,
13 inertia in their velocities, that mirrors the animals’ natural movement. This stands in contrast to replay
14 events during sleep which seem to follow Brownian motion without such momentum. Lastly, interpreting the
15 replay trajectories in the context of navigational planning revealed that similar past analyses were biased by
16 the heuristic SWR sub-selection. Overall, our approach provides a more complete characterization of the
17 spatio-temporal dynamics within SWRs, highlights qualitative differences between sleep and awake replay,
18 and ought to support future, more detailed, and less biased analysis of the role of awake replay in
19 navigational planning.

20 Introduction

21 Planning through mental simulations, or the anticipation of future action-outcome sequences, is a powerful
22 mechanism for improving action selection (Sutton and Barto, 1998). Strikingly, rodents performing spatial
23 navigation tasks appear to perform such mental simulations (Buzsáki, 2015; Carr et al., 2011). While the
24 animal is moving, hippocampal place cells exhibit spatially localized firing patterns, such that a population of
25 place cells represents the animal's current location in the environment (O'Keefe and Dostrovsky, 1971; Fig.
26 1a/b). During a fraction of sharp-wave ripples (SWRs), which are population bursts of activity associated
27 with brief pauses in the animal's movement, these place cells appear to shift to generating "simulated"
28 trajectories through the environment (Davidson et al., 2009; Diba and Buzsáki, 2007; Tingley and Peyrache,
29 2020; Fig. 1a-c). These "awake hippocampal replay events" have been proposed to support a range of
30 computational functions, such as memory storage (Wilson et al., 1994), recall (Gillespie et al., 2021; Shin et
31 al., 2019; Xu et al., 2019), and planning (Jadhav et al., 2012; Pfeiffer and Foster, 2013; Singer et al., 2013).
32 However, determining the precise computational function of these events has been challenging (Joo and
33 Frank, 2018; Mattar and Daw, 2018), and it remains unclear why they constitute only a small fraction of
34 SWRs (Tingley and Peyrache, 2020). A critical foundation for assessing the computational role of SWRs is
35 a comprehensive and systematic characterization of the spatio-temporal dynamics of the trajectories they
36 encode, and how these dynamics relate to natural movement in the environment.

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

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

67 differ in both their engaged neural mechanisms and their functional role. Fourth, almost all SWRs encoding
68 trajectories implies that previous work analyzing these trajectories might have been biased by discarding
69 the majority of them — as we will demonstrate in the context of navigational planning further below. Lastly,
70 our approach is generally less sensitive to noise than the established heuristics, and thus should lead to
71 cleaner, less biased, decoded trajectories that can inform further work on their role in navigational planning.

72 While previous work has applied probabilistic methods to the analysis of replay events, it has done
73 so in a different context. Location decoding from place cell population activity at isolated time points, for
74 example, commonly relies on maximum likelihood approaches (Johnson and Redish, 2007; Zhang et al.,
75 1998). Other past studies have applied probabilistic methods to identifying replay events, but did not model
76 the spatial dynamics of the trajectories (Linderman et al., 2016; Maboudi et al., 2018). Recent work
77 (Denovellis et al., 2020) used a related probabilistic formulation of the position sequence encoded within an
78 SWR, but restricted itself to W-maze environments, and did not distinguish between diffusion and
79 momentum dynamics. We instead perform a systematic characterization of the spatial dynamics of replay
80 trajectories in unconstrained 2D environments, and use a rich class of considered dynamics models for
81 such environments that yields novel insights into the dynamics underlying SWRs.

82 Results

83 State-space models characterize the spatio-temporal structure of SWRs

84 We analyzed the spatio-temporal structure of SWRs in the dataset of Pfeiffer and Foster (2013,
85 2015), which consists of tetrode recordings from hippocampal CA1 collected while rats foraged around a 2m
86 x 2m open field environment for hidden food reward (Fig. 1a/b). The rats' ability to freely roam an open field,
87 unconstrained by the topology of a maze, was essential to our main objective of characterizing the replay
88 events' spatio-temporal structure. Each session consists of 80-263 single units from recording sessions of
89 37-66 minutes ($n=8$ sessions, 4 rats for 2 sessions each). For our analysis we used the 2956 sharp-wave
90 ripples (SWRs; Fig. 1b) that Pfeiffer and Foster (2015) identified by their ripple-band power of the local-field
91 potential (mean \pm SD = 372 ± 42 SWRs per session, see Methods).

92 We used Bayesian model comparison applied to state-space models to distinguish between different
93 spatio-temporal dynamics encoded by each SWR. Our state-space models assume that the place cell
94 activity sequence underlying each SWR encodes a sequence of positions (i.e., the latent state sequence) in
95 the open field environment, and allow us to distinguish different dynamics of these position sequences from
96 recorded spike data (Fig. 1c). Their two components are (i) a spiking model that determines, for each latent
97 position in the environment, the resulting place cell activity encoding this position, and (ii) a dynamics
98 model, describing the presumed dynamical structure of each sequence of latent positions. We represent the
99 latent position \mathbf{z}_t at each time point t by discretizing the environment into a 50 x 50 grid (4cm x 4cm bins).
100 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
101 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
102 denotes the vector of spike counts emitted by each place cell in the t^{th} time bin. As in previous work
103 (Davidson et al., 2009; Pfeiffer and Foster, 2015), we assume that place cell activity encodes latent position
104 during SWRs as they do during periods of active movement. Thus, we estimated each cell i 's place field $f_i(\mathbf{z})$
105 from its spiking activity during movement, and in turn assumed that the observed spikes during SWRs were,
106 for some latent position \mathbf{z} , generated by draws from a Poisson distribution, independent across cells, with
107 spike rate $f_i(\mathbf{z})$ for place cell i (see Methods).

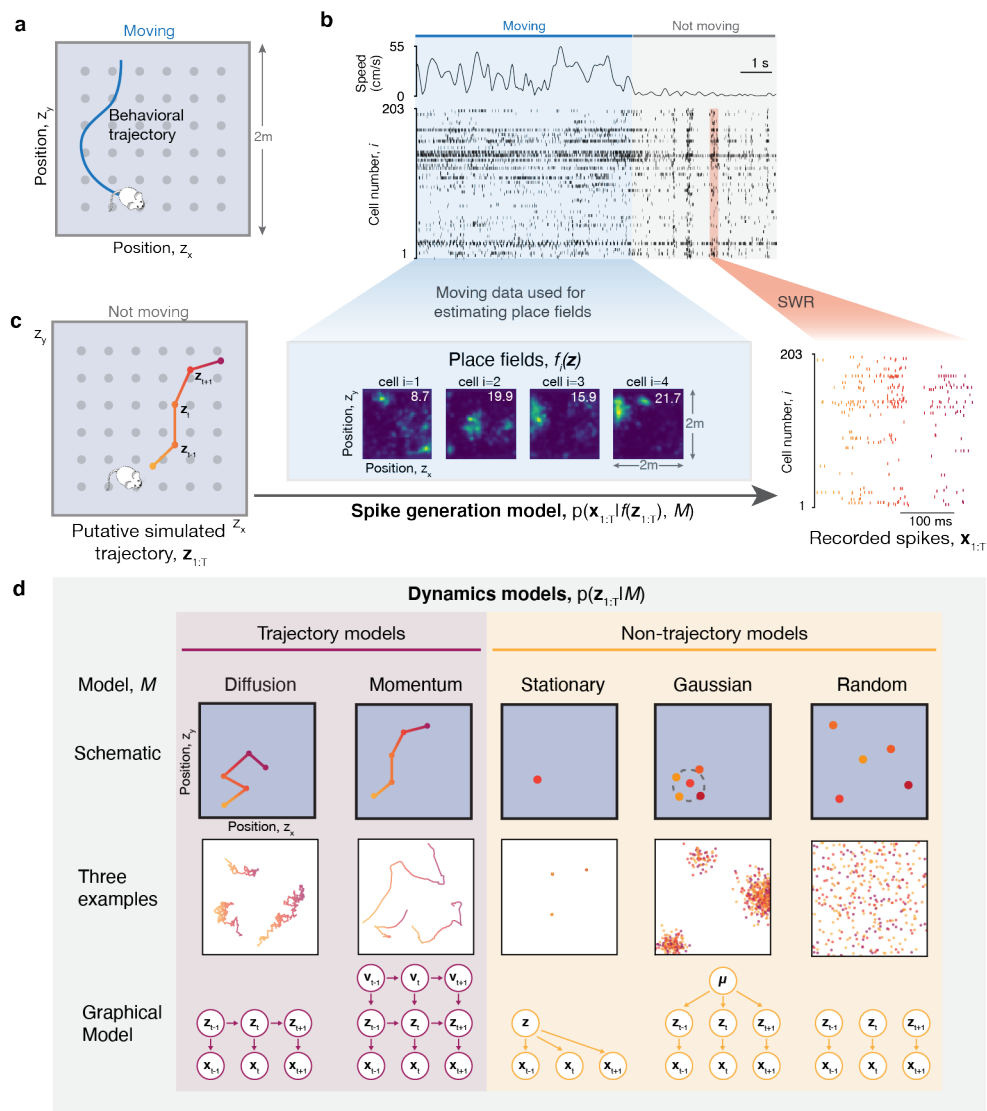


Fig. 1. State-space models for characterizing the spatio-temporal structure of SWRs.

a. Rats foraged in a 2m x 2m 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 the rat 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, $\mathbf{z}_{1:T}$, encoded by the example SWR (left panel, z_x/z_y = encoded x/y position) and the recorded spikes, $\mathbf{x}_{1:T}$ (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 (4 example cells; top-right number = max. 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 spatio-temporal 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 Methods for detailed description).

We then define a set of candidate dynamics models that describe the hypothesized spatio-temporal structure of how the latent position, \mathbf{z}_t , evolves through time and space (Fig. 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

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

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

151 We applied model comparison separately to each SWR (Fig. 2). That is, we did *not* assume all
152 SWRs to encode the same latent state dynamics, but instead asked for each SWR in the dataset separately
153 how likely each model was to have generated the recorded sequence of spikes. To apply the model
154 comparison, we chose to bin observed spikes into non-overlapping time bins of 3ms, but found comparable
155 results for other choices of the time bin size (Fig. S1). Combining the spike generation model with the
156 different dynamics models, and averaging over all possible latent position sequences, allowed us to
157 compute the likelihood $p(\mathbf{x}_{1:T}|M)$ of observing the given spike sequence for each dynamics model (see
158 Methods). Assuming each dynamics model is a-priori equally likely, and applying Bayes' rule to this
159 likelihood, in turn provides a posterior distribution, $p(M|\mathbf{x}_{1:T})$, of how likely each of the dynamics models is for
160 this particular SWR (see Methods). This posterior both tells us which model described the spiking data best
161 — the model with the highest likelihood $p(\mathbf{x}_{1:T}|M)$ — and also gives us a measure of relative certainty across
162 the dynamics models (Fig. 2b). Importantly, this procedure implicitly penalized more complex dynamics
163 models (MacKay, 1995). For example, the latent position sequence arising from a stationary model would
164 also be compatible with a stationary Gaussian model with a small variance. However, the higher complexity
165 of the latter makes our model comparison prefer the former, if compatible with the observed spike
166 sequence.

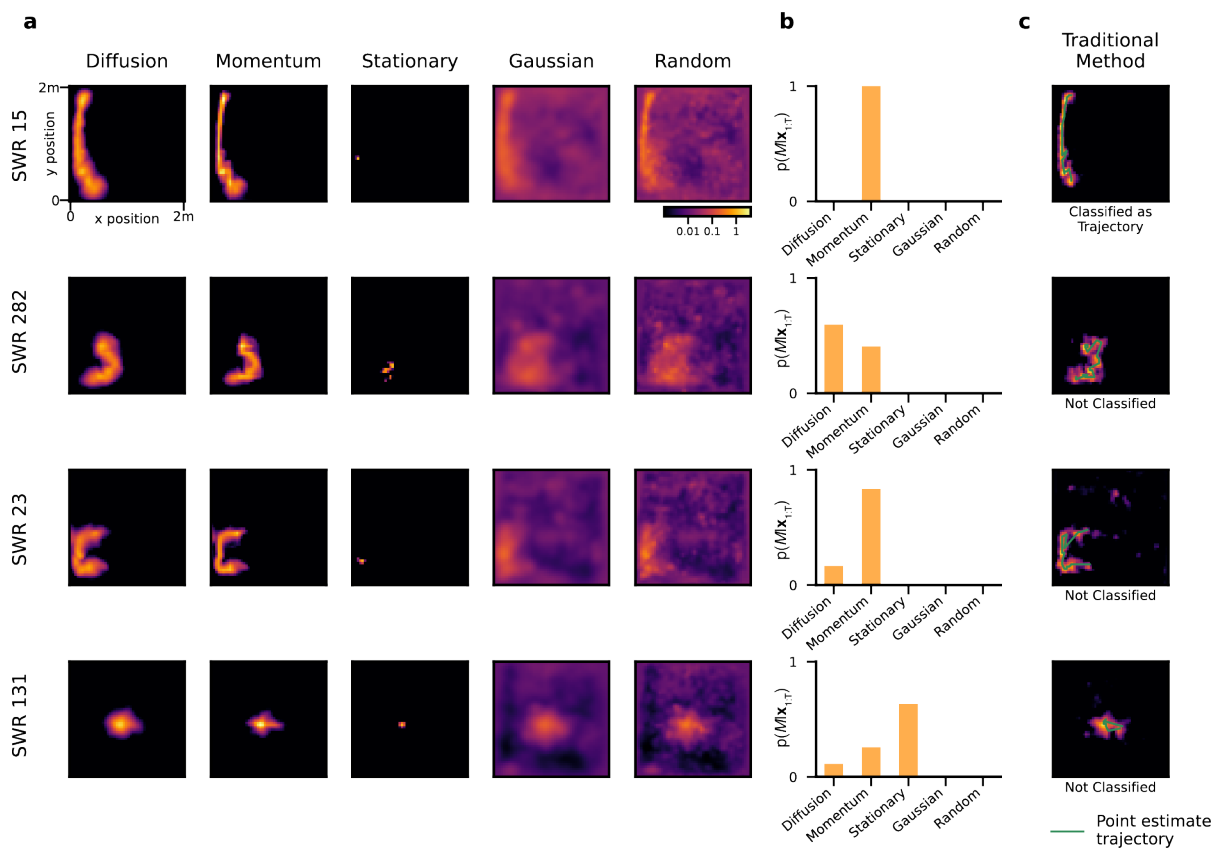


Fig. 2. State-space models reveal the spatio-temporal dynamics within individual SWRs.

The shown examples 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 to 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.

Most awake SWRs feature trajectories with momentum

Applying our model comparison to all SWRs across all sessions confirmed that almost all (92.8%, see Fig. 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 Fig. 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 282 in Fig. 2 are best described by the momentum and diffusion models, respectively, though they failed the previous classification criteria due to a too small start-to-end distance of the extracted trajectory. Indeed, the *traditional method* only classified 23.7% of the 2956 SWRs as trajectories, which we will refer to as *previously classified* trajectories (see Methods). Furthermore, few SWRs were best described by one of the non-trajectory models, as, for example, the stationary model for SWR 131.

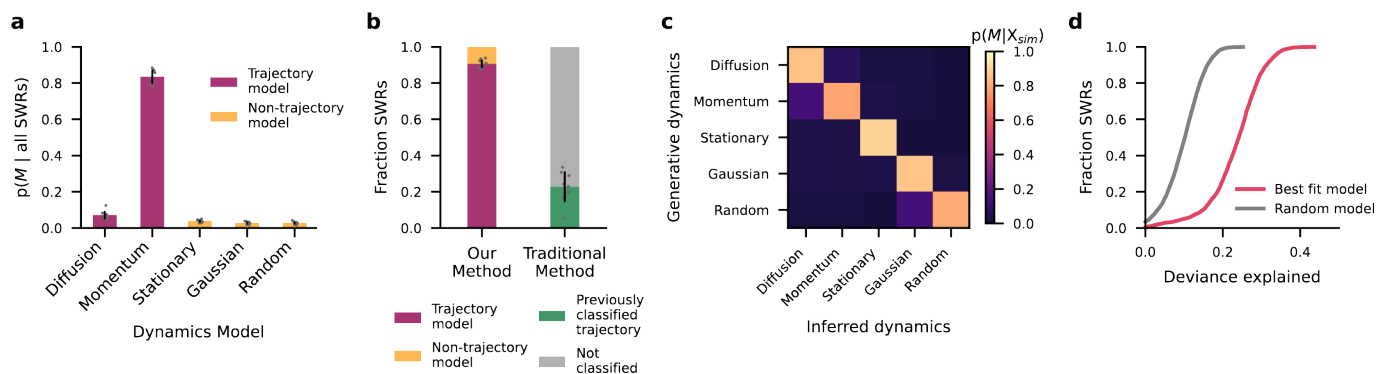


Fig. 3. The vast majority of SWRs feature trajectories with momentum.

a. Inferred distribution of dynamics models underlying the observed SWRs (mean \pm SEM across sessions), computed by random effects model comparison (see main text).

b. Comparison of our method to the traditional method, grouping the diffusion and momentum dynamics models into trajectory models, and the remaining models into non-trajectory models (mean \pm SEM 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, $p(M | \text{all simulated SWRs } X_{\text{sim}})$, 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-fit model (pink line) and the random model (gray line). 0% deviance explained = predicting same avg. spike firing rate across all neurons/SWRs, 100% deviance explained = predicts correct spike count in each time bin.

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; Fig. S3). Across sessions, random effects analysis (Fig. 3a) revealed that $83.6\% \pm 3.8\%$ (mean \pm SEM across sessions) of SWRs are generated by dynamics containing momentum (exceedance probability of momentum model ≈ 1 for all sessions; exceedance probability = probability that 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, $7.1\% \pm 2.4\%$ of SWRs are generated by diffusion dynamics without momentum. The remaining $9.3\% \pm 2.2\%$ are consistent with non-trajectory dynamics (stationary: $3.8\% \pm 0.8\%$, stationary gaussian: $2.8\% \pm 0.8\%$, random: $2.8\% \pm 0.8\%$). Overall, this implies that $90.7\% \pm 2.2\%$ of SWR were identified to encode trajectories, and thus contain temporal structure (Fig. 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 (Fig. 3b).

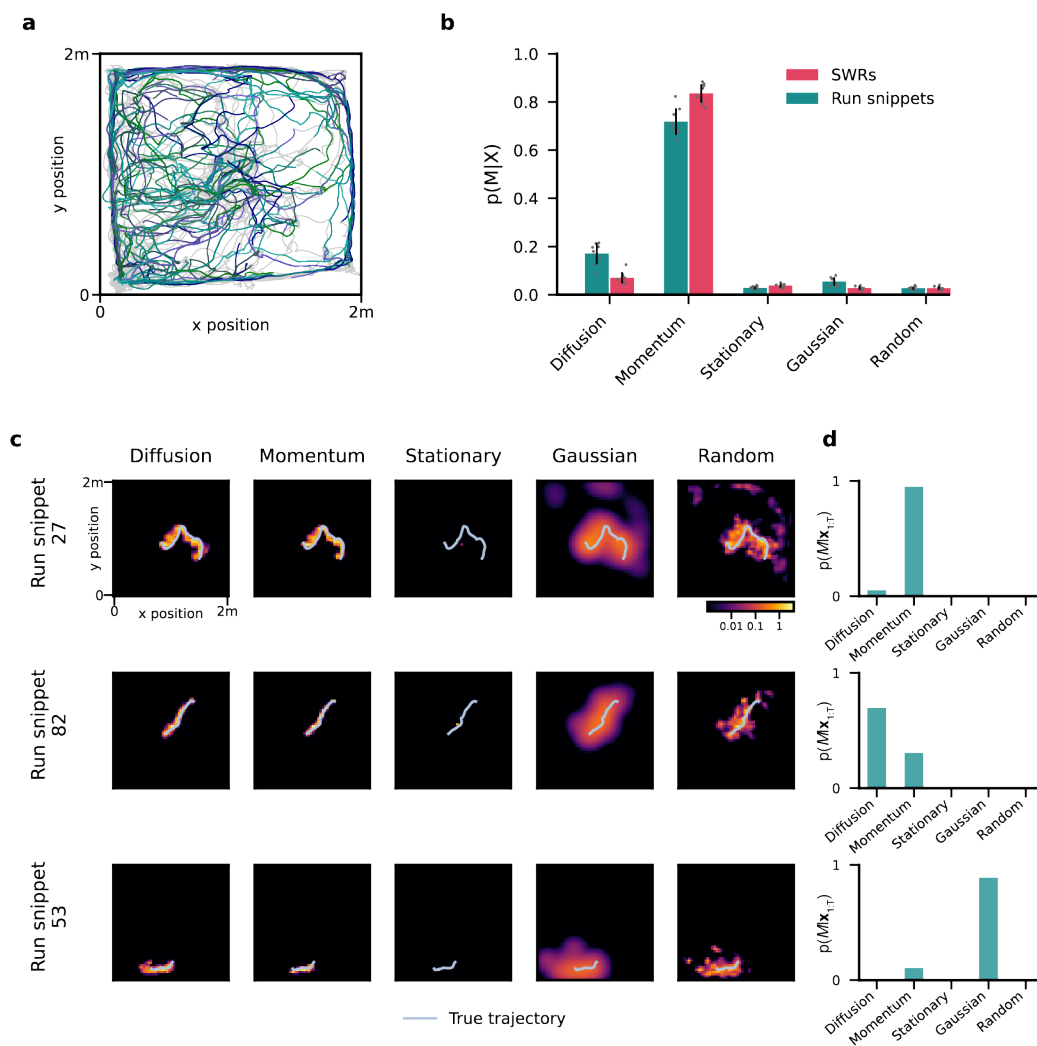
To ensure that our method did not erroneously identify temporal structure where there was none, we generated simulated spiking data under the dynamics of each of our five models and checked how reliably our method could recover the model that we used to generate the data. To match the statistics of true data, we used parameters similar to those recovered from data, and generated spikes using the estimated place fields (see 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 (Fig. 3c). This revealed a reliable discrimination between trajectory and non-trajectory models (F-score=0.93), thus making the spurious detection of temporal structure unlikely. Within trajectory models, SWRs following momentum dynamics are more likely misclassified as diffusion (13.9%) than the reverse (7.4%). 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.

Lastly, we asked how much of the variance in spiking activity our best-fit 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: $0.231 \pm .07$, Fig. 3d), significantly more than the random model (mean \pm SD $0.13 \pm .02$ greater than random model, paired t-test $t(2882)=339.5$, two-sided $p < 1 \times 10^{-6}$). Thus,

235 the identified latent dynamics structured the observed activity in a meaningful way. Taken together, our
 236 findings reveal that almost all SWRs contain dynamics with continuous spatial structure, and consequently
 237 should be included in the analysis and interpretation of the computational role of replay events.

238 Neural activity during SWRs resembles neural activity during behavior

239 If replay events are indeed involved in processing past, or planning for future, behavior, we
 240 speculate the position dynamics they encode to mimic those of real behavior. To test this, we compared the
 241 dynamics inferred from these replay events to dynamics inferred from place cell data of moving animals. We
 242 did so by applying our analysis to randomly selected snippets of place cell activity during periods of
 243 movement (Fig. 4a). For a fair comparison to replay events, we matched the distribution of distances
 244 traversed in these run snippets to the distribution of distances traversed within individual SWRs (see
 245 Methods, Fig. S4).
 246



247 **Fig. 4. Place cell activity during behavior is also best described by momentum dynamics.**

248 **a.** Run snippets selected for model comparison analysis for an example session. The gray trace indicates the animal's behavior
 249 throughout the entire example session. Each colored trace represents a selected run snippet (see Methods and Fig. S4).

250 **b.** Inferred distribution of dynamics models underlying the generation of neural data during behavioral run snippets (teal). The same
 251 distribution is shown for SWRs (pink; same as Fig. 3a) for comparison.

252 **c.** Examples of individual run snippets. Heatmaps as in Fig. 2a, with true trajectory overlaid (light blue line).

253 **d.** The relative likelihood of each model to generate the recorded spikes within the example run snippets shown in **a**.

254
 255
 256 Applying our random effects model comparison to place cell activity during movement yields a very
 257 similar distribution across dynamics models as for SWRs (Fig. 4b): most snippets appeared to feature

258 momentum dynamics, some diffusion dynamics, and few non-trajectory dynamics (mean \pm SD diffusion:
259 17.1% \pm 4.4%, momentum: 71.8% \pm 5.5%, stationary: 2.9% \pm 0.7%, stationary gaussian: 5.5% \pm 1.7%,
260 random: 2.7% \pm 0.7%; example run snippets in Fig. 4c/d). The close match to the dynamics model
261 distribution recovered from SWRs indicates that the statistical structure of spiking activity during SWRs
262 mirrors that during real movement. Given that the rats' actual movement is known to feature momentum
263 (Stella et al., 2019), it might appear surprising that the place cell activity associated with any of the run
264 snippets appeared to feature non-momentum dynamics. However, it only confirms what we have already
265 shown in simulations (Fig. 3c): the ambiguity inherent in limited and noisy spiking data causes our model
266 comparison to mistake trajectories with momentum as only featuring diffusion, and rarely even as non-
267 trajectories. Thus, it is also bound to underestimate the proportion of SWRs with momentum dynamics — in
268 particular for short SWRs with few spikes (Fig. S5). As a consequence, we expect an even larger majority of
269 SWRs than estimated to contain spatial trajectories with momentum.

270 Despite the good match in recovered model distributions across SWRs and run snippets, we
271 observed quantitative differences in the underlying neural activity. SWRs are associated with bursts in
272 population activity, and thus feature more spikes per second across neurons than in similar time periods
273 while the animal is moving (Nádasy et al., 1999; Fig. S6). Even after lengthening movement snippets to
274 match the SWRs spike counts, the trajectories decoded from SWRs traverse further distances than those
275 decoded during movement (Fig. S4). For the aforementioned model comparison, we further increased the
276 movement snippet durations to achieve matching trajectory distance distributions. Overall, SWRs encoded
277 trajectories with more spikes within the same time period, and traversed larger distances with the same
278 number of spikes when compared to place cell activity during movement.

279 Decoded trajectories confirm spatio-temporal momentum dynamics

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

289 To ensure that this discrepancy did not arise from the difference in applied methodology, we
290 replicated our analysis using the method of Stella et al. (2019). First, for each SWR we inferred the most
291 likely trajectory by the Viterbi algorithm. This approach takes temporal continuity of latent positions into
292 account, and is thus more robust to decoding noise than the traditional approach of concatenating the most
293 likely position across individual time bins (Fig. 5a/b, Methods). To avoid confounding our analysis with non-
294 trajectory SWRs, we only considered SWRs that are best described by one of the trajectory models (2600
295 SWRs, 81.5% of total dataset; Fig. S3). Separately, we applied the same approach to infer trajectories from
296 the neural data during movement, using the same run snippets we used for our model comparison analysis
297 further above. The most likely position sequence aligned well with the marginal position estimates for SWRs
298 (Fig. 5a), and with both the rat's behavioral trajectory and the marginal position estimates for run snippets
299 (Fig. 5b; see Fig. S7 for run snippet decoding accuracy).

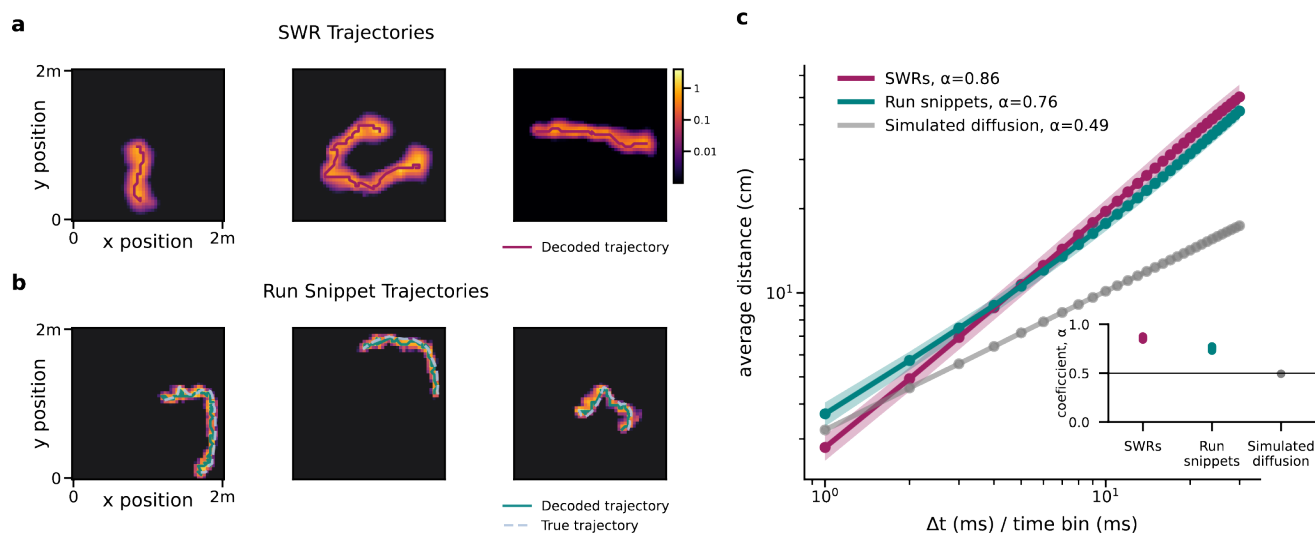


Fig. 5. Trajectories decoded from awake SWRs and behavior are not consistent with Brownian motion.

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

Plotting the distance traveled within individual, inferred trajectories over elapsed time (Fig. 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 $p < 1 \times 10^{-6}$, Fig. 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 $p = 0.80$). 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 "home" well that was consistent across trials within a session (see Fig 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, or finding the "goal" well, whose location is not. Following Pfeiffer and Foster (2013), we group SWRs into "home events" and "away events" based on the current location of the animal — corresponding to being either at the home well, or elsewhere in the environment, respectively. In contrast to Pfeiffer and Foster (2013), who focused on the 23.7% of SWRs classified as replay events by the traditional method, we analyzed the majority of SWRs, namely those classified as best-fit by one of our trajectory models.

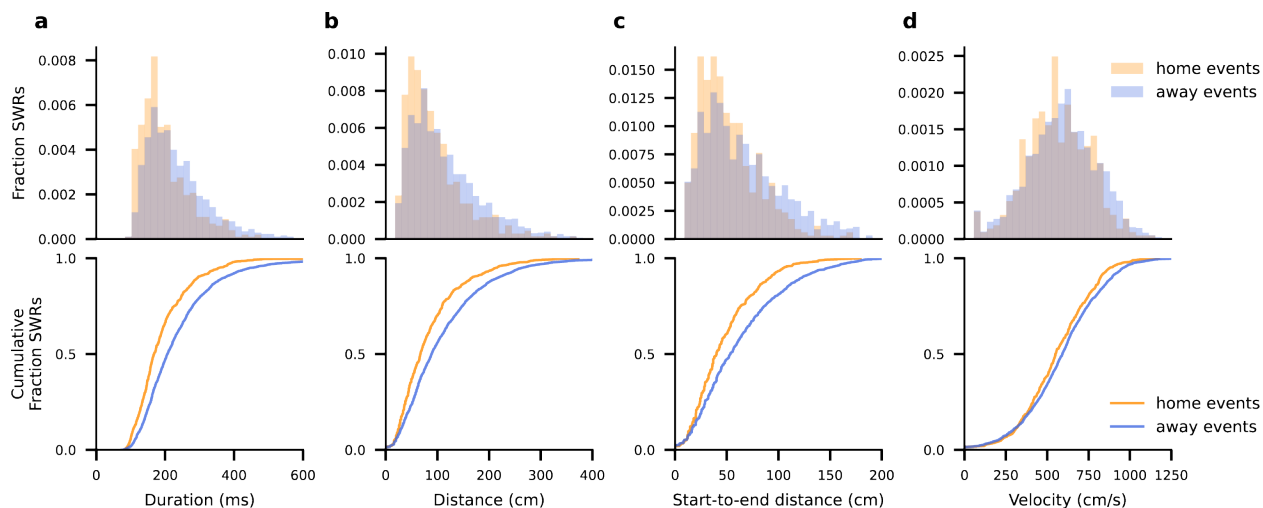


Fig. 6. Decoded trajectories of home events tend to be shorter and slower than of away events.

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.

First, we compared decoded trajectories between these two trial types by 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 (Fig. 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; $t(2352)=7.2$, two-sided, Bonferroni corrected $p<1\times 10^{-6}$), total distance (independent t-test; $t(2352)=6.4$, two-sided, Bonferroni corrected $p<1\times 10^{-6}$), and start-to-end distance (independent t-test; $t(2352)=7.3$, 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; $t(2352)=2.7$, two-sided, Bonferroni corrected $p=0.020$). These differences increased for a more restrictive definition of “away events” that only considered replay events while the animal was at the goal location (Fig. S8a), even after controlling for the difference in trajectory durations (Fig. S8b). However, they vanished for all but duration once we only considered the subset of trajectories extracted with the traditional method for trajectory classification (Fig. S9), and thus have eluded discovery so far. Lastly, the differences were not present in the rodents’ movement (Fig. S10), 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 only considering the subset of previously classified SWRs, while decoding trajectories using our probabilistic model (Fig. 7b, green lines). Once we included all SWRs deemed as encoding trajectories, the effects remained for “away events” but vanished for “home events” (Fig. 7b, purple lines), and similar results were observed for the more restrictive definition of “away events” (Fig. S11).

Though 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.

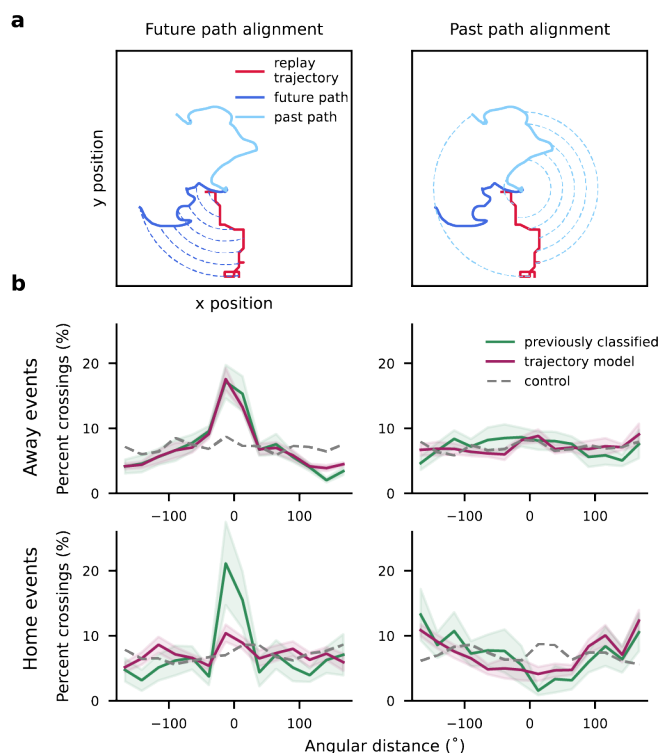


Fig. 7. Away events are more biased to the future path than home events.

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 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 "home events" (top) and "away events" (bottom), and for past (left) and future (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 Methods).

Discussion

Applying a novel, probabilistic method to classify SWRs, we found that the large majority 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 only identified 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.

Though 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 (Yu et al., 2017), 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, 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 (Figs. 4b & 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 Methods/SI). 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.

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 (2019), 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. While 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 (Fig. 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. While 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 spatio-temporal structure.

Treating the identification of latent position dynamics as a probabilistic inference problem has several benefits. First of all, 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, due to 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. Lastly, 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.

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). While data sparsity makes

446 it impossible to precisely quantify such correlations, we do not expect that taking them into account would
447 qualitatively change our results. Furthermore, our model comparison approach is restricted to the five
448 dynamics models we have considered, and trajectories that match neither dynamics model would
449 erroneously be attributed to one of them. Given the high flexibility with which we formulated the dynamics
450 models, we don't expect this to have confounded our results. Nonetheless, an interesting avenue for future
451 work could be to model the latent location sequence as a hierarchical Gaussian process, as in (Wu et al.,
452 2018), in which case structural assumptions are flexibly captured by the covariance matrix governing
453 location evolution. This would allow for even greater flexibility in defining possible transition structures of
454 interest. Lastly, we have focused on identifying the macrodynamics of replay trajectories, which does not
455 capture previously identified fine scale dynamics within replay events in which the position "jumps" between
456 successive positions (Pfeiffer and Foster, 2015).

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

465 Acknowledgements

466 We thank Brad Pfeiffer and David Foster for sharing their data with us. Further, we would like to thank Matt
467 Wilson, Sam Gershman, and members of the Drugowitsch lab for feedback on the work, and Anna
468 Kutschireiter, Christopher Harvey, Johannes Bill, and John Vestola for comments on an early draft of this
469 manuscript. The work was supported by a James S. McDonnell Foundation Scholar Award for
470 Understanding Human Cognition (J.D., grant# 220020462) and a National Defense Science and
471 Engineering Graduate Fellowship (E.K.).
472

473 References

- 474 Averbeck, B.B., and Costa, V.D. (2017). Motivational neural circuits underlying reinforcement learning. *Nat.*
475 *Neurosci.* *20*, 505–512.
- 476 Bishop, C. (2006). *Pattern Recognition and Machine Learning* (New York: Springer-Verlag).
- 477 Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and
478 planning: HIPPOCAMPAL SHARP WAVE-RIPPLE. *Hippocampus* *25*, 1073–1188.
- 479 Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential
480 substrate for memory consolidation and retrieval. *Nat. Neurosci.* *14*, 147–153.
- 481 Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal Replay of Extended Experience.
482 *Neuron* *63*, 497–507.
- 483 Denovellis, E.L., Gillespie, A.K., Coulter, M.E., Sosa, M., Chung, J.E., Eden, U.T., and Frank, L.M. (2020).
484 Hippocampal replay of experience at real-world speeds (*Neuroscience*).
- 485 Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples.
486 *Nat. Neurosci.* *10*, 1241–1242.
- 487 Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells
488 during the awake state. *Nature* *440*, 680–683.
- 489 Gillespie, A.K., Astudillo Maya, D.A., Denovellis, E.L., Liu, D.F., Kastner, D.B., Coulter, M.E., Roumis, D.K.,
490 Eden, U.T., and Frank, L.M. (2021). Hippocampal replay reflects specific past experiences rather than a
491 plan for subsequent choice (*Neuroscience*).
- 492 Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake Hippocampal Sharp-Wave
493 Ripples Support Spatial Memory. *Science* *336*, 1454–1458.
- 494 Johnson, A., and Redish, A.D. (2007). Neural Ensembles in CA3 Transiently Encode Paths Forward of the
495 Animal at a Decision Point. *J. Neurosci.* *27*, 12176–12189.
- 496 Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave–ripple in memory retrieval for immediate
497 use and consolidation. *Nat. Rev. Neurosci.* *19*, 744–757.
- 498 Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. *Nat.*
499 *Neurosci.* *12*, 913–918.
- 500 Kohn, A., Coen-Cagli, R., Kanitscheider, I., and Pouget, A. (2016). Correlations and Neuronal Population
501 Information. *Annu. Rev. Neurosci.* *39*, 237–256.
- 502 Kropff, E., Carmichael, J.E., Moser, M.-B., and Moser, E.I. (2015). Speed cells in the medial entorhinal
503 cortex. *Nature* *523*, 419–424.
- 504 Leutgeb, S. (2004). Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science* *305*, 1295–
505 1298.
- 506 Linderman, S.W., Johnson, M.J., Wilson, M.A., and Chen, Z. (2016). A Bayesian nonparametric approach
507 for uncovering rat hippocampal population codes during spatial navigation. *J. Neurosci. Methods* *263*, 36–
508 47.
- 509 Maboudi, K., Ackermann, E., de Jong, L.W., Pfeiffer, B.E., Foster, D., Diba, K., and Kemere, C. (2018).
510 Uncovering temporal structure in hippocampal output patterns. *ELife* *7*, e34467.

- 511 MacKay, D.J.C. (1995). Information Theory, Inference, and Learning Algorithms.
- 512 Mattar, M.G., and Daw, N.D. (2018). Prioritized memory access explains planning and hippocampal replay.
513 *Nat. Neurosci.* *21*, 1609–1617.
- 514 Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of
515 recurring spike sequences in the hippocampus. *J. Neurosci.* *19*, 9497–9507.
- 516 O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit
517 activity in the freely-moving rat. *Brain Res.* *34*, 171–175.
- 518 Penny, W.D., Stephan, K.E., Daunizeau, J., Rosa, M.J., Friston, K.J., Schofield, T.M., and Leff, A.P. (2010).
519 Comparing Families of Dynamic Causal Models. *PLoS Comput. Biol.* *6*, e1000709.
- 520 Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to
521 remembered goals. *Nature* *497*, 74–79.
- 522 Pfeiffer, B.E., and Foster, D.J. (2015). Autoassociative dynamics in the generation of sequences of
523 hippocampal place cells. *Science* *349*, 180–183.
- 524 Rudnick, J., and Gaspari, G. (2004). Elements of the Random Walk: An introduction for Advanced Students
525 and Researchers (Cambridge University Press).
- 526 Shin, J.D., Tang, W., and Jadhav, S.P. (2019). Dynamics of Awake Hippocampal-Prefrontal Replay for
527 Spatial Learning and Memory-Guided Decision Making. *Neuron* *104*, 1110-1125.e7.
- 528 Singer, A.C., Carr, M.F., Karlsson, M.P., and Frank, L.M. (2013). Hippocampal SWR Activity Predicts
529 Correct Decisions during the Initial Learning of an Alternation Task. *Neuron* *77*, 1163–1173.
- 530 Stella, F., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Hippocampal Reactivation of Random
531 Trajectories Resembling Brownian Diffusion. *Neuron* *102*, 450-461.e7.
- 532 Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., and Friston, K.J. (2009). Bayesian model selection
533 for group studies. *NeuroImage* *46*, 1004–1017.
- 534 Sutton, R.S., and Barto, A.G. (1998). Reinforcement learning: an introduction (Cambridge, Mass: MIT
535 Press).
- 536 Taube, J.S. (2007). The Head Direction Signal: Origins and Sensory-Motor Integration. *Annu. Rev.*
537 *Neurosci.* *30*, 181–207.
- 538 Tingley, D., and Peyrache, A. (2020). On the methods for reactivation and replay analysis. *Philos. Trans. R.*
539 *Soc. B Biol. Sci.* *375*, 20190231.
- 540 Wilson, M.A., McNaughton, B.L., and others (1994). Reactivation of hippocampal ensemble memories
541 during sleep. *Science* *265*, 676–679.
- 542 Wu, A., Roy, N.G., Keeley, S., and Pillow, J.W. (2018). Gaussian process based nonlinear latent structure
543 discovery in multivariate spike train data. *Adv. Neural Inf. Process. Syst.* *30*, 3499–3508.
- 544 Xu, H., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1
545 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze.
546 *Neuron* *101*, 119-132.e4.

547 Yu, J.Y., Kay, K., Liu, D.F., Grossrubatscher, I., Loback, A., Sosa, M., Chung, J.E., Karlsson, M.P., Larkin,
548 M.C., and Frank, L.M. (2017). Distinct hippocampal-cortical memory representations for experiences
549 associated with movement versus immobility. *ELife* 6, e27621.

550 Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population
551 activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.*
552 79, 1017–1044.

553

558 Methods

559 Experimental Details

560 Behavioral Task and Data Acquisition

561 The dataset used in this study has been described in detail previously in Pfeiffer & Foster (2013, 2015). All procedures
562 were approved by the Johns Hopkins University Animal Care and Use Committee and followed US National Institutes of
563 Health animal use guidelines.

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

570 The animal's position was determined by two distinctly colored, head-mounted LEDs, recorded at 60Hz by an over-
571 head video system, and down-sampled to 30Hz before analysis. The animal's speed was calculated as the distance
572 between successive recorded positions over the recording resolution rate, and the data was split in to "moving" and "not
573 moving" periods by a threshold of 5 cm/s on the running speed.

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

578 Estimating Place Fields

579 Place fields were fit using spiking data from periods in which the rat was considered moving (see above). The animal's
580 position, $\mathbf{z} = [z_x, z_y]$, was binned into a 50 x 50 grid (4 cm x 4 cm bins), leading to 2,500 unique positions within the
581 arena. For each cell, i , the corresponding place field $f_i(\mathbf{z})$ was defined as the maximum a-posteriori estimate assuming
582 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
583 spatial bin, $\mathbf{z}_t = k$,

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

584 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
585 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
586 was used with parameter settings $\alpha = 1.01$ and $\beta = 0.01$, such that the maximum a-priori firing rate is $(\alpha - 1)/\beta = 1$
587 spikes/s. After estimation, place fields $f_i(\mathbf{z}_t)$ were smoothed with a Gaussian kernel (4 cm standard deviation). A cell
588 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
589 as place cells were used in further analysis.

590 SWR Detection

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

598 Extracting Population Bursts from SWRs

599 To ensure low firing-rate periods flanking the population activity burst of interest within an SWR did not impact our
600 analysis, we sub-selected a time period of each SWR in which the average firing rate across all neurons was above 2
601 spikes/s per neuron. Specifically, for each SWR, the spiking activity was binned using a 3ms time bin (unless stated
602 otherwise, Fig. S1), and the "population burst" was defined as the period of time from the first upward-crossing to the
603 last downward-crossing of the firing rate threshold within the SWR (Fig. S12a). Population bursts shorter than 30ms
604 were excluded from the analysis (73 SWRs, or 2.5% of the 2956 total SWRs were excluded). This resulted in population

605 bursts that were shorter than the original SWRs, but had roughly the same number of spikes (Fig. S12b/c). Throughout
 606 all further analyses we only analyzed neural data within the extracted population burst, but referred to the extracted
 607 population bursts as SWRs, for simplicity.

608 Description of State-Space Models used to Characterize SWRs

609 All five state-space models that we use to distinguish between the spatio-temporal dynamics of the SWRs consist of two
 610 components: (i) a spike generation model, which is common across all models, and (ii) a dynamics model, which differed
 611 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,
 612 $\mathbf{z}_{1:T} = \mathbf{z}_1 \dots \mathbf{z}_T$ 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
 613 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
 614 stated otherwise, see Fig. S1). The dynamics model, $p(\mathbf{z}_{1:T}|M)$ differs across models, M , and describes the probability
 615 of observing a specific position sequence $\mathbf{z}_{1:T}$ for each of these models. Together, they result in the joint probability
 616 across spikes and latent positions for state-space model M to be given by

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

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

619 Spike Generation Model

620 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,
 621 that is $p(\mathbf{x}_{1:T}|\mathbf{z}_{1:T}) = \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{z}_t)$. Within each time bin, spikes are assumed to be drawn from a Poisson distribution,
 622 independent across cells, with rates determined by the place fields fit during movement (see Estimating Place Cells),
 623 that is

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

624 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
 625 factor (described below). This is equivalent to assuming that the spike sequence of each cell is drawn according to a
 626 time-discretized inhomogeneous Poisson process with instantaneous spike rate $f_i(\mathbf{z}_t)\gamma_{\text{spike}}$ that varies across positions
 627 \mathbf{z}_t , and thus time.

628 We accounted for the increase in firing rate within SWRs as compared to movement by calculating a spike count
 629 scaling factor, γ_{spike} (as described in Fig. S6). Average firing rate was calculated for each unit as the total number of
 630 spikes emitted over the total duration within periods that the rat was either considered “moving”, $\tilde{f}_{i,\text{movement}}$, or within
 631 population bursts, $\tilde{f}_{i,\text{popburst}}$ for each cell i . A linear regression was performed over the set of $(\tilde{f}_{i,\text{movement}}, \tilde{f}_{i,\text{popburst}})$
 632 pairs per session, and the average slope, or average population activity scaling, across all sessions, $\gamma_{\text{spike}} = 2.9$, was
 633 used as the spike count scaling factor.

634 Dynamics Models

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

$$p(\mathbf{z}_{1:T}|M) = p(\mathbf{z}_1|M) \prod_{t=2}^T p(\mathbf{z}_t|\mathbf{z}_{1:t-1}, M). \quad (4)$$

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 (Fig. 1d, bottom row). In continuous time, the dynamics of latent position \mathbf{z}_j in the diffusion and momentum models is described by

$$\text{Diffusion:} \quad \frac{dz_j}{dt} = \sigma_d \eta_{j,d}(t), \quad (5)$$

$$\text{Momentum:} \quad \frac{dz_j}{dt} = v_j(t), \quad \frac{dv_j}{dt} = -\lambda_m v_j(t) + \sigma_m \eta_{j,m}(t), \quad (6)$$

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

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

$$\text{Diffusion:} \quad p(\mathbf{z}_t | \mathbf{z}_{t-1}) = \mathcal{N}(\mathbf{z}_t | \mathbf{z}_{t-1}, \sigma_d^2 \delta t \mathbf{I}), \quad (7)$$

$$\text{Momentum:} \quad p(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{z}_{t-2}) = \mathcal{N}\left(\mathbf{z}_t | (1 + e^{-\lambda_m \delta t}) \mathbf{z}_{t-1} - e^{-\lambda_m \delta t} \mathbf{z}_{t-2}, \frac{\sigma_m^2 \delta t^2}{2\lambda_m} (1 - e^{-2\lambda_m \delta t}) \mathbf{I}\right), \quad (8)$$

$$p(\mathbf{z}_2 | \mathbf{z}_1) = \mathcal{N}(\mathbf{z}_2 | \mathbf{z}_1, \sigma_{m0}^2 \delta t \mathbf{I}). \quad (9)$$

640 In the above, all normal distributions $\mathcal{N}(\cdot)$ are bivariate normal distributions over both spatial dimensions with diagonal
 641 covariance matrices, and discretized and appropriately normalized over the spatial grid. The detailed derivation for the
 642 momentum model is provided in Supplementary Information. It results in a second-order Markov model, where the first
 643 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
 644 initial velocity, $v_j(0)$.

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

$$\text{Stationary:} \quad p(\mathbf{z}_t | \mathbf{z}_{1:t-1}) = p(\mathbf{z}_t) = \delta_{\mathbf{z}_t, \mathbf{z}_s}, \quad (10)$$

$$\text{Gaussian:} \quad p(\mathbf{z}_t | \mathbf{z}_{1:t-1}) = p(\mathbf{z}_t) = \mathcal{N}(\mathbf{z}_t | \boldsymbol{\mu}, \sigma_g^2 \mathbf{I}), \quad (11)$$

$$\text{Random:} \quad p(\mathbf{z}_t | \mathbf{z}_{1:t-1}) = p(\mathbf{z}_t) = \mathcal{U}(\mathbf{z}_t), \quad (12)$$

645 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
 646 normal distribution $\mathcal{N}(\mathbf{z}_t | \boldsymbol{\mu}, \sigma_g^2 \mathbf{I})$ is, as before, a bivariate normal distribution over both spatial dimensions with a diagonal
 647 covariance matrix, and $\mathcal{U}(\mathbf{z}_t) = 1/K$ is the uniform distribution across all $K = 2,500$ spatial bins. The non-trajectory
 648 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,
 649 respectively. The non-trajectory and trajectory models relate to each other through their parameters: the standard
 650 deviation, σ_g varies the Gaussian model from the stationary model to the random model, and the momentum model with
 651 a large decay, λ_m , becomes equivalent to the diffusion model.

652 The model parameters $\sigma_d, \sigma_m, \lambda_m, \sigma_g, \mu_g$, and \mathbf{z}_s were fit over a grid and marginalized out for model comparison (see
 653 next section for details). σ_{m0} in the momentum model, which is the standard deviation of the diffusion process on the first
 654 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
 655 the trajectory. For all models a uniform prior was used for the initial position, $p(\mathbf{z}_1 | M) = \mathcal{U}(\mathbf{z}_1) = 1/K$, where $K = 2,500$
 656 is the number of spatial bins.

657 Marginalizing over Latent Position Sequences and Model Parameters

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

$$p(\mathbf{x}_{1:T} | M, \theta_M) = \sum_{\mathbf{z}_{1:T}} p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T} | M, \theta_M) = \sum_{\mathbf{z}_{1:T}} p(\mathbf{x}_1 | \mathbf{z}_1) p(\mathbf{z}_1 | M) \prod_{t=2}^T p(\mathbf{x}_t | \mathbf{z}_t) p(\mathbf{z}_t | \mathbf{z}_{1:t-1}, M, \theta_M), \quad (13)$$

660 where θ_M are the parameters of model M . This provides a relative comparison of how likely each model is to have
 661 generated the spiking data per SWR, for a fixed set of model parameters θ_M . For both the diffusion and momentum
 662 models, we used the forward pass of the forward-backward algorithm to calculate the data likelihood (Bishop, 2006).

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

$$p(\mathbf{x}_{1:T} | M) = \sum_{\theta_M} p(\mathbf{x}_{1:T} | M, \theta_M) p(\theta_M | M). \quad (14)$$

665 For $\boldsymbol{\mu}$ and \mathbf{z}_s we used a uniform grid across all spatial locations. For the remaining parameters, $\sigma_d, \sigma_m, \lambda_m$, and σ_g , we
 666 determined the prior in two steps. First, we evaluated the data likelihood, $p(\mathbf{x}_{1:T} | M, \theta_M)$ over a uniform grid in log-space:
 667 30 bins from 2 to 100cm for σ_d , 30 bins from 40 to 400m/s for σ_m , 10 bins from 1 to 4000 for λ_m , and 30 bins from 6 to
 668 200cm for σ_g . We chose this grid to ensure that most mass of the parameter likelihood $p(\mathbf{x}_{1:T} | M, \theta_M)$ lies within this grid.
 669 Second, we found for each session and relevant model M the maximum likelihood parameter fits over the parameter
 670 grid, $\hat{\theta}_{M,ML} = \operatorname{argmax}_{\theta_M} p(\mathbf{x}_{1:T} | M, \theta_M)$, for each SWR, and used the distribution of $\hat{\theta}_{M,ML}$'s across SWRs to fit an
 671 appropriate prior, $p(\theta_M | M)$. Specifically, for the standard deviation parameters in the diffusion, and Gaussian models, σ_g

672 and σ_d , we used an inverse Gamma distribution as the prior. For the Gaussian model, the mean μ was marginalized out
673 before performing these fits. For the standard deviation and decay parameter in the momentum model, σ_m and λ_m , we
674 fit a multivariate log-normal distribution as the prior. As mentioned before, we assumed a uniform prior over z_s for the
675 stationary model, such that no prior fitting was required.

676 Computing Fraction Deviance Explained

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

$$D_{model} = 2(\log p(\mathbf{x}_{1:T}|M_{sat}) - \log p(\mathbf{x}_{1:T}|M)) \quad (15)$$

$$D_{null} = 2(\log p(\mathbf{x}_{1:T}|M_{sat}) - \log p(\mathbf{x}_{1:T}|M_{null})) \quad (16)$$

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

683 Best-Fitting Models, and Fixed and Random Effects Model Comparison of Dynamics Models

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

689 Specifically, random effects model comparison assumes a prior over models $p(M|\alpha) = r_M$, where r_M is the proba-
690 bility of an SWR being generated according to dynamics model M . Across models, the r_M 's form the vector r which has
691 a Dirichlet prior $\text{Dir}(r|\alpha)$ with concentration parameters α . We set α to get a weak uniform distribution across models,
692 equivalent to about 5% of data points (each element of α is set to 15 for analysis of real data, and 5 for simulated data).
693 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
694 $\mathbf{x}_{1:T_n}^{(n)}$ is the spiking activity within SWR n and N is the total number of SWRs in the session, we compute the posterior
695 $p(M|X)$ using Gibbs sampling, as described in Penny et al. (2010).

696 For comparison, we also calculate the inferred distribution of models using the more standard fixed effects analysis,
697 which assumes each SWR is drawn from the same model. Here the likelihood of the spiking data across SWRs within
698 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
699 prior on M to obtain the posterior distribution over models, $p(M|X)$.

700 For per-SWR analyses, each SWR was “classified” as the model M that had the highest (marginalized) log-likelihood,
701 $\log p(\mathbf{x}_{1:T}|M)$. The fraction trajectory (non-trajectory) models is the fraction SWRs best described by the diffusion or
702 momentum (stationary, Gaussian, or random) models.

703 Model Recovery from Simulated Data

704 For each model M , we simulated 100 position sequences in a 2D continuous space (2m x 2m) with a time bin size of
705 $\delta t = 1\text{ms}$, according to the dynamics of the respective model. For models with parameters, we drew these parameters for
706 each trajectory from the prior, $p(\theta_M|M)$, fitted as described above. For all models, we generated neural activity from the
707 simulated trajectory by binning the trajectory into the same 50 x 50 grid (with 4cm x 4cm bins) used for estimating place
708 fields, and, for each place cell, drawing their activity within each time bin from a Poisson distribution with a rate given
709 by the cell's place field for the specific position in that time bin (Eq. 3) with γ_{spike} set to 2.9 (Fig. S6). For each set of
710 trajectories generated for a fixed model M , we applied the same Bayesian model comparison and random effects analysis
711 to the simulated neural data, X_{sim} , as we describe above for the real neural data, resulting in an inferred distribution
712 across models $p(M|X_{sim})$. Even though we simulated data in time steps of 1ms, for model recovery we used time bins
713 of $\delta t = 3\text{ms}$, as for the real neural data. F-score (computed as the harmonic mean of the precision, TP/(TP+FP), and
714 sensitivity, TP/(TP+FN), where TP=true positive rate, FP=false positive rate, and FN=false negative rate) was used to
715 quantify our model recovery results. To quantify the reliability of distinguishing trajectory vs. non-trajectory dynamics, we
716 assumed true/false corresponded to trajectory/non-trajectory, respectively, and to quantify the reliability of distinguishing
717 momentum from diffusion dynamics we assumed that true/false corresponded to momentum/diffusion, respectively.

718 Extracting Maximum Likelihood Trajectories

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

$$\hat{z}_{1:T,ML} = \arg \max_{z_{1:t}} [p(\mathbf{x}_{1:T}, z_{1:T})] = \arg \max_{z_{1:t}} [p(\mathbf{x}_{1:T}|z_{1:T})p(z_{1:T}|M)] \quad (17)$$

721 using the max-sum algorithm applied to our diffusion model, $p(z_{1:T}|M = \text{diffusion})$ and spike generation model $p(\mathbf{x}_{1:T}|z_{1:T})$.
722 We chose the diffusion model to take into account temporal continuity, but not bias the trajectories to contain momen-
723 tum dynamics. The Viterbi algorithm is preferred to simply taking the maximum posterior estimate at each time bin,
724 $\hat{z}_{t,ML} = \arg \max_z p(z_t|\mathbf{x}_{1:T})$, as the sequence of individually most likely positions does not imply that their sequence is
725 the most likely.

726 Traditional Method for Classifying SWRs

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

736 Analysis of Place Cell Activity During Movement

737 To evaluate how the spatio-temporal dynamics of neural activity during SWRs compare to the spatio-temporal activity of
738 place cells during movement, we applied our Bayesian model comparison analysis described above to snippets of neural
739 activity during movement (Fig. 4). Run snippets were selected to approximately match the distribution of total distance
740 traveled within the trajectories decoded from SWRs (see Extracting Maximum Likelihood Trajectories and Fig. S4).
741 Specifically, we first determined a velocity scaling factor, γ_v that represented, on average, how much faster trajectories
742 within SWRs evolved in comparison to real movement of the animal. For each session, we calculated the velocity of each
743 replay trajectory (total distance/duration) and each run period (a continuous period of time within a session in which the
744 animal was considered “moving” for at least 2s). We calculated the scaling factor between the mean replay trajectory
745 velocity and mean run period velocity for each session, and defined the velocity scaling factor γ_v as the average of the
746 scaling factors across all sessions. This procedure found that the velocity of replay trajectories was on average 19.7
747 times higher than real movement.

748 Next, we selected the run snippets. For each session, the duration distribution of run snippets was determined by
749 up-scaling the duration of SWRs by the velocity scaling factor. We selected run snippets by randomly sampling a run
750 period and then randomly sampling a start time within that run period, such that the entire duration of the run snippet fell
751 within the run period. In order to match the distribution of sequence lengths T between SWRs and run snippets, we also
752 up-scaled the time bin size used for run snippet analysis by the velocity scaling factor, to $\delta t = 60\text{ms}$ (Fig. S4). We then
753 applied the model comparison analysis to the neural data within the set of run snippets exactly as outlined above, except
754 that in Eq. (3) we set $\gamma_{spike} = 1$ to reflect the fact that neural activity came from the same “movement” periods that were
755 used to estimate place fields.

756 Diffusion Coefficient Analysis

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

$$\sqrt{\langle d(t)^2 \rangle} \propto t^\alpha, \quad (18)$$

759 where d is the distance between two points in time within a replay trajectory, and t is the time elapsed between those
760 two time points. An α -value of 0.5 corresponds to Brownian motion, or a random walk, while deviations from 0.5 are
761 inconsistent with Brownian motion. We evaluate if replay trajectories evolve according to Brownian motion by plotting
762 the time-distance relationship in log-space, and use linear regression applied to this log-log plot to find the coefficient α ,
763 which corresponds to the slope of the log-log time-distance relationship. To avoid confounds, we restricted this analysis

764 to SWRs best-fit by a trajectory model (either the diffusion or momentum model). Replay trajectories were decoded from
765 neural activity as described in “Extracting maximum likelihood trajectories”. For each replay trajectory within a session
766 we found the squared distance d_j^2 between all pairs of decoded positions separated by all multiples of the time bin used
767 for trajectory decoding $\Delta t/\delta t$, where Δt is the time elapsed between decoded positions and δt is the decoding time bin.
768 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
769 session. We then found the mean distance, \bar{d}^2 for each multiple of the decoding time bin, fit a linear regression model per
770 session to $(\log(\sqrt{\bar{d}^2}), \log(\Delta t/\delta t))$, which gave the estimate of coefficient α for that session as the slope of the regression.

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

779 Splitting SWRs into “Home” and “Away” Events

780 SWRs were split into “home” and “away” events according to the current location of the animal at SWR onset: a “home”
781 event if the animal is less than 5cm from the home well, and an “away” event if the animal is elsewhere in the environment.
782 In Fig. S8, and Fig. S11 “away” event is redefined as if the animal is less than 5cm from the goal well, rather than
783 anywhere in the environment other than within 5cm of the home well.

784 Analysis of Correlation Between Replay Trajectories and Behavior

785 To quantify the correlation between replay trajectories and behavior, we implemented the method described in Pfeiffer
786 and Foster (2013). Replay trajectories were extracted from SWRs as described in “Extracting maximum likelihood tra-
787 jectories”. The future/past path of the animal was defined as the immediate future/past path of the animal for either
788 10s after/before the SWR or until a direct distance of 75cm from the current location of the animal during the SWR was
789 reached, whichever threshold gives the greater direct distance. To calculate the correlation between the replay trajectory
790 and the behavioral path, the angular distance between the replay trajectory and behavioral trajectory was found at a se-
791 ries of progressively larger radii from the current location of the animal during the SWR. Specifically, the crossing points
792 between the behavioral trajectory and the replay trajectory with a circle centered on current location of the animal was
793 found for each radius, and the angular distance between these two crossing points was calculated (as visualized in Fig.
794 7a). The minimum radius was 5 cm, increased in increments of 3 cm until either 75 cm or the end of the replay trajectory
795 was reached. For each session, a histogram of angular displacements was calculated using all crossings across SWRs
796 within the session, with SWRs split into “home” and “away” events as described in “Splitting SWRs into “home” and
797 “away” events”. This analysis was applied separately to either all SWRs classified as a trajectory model or all SWRs
798 previously classified by the traditional method. Chance correlation was found by 2000 shuffles in which a behavioral path
799 and SWR was selected at random from any session, the behavioral path was shifted to the current location of the animal
800 during the SWR, and the angular distance was calculated using the same method described above.

801 Figure Details

802 **Figure 1.** Panel **b** shading visualizes “moving” and “not moving” periods as described in Behavioral Task and Data
803 Acquisition, and the example place fields in panel **c** were computed using the place field data from “moving” periods
804 as described in Estimating Place Fields. The spike generation model in panel **c** and the dynamics models in panel **d**
805 are described in the Spike Generation Model and Dynamics Models sections. The graphical models in **d** describe the
806 statistical relationship between variables in each dynamics model, whose details are given in Dynamics Models.

807 **Figure 2.** The heatmaps in panel **a** were computed from the marginal decoded position per time bin $p(z_t|\mathbf{x}_{1:T}, M)$ for
808 each dynamics model described in Description of State-Space Models. For the diffusion and momentum models, the
809 marginal decoded positions were found by the forward-backward algorithm (Bishop, 2006). The other, non-trajectory,
810 models had a simpler form that supported a simpler computation of the position estimates. We normalized these posi-
811 tion estimate separately for each time bin, $p(\tilde{z}_t|\mathbf{x}_{1:T}, M) = p(z_t|\mathbf{x}_{1:T}, M) / \sum_{z_t} p(z_t|\mathbf{x}_{1:T}, M)$, summed them over time
812 $p(\tilde{z}_{1:T}|\mathbf{x}_{1:T}, M) = \sum_{t=1}^T p(\tilde{z}_t|\mathbf{x}_{1:T}, M)$, and plotted this sum. While the sum is no longer a real probability distribution,
813 it is nonetheless useful for visualization. In panel **b**, the distribution over models per SWR was computed as described

814 in Marginalizing over Latent Position Sequences and Model Parameters. In panel **c**, the heatmaps were computed by
815 obtaining the decoded position per time bin $p(z_t|x_t)$ as calculated by the traditional method described in Traditional
816 Method for Classifying SWRs, then normalized, summed, and plotted as described for panel **a**. The trajectory plotted is
817 as described in Traditional Method for Classifying SWRs.

818 **Figure 3.** The distribution over models in panel **a** was inferred by random effects analysis as explained in Best-Fitting
819 Models, and Fixed and Random Effects Model Comparison of Dynamics Models, the comparison to the traditional method
820 as described in Traditional Method for Classifying SWRs, the simulated data model recovery as described in Model
821 Recovery from Simulated Data, and the deviance explained as described in Computing Fraction Deviance Explained.

822 **Figure 4.** The run snippets visualized in panel **a** were selected as described in Analysis of Place Cell Activity During
823 Movement for example session 1 (rat 1, day 1). The distribution over models was computed by random effects analysis as
824 described in Best-Fitting Models, and Fixed and Random Effects Model Comparison of Dynamics Models. The example
825 run snippets decoded positions are visualized following the same procedure outlined for Figure 2a (above), and the
826 distribution over models per-run snippet was computed as described in Marginalizing over Latent Position Sequences
827 and Model Parameters.

828 **Figure 5.** The heatmaps in panels **a** and **b** are computed as described for Figure 2a (above). For panel **c**, the time-
829 distance relationship, estimated diffusion coefficient in panel, and bootstrap significance testing (inset) are described in
830 Diffusion Coefficient Analysis.

831 **Figure 6.** The data presented here agglomerates SWRs best fit by a trajectory model, as described in Best-Fitting
832 Models, and Fixed and Random Effects Model Comparison of Dynamics Models, across all sessions, and splits SWRs
833 into “home” and “away” events as described in Splitting SWRs into “Home” and “Away” Events.

834 **Figure 7.** “Trajectory model” SWRs are selected as described in Best-Fitting Models, and Fixed and Random Effects
835 Model Comparison of Dynamics Models, “previously classified” SWRs are selected as described in Traditional Method
836 for Classifying SWRs, and SWRs are split into “home” and “away” events as described in Splitting SWRs into “Home” and
837 “Away” Events. The angular distance was quantified as described in Analysis of Correlation Between Replay Trajectories
838 and Behavior.