

# Uncovering temporal structure in hippocampal output patterns

Kourosh Maboudi<sup>1, 2†</sup>, Etienne Ackermann<sup>3†</sup>, Laurel Watkins de Jong<sup>1, 2</sup>, Brad Pfeiffer<sup>4</sup>, David Foster<sup>5</sup>, Kamran Diba<sup>1, 2\*†</sup>, Caleb Kemere<sup>3\*†</sup>

**\*For correspondence:**

[kdiba@umich.edu](mailto:kdiba@umich.edu) (KD);  
[caleb.kemere@rice.edu](mailto:caleb.kemere@rice.edu) (CK)

<sup>†</sup>These authors contributed equally to this work

<sup>1</sup>Department of Anesthesiology, University of Michigan, Ann Arbor, MI; <sup>2</sup>Department of Psychology, University of Wisconsin-Milwaukee, Milwaukee, WI; <sup>4</sup>Department of Neuroscience, University of Texas Southwestern, Dallas, TX; <sup>5</sup>Department of Psychology and Helen Wills Neuroscience Institute, University of California, Berkeley, CA; <sup>3</sup>Department of Electrical and Computer Engineering, Rice University, Houston, TX

**Abstract** Place cell activity of hippocampal pyramidal cells has been described as the cognitive substrate of spatial memory. Replay is observed during hippocampal sharp-wave-ripple-associated population burst events (PBEs) and is critical for consolidation and recall-guided behaviors. PBE activity has historically been analyzed as a phenomenon subordinate to the place code. Here, we use hidden Markov models to study PBEs observed in rats during exploration of both linear mazes and open fields. We demonstrate that estimated models are consistent with a spatial map of the environment, and can even decode animals' positions during behavior. Moreover, we demonstrate the model can be used to identify hippocampal replay without recourse to the place code, using only PBE model congruence. These results suggest that downstream regions may rely on PBEs to provide a substrate for memory. Additionally, by forming models independent of animal behavior, we lay the groundwork for studies of non-spatial memory.

## Introduction

Large populations of neurons fire in tandem during hippocampal sharp-waves and their accompanying CA1 layer ripple oscillations (*Buzsáki, 1986*). By now, multiple studies have shown that during many sharp-wave ripple-associated population burst events (PBEs), hippocampal "place cells" (*O'Keefe, 1976*) fire in compressed sequences that reflect the firing order determined by the sequential locations of their individual place fields (*Diba and Buzsáki, 2007; Foster and Wilson, 2006; Lee and Wilson, 2002; Nádasdy et al., 1999*). While the firing patterns during active exploration are considered to represent the brain's global positioning system and provide a substrate for spatial and episodic memory, instead it is the synchronized activity during PBEs that is most likely to affect cortical activity beyond the hippocampus (*Buzsáki, 1989; Carr et al., 2011; Diekelmann and Born, 2010; Siapas and Wilson, 1998*). Likewise, widespread activity modulation is seen throughout the brain following these sharp-wave ripple population bursts (*Logothetis et al., 2012*).

The literature on PBEs has largely focused on developing templates of firing patterns during active behavior and evaluating the extent to which these templates' patterns are reprised during subsequent PBEs. But what if the fundamental mode of the hippocampus is not the re-expression of place fields, but rather the PBE sequences during sharp-wave ripples (SWRs)? PBE sequences are enhanced during exploration of novel environments (*Cheng and Frank, 2008; Foster and Wilson, 2006*), they presage learning-related changes in place fields (*Dupret et al., 2010*), and appear to be critical to task learning (*Ego-Stengel and Wilson, 2010; Girardeau et al., 2009; Jadhav et al., 2012*).

42 Here, we examine the information provided by CA1 and CA3 pyramidal neurons, the output nodes  
43 of the hippocampus, through the looking glass of PBE firing patterns.

44 We developed a technique to build models of PBE sequences strictly outside of active exploration  
45 and independent of place fields and demonstrate that this nevertheless allows us to uncover spatial  
46 maps. Furthermore, these models can be used to detect congruent events that are consistent  
47 with replay but without any explicit place cell template. Our technique therefore provides new  
48 possibilities for evaluating hippocampal output patterns in single-trial and other fast learning  
49 paradigms, where a reliable sequential template pattern is not readily available. Overall, our work  
50 suggests that a sequence-first approach can provide an alternative view of hippocampal activity  
51 that may shed new light on how memories are formed, stored, and recalled.

## 52 Results

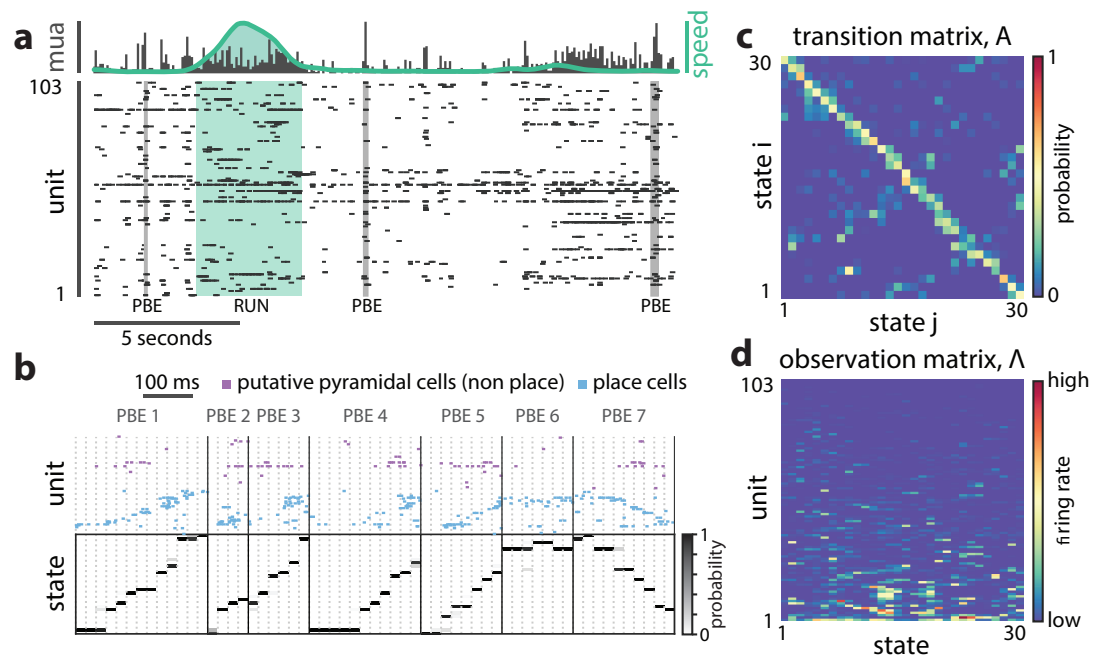
### 53 Awake population burst events

54 We began by analyzing the activity of large numbers of individual neurons in areas CA1 and CA3  
55 of the dorsal hippocampus as rats navigated linear mazes for water reward (linear track:  $n = 3$   
56 rats,  $m = 18$  sessions; previously used by *Diba and Buzsáki (2007)*). Using pooled multiunit activity,  
57 we detected PBEs during which many neurons were simultaneously active. The majority of these  
58 events occurred when animals paused running (speed  $< 5$  cm/s, corresponding to  $54.0\% \pm 20.1\%$  sd  
59 of events) to obtain reward, groom, or survey their surroundings (*Buzsáki et al., 1983*), and were  
60 accompanied by SWR complexes, distinguished by a burst of oscillatory activity in the 150–250 Hz  
61 band of the CA1 local field potential (LFP). Because we are interested in understanding internally  
62 generated activity during PBEs, we included only these periods without active behavior, ensuring  
63 that theta sequences would not bias our results. While we identified active behavior using a speed  
64 criterion, we found similar results when we instead used a theta-state detection approach (not  
65 shown). We did not add any other restrictions on behavior, LFPs, or the participation of place cells.  
66 We found that inactive PBEs occupied an average of 1.8% of the periods during which animals were  
67 on the linear track ( $16.9 \pm 15.1$  s of  $832.6 \pm 390.5$  s). In comparison, classical Bayesian approaches to  
68 understand PBE activity require the 34.8% of time animals are running (speed  $> 10$  cm/s) on the  
69 track ( $254.4 \pm 106.6$  s of  $832.6 \pm 390.5$  s) to build models of place fields.

### 70 Learning hidden Markov models from PBE data

71 Activity during PBEs is widely understood to be internally-generated in the hippocampal-entorhinal  
72 formation, and likely to affect neuronal firing in downstream regions (*Buzsáki, 1989; Chrobak and*  
73 *Buzsáki, 1996; Logothetis et al., 2012; Yamamoto and Tonegawa, 2017*). Given the prevalence of  
74 PBEs during an animal's early experience, we hypothesized that the neural activity during these  
75 events would be sufficient to train a machine learning model of sequential patterns—a hidden  
76 Markov model—and that this model would capture the relevant spatial information encoded in the  
77 hippocampus independent of exploration itself.

78 Hidden Markov models have been very fruitfully used to understand sequentially structured  
79 data in a variety of contexts. A hidden Markov model captures information about data in two  
80 ways. First, it clusters observations into groups ("states") with shared patterns. In our case, this  
81 corresponds to finding time bins in which the same sets of neurons are co-active. This is equivalent  
82 to reducing the dimension of the ensemble observations into a discretized latent space or manifold.  
83 Second, it models the dynamics of state transitions. This model is Markovian because it is assumed  
84 that the probability to transition to the next state only depends on the current state. Critically,  
85 these operations of clustering and sequence modeling are jointly optimized, allowing the structure  
86 of ensemble firing corresponding to each of the final states to combine information over many  
87 observations. Given the role of the hippocampus in memory, in our hidden Markov models (HMMs),  
88 the unobserved latent variable presumably corresponds to the temporal evolution of a memory  
89 trace that is represented by co-active ensembles of CA1 and CA3 neurons. The full model will



**Figure 1.** A hidden Markov model of ensemble activity during population burst events. **a.** Examples of three PBEs and a run epoch. **b.** Spikes during 7 example PBEs (top) and their associated (30 state HMM-decoded) latent space distributions (bottom). The place cells are ordered by their place fields on the track, whereas the non-place cells are unordered. The latent states are ordered according to the peak densities of the latent-state place fields (IsPFs, see Materials and Methods). **c.** The transition matrix models the dynamics of the unobserved internally-generated state. The sparsity and banded-diagonal shape are suggestive of sequential dynamics. **d.** The observation model of our HMM is a set of Poisson probability distributions (one for each neuron) for each hidden state. Looking across columns (states), the mean firing rate is typically elevated for only a few of the neurons and individual neurons have elevated firing rates for only a few states.

**Figure 1-Figure supplement 1.** Hidden Markov models capture state dynamics beyond pairwise co-firing.

90 correspond to the structure which connects all the memory traces activated during PBEs.

91 The parameters of our model that are fit to data include the observation model (the cluster  
 92 descriptions, or predicted activity of each excitatory neuron within the CA1/CA3 ensemble for a  
 93 given state), the state transition model (the probability that the CA1/CA3 ensemble will transition  
 94 from a start state to a destination state in the next time bin), and the initial state distribution (the  
 95 probability for sequences to start in each given state). In prior work using HMMs to model neural  
 96 activity, a variety of statistical distributions have been used to characterize ensemble firing during a  
 97 specific state (the observation model, *Chen and Wilson (2017); Chen et al. (2012, 2014); Deppisch*  
 98 *et al. (1994); Kemere et al. (2008); Radons et al. (1994)*). We opted for the Poisson distribution to  
 99 minimize the number of parameters per state and per neuron (see Materials and Methods). We  
 100 used the standard iterative expectation-maximization (EM) algorithm (*Rabiner, 1989*) to learn the  
 101 parameters of an HMM from binned PBE data (20 ms bins). **Figure 1** depicts the resultant state  
 102 transition matrix and observation model for an example linear-track session.

103 Using separate training- and test-datasets (cross-validation) mitigates over-fitting to training  
 104 data, but it is still possible for the cross-validated goodness-of-fit to increase with training without  
 105 any underlying dynamics, e.g., if groups of neurons tend to activate in a correlated fashion. Does the  
 106 model we have learned reflect underlying sequential structure of memory traces beyond pairwise  
 107 co-firing? To answer this question, we cross-validated the model against both real “test” data  
 108 and against surrogate “test” data derived from shuffling each PBE in two ways: one in which the  
 109 binned spiking activity was circularly permuted across time for each neuron independently of the  
 110 other neurons (“temporal shuffle”, which removes co-activation), and one in which the order of  
 111 the binned data was scrambled coherently across all neurons (“time-swap”, which maintains co-

112 activation). Note that the second shuffle preserves pairwise correlations while removing the order  
 113 of any sequential patterns that might be present. Using five-fold cross-validation, we compared  
 114 learned models against both actual and surrogate test data and found that the model likelihood  
 115 was significantly greater for real data (vs. temporal shuffle,  $p < 0.001$ , vs. time-swap,  $p < 0.001$ ,  $n = 18$   
 116 sessions, Wilcoxon signed-rank test, **Figure 1–Figure Supplement 1**).

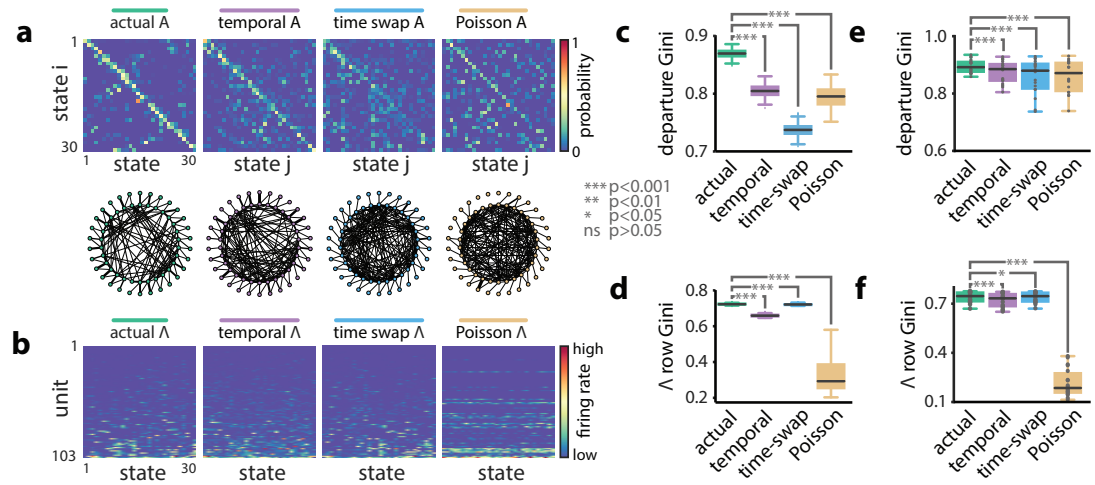
### 117 **What do the learned model parameters tell us about PBEs?**

118 To begin to understand what structure we learn from PBE activity, we compared our HMMs (trained  
 119 on real data) against models trained on multiple different surrogate datasets (**Figure 2a,b**). These  
 120 surrogate datasets were obtained from actual data following: 1) temporal shuffles and 2) time-  
 121 swaps, as above, and 3) by producing a surrogate PBE from independent Poisson simulations  
 122 according to each unit's mean firing rate within the original PBEs. First, we investigated the sparsity  
 123 of the transition matrices using the Gini coefficient (see Materials and Methods and **Figure 2–Figure**  
 124 **Supplement 1**). A higher Gini coefficient corresponds to higher sparsity. Strikingly, the actual data  
 125 yielded models in which the state transition matrix was sparser than in each of the surrogate  
 126 counterparts ( $p < 0.001$ , **Figure 2c**), reflecting that each state transitions only to a few other states.  
 127 Thus, intricate yet reliable details are captured by the HMMs. Next, we quantified the sparsity of the  
 128 observation model. We found that actual data yielded mean firing rates which were highly sparse  
 129 (**Figure 2d**), indicating that individual neurons were likely to be active during only a small fraction  
 130 of the states. Using a graph search algorithm (see Materials and Methods), we simulated paths  
 131 through state space generated by these transition matrices, and found that this increased sparsity  
 132 accompanied longer trajectories (**Figure 2–Figure Supplement 3**) through the state space of the  
 133 model. Thus, the state transition matrices we learn are suggestive of dynamics in which each sparse  
 134 state is preceded and followed by only a few other, in turn, sparse states, providing long sequential  
 135 paths through state space-consistent with spatial relationships in the environment in which the  
 136 animal was behaving, but generated from PBEs. The increased sparsity of the observation model  
 137 and transition matrix in the example session was representative of a significant increase over all  
 138 remaining sessions ( $p < 0.05$ ,  $n = 18$  sessions, Wilcoxon signed-rank tests, **Figure 2e,f**).

139 These observations indicate that PBEs inform an HMM about extant spatial relationships within  
 140 the environment. So, next we asked how the firing patterns of neurons during actual behavior  
 141 project into the learned latent spaces. To observe the evolution of the latent states during behavior,  
 142 we used our model to determine the most likely sequence of latent states corresponding to decode  
 143 the neural activity observed in 100 ms bins during epochs that displayed strong theta oscillations  
 144 (exclusive of PBEs) when rats were running (speed  $> 10$  cm/s; see Materials and Methods). If the  
 145 learned model was distinct from ensemble patterns during behavior, we might expect the resulting  
 146 state space probability distributions at each point in time to be randomly spread among multiple  
 147 states. Instead, we found distributions that resembled sequential trajectories through the latent  
 148 space (**Figure 3a**) in parallel with the physical trajectories made by the animal along the track, further  
 149 demonstrating that the latent state dynamics learned from PBEs corresponds to an internalized  
 150 model of physical space.

151 To better understand the relationship between the latent space and physical space, we used  
 152 the latent state trajectories decoded during running to form an estimate of the likelihood of each  
 153 state as a function of location on the track (see Materials and Methods). These “latent-state  
 154 place fields” (lsPFs, **Figure 3b**) in many ways resembled neuronal place fields and similarly tiled  
 155 the extent of the track. This spatial localization went away when we re-estimated the lsPFs with  
 156 shuffled positions (**Figure 3c**). To quantify how informative the latent states were about position,  
 157 we used the lsPFs to map decoded state sequences to position during running periods (**Figure 3d**).  
 158 In our example session, decoding through the latent space resulted in a median accuracy of 5 cm,  
 159 significantly greater than the 47 cm obtained from shuffled lsPFs ( $p < 0.001$ , Wilcoxon signed-rank  
 160 test, **Figure 3d**). When we evaluated decoding error over our full set of sessions, we observed a  
 161 similar result ( $p < 0.001$ , Wilcoxon signed-rank test, **Figure 3e, Figure 3–Figure Supplement 1**). As



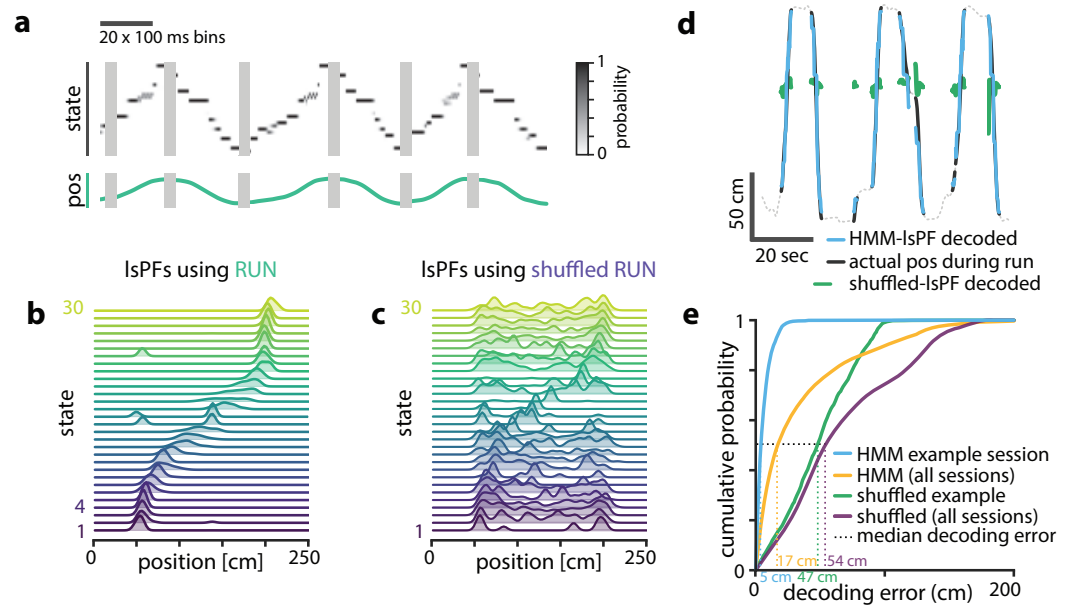


**Figure 2.** Models of PBE activity are sparse. We trained HMMs on neural activity during PBEs (in 20 ms bins), as well as on surrogate transformations of those PBEs. **a.** (top) The transition matrices for the actual and surrogate PBE models with states ordered to maximize the transition probability from state  $i$  to state  $i + 1$ . (bottom) Undirected connectivity graphs corresponding to the transition matrices. The nodes correspond to states (progressing clockwise, starting at the top). The weights of the edges are proportional to the transition probabilities between the nodes (states). The transition probabilities from state  $i$  to every other state except  $i + 1$  are shown in the interior of the graph, whereas for clarity, transition probabilities from state  $i$  to itself, as well as to neighboring state  $i + 1$  are shown between the inner and outer rings of nodes (the nodes on the inner and outer rings represent the same states). **b.** The observation matrices for actual and surrogate PBE models show the mean firing rate for neurons in each state. For visualization, neurons are ordered by their firing rates. **c.** We quantified the sparsity of transitions from one state to all other states using the Gini coefficient of rows of the transition matrix for the example session in **a.**. Actual data yielded sparser transition matrices than shuffles. **d.** The observation models—each neuron’s expected activity for each state—learned from actual data for the example session are significantly sparser than those learned after shuffling. This implies that as the hippocampus evolves through the learned latent space, each neuron is active during only a few states. **e.** Summary of transition matrix sparsity and **f.** Observation model sparsity with corresponding shuffle data pooled over all sessions/animals. (\*\*\*:  $p < 0.001$ , \*:  $p < 0.05$ ; single session comparisons:  $n = 250$  realizations, Welch’s t-test; aggregated comparisons -  $n = 18$  sessions, Wilcoxon signed-rank test).

**Figure 2-Figure supplement 1.** PBE model states typically only transition to a few other states.

**Figure 2-Figure supplement 2.** Each neuron is active in only a few model states.

**Figure 2-Figure supplement 3.** The sparse transitions integrate into long sequences through the state space.



**Figure 3.** Latent states capture positional code. **a.** Using the model parameters estimated from PBEs, we decoded latent state probabilities from neural activity during periods when the animal was running. An example shows the trajectory of the decoded latent state probabilities during six runs across the track. **b.** Mapping latent state probabilities to associated animal positions yields latent-state place fields (IsPFs) which describe the probability of each state for positions along the track. **c.** Shuffling the position associations yields uninformative state mappings. **d.** For an example session, position decoding during run periods through the latent space gives significantly better accuracy than decoding using the shuffled tuning curves. The dotted line shows the animal's position during intervening non run periods. **e.** The distribution of position decoding accuracy over all sessions ( $n = 18$ ) was significantly greater than chance. ( $p < 0.001$ ).

**Figure 3-Figure supplement 1.** Latent states capture positional code over wide range of model parameters.

162 our method required discretizing the state space, a potential caveat is that the number of latent  
 163 states is a relevant parameter, which we arbitrarily chose to be 30. However, latent-state place  
 164 fields were informative of position over a wide range of values of this parameter (**Figure 3-Figure**  
 165 **Supplement 1**). Note that decoding into the latent space and then mapping to position resulted  
 166 in slightly higher error than simply performing Bayesian decoding on the neural activity during  
 167 behavior. This suggests that the latent space we learn from PBEs may not capture all the information  
 168 about space that is present in hippocampal activity during behavior, though this may also reflect  
 169 the limited number of PBEs from which we can learn.

### 170 HMM-congruent PBEs capture sequence replay

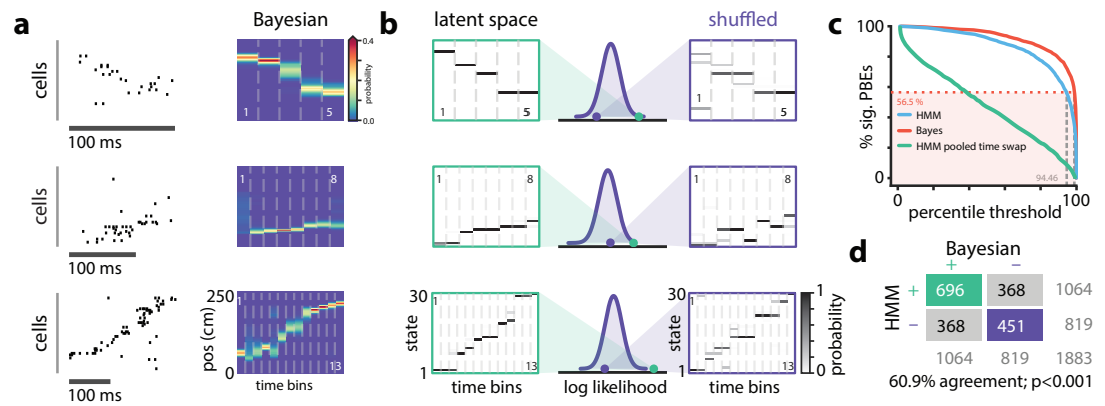
171 We and others have previously described how the pattern of place cell firing during many PBEs  
 172 recapitulates the order in which they are active when animals run on the track (**Figure 4a**). We  
 173 employed the versatile and widely-used Bayesian decoding method to ascribe a replay score to  
 174 sequential patterns during PBEs. Briefly, for each PBE, we used place-field maps to estimate a  
 175 spatial trajectory (an *a posteriori* distribution of positions) in 20 ms bins. We generated surrogate  
 176 data via a column-cycle shuffle (i.e., a circular shift across positions for each time bin (**Davidson**  
 177 **et al., 2009**)) of the *a posteriori* distributions during PBEs. The real and surrogate trajectories were  
 178 scored (see Materials and Methods), and we defined replay events as those for which the score of  
 179 the actual trajectory was larger than a threshold fraction of the null distribution generated by the  
 180 surrogate scores. Using this approach, we found that 57% of PBEs (1064 of 1883) were identified as  
 181 replay beyond a threshold of 99% (median across datasets 54.2%, interquartile range = 32.8–61.0%,  
 182 **Figure 4-Figure Supplement 1**). Thus, as has been reported many times (**Davidson et al., 2009; Diba**

183 *and Buzsáki, 2007; Foster and Wilson, 2006; Karlsson and Frank, 2009*), only a fraction of PBEs (but  
 184 many more than expected by chance) represent statistically significant replay. Given that we use all  
 185 PBEs for model learning and our models capture the structure of the environment and the patterns  
 186 expressed by place cells during exploration, we were interested in understanding whether we could  
 187 also use our latent-space models to find these replay events. Indeed, for many events when we  
 188 decode trajectories through state space, they resemble the sequential patterns observed when  
 189 we decode position using Bayesian techniques and the place cell map (**Figure 4b**, left). However,  
 190 given previous evidence for replay of environments not recently experienced (*Gupta et al., 2010;*  
 191 *Karlsson and Frank, 2009*), we hypothesized that some PBEs might contain ensemble neural activity  
 192 which is unstructured and thus unrelated to the learned model, and that these would correspond  
 193 to the “non-replay” events found using traditional methods.

194 To assess how well the pattern of ensemble activity during individual PBEs related to the overall  
 195 state-space model learned from PBE activity (“congruence”), we developed a statistical approach  
 196 for identifying the subset of strongly structured PBEs. Specifically, rather than comparing real and  
 197 surrogate PBEs, we compared the goodness-of-fit for each event to a null distribution generated  
 198 via a computationally-efficient manipulation of the transition matrix of the model (**Figure 4b**);  
 199 we row-wise shuffled the non-diagonal elements of the transition matrix to assess whether an  
 200 individual PBE is a more ordered sequence through state space than would be expected by chance.  
 201 Maintaining the diagonal avoids identifying as different from chance sequences which consist of few  
 202 repeated states, marked by transitions between state  $i$  and itself. As described above, the fraction  
 203 of events identified as replay using Bayesian decoding is strongly tied to how the null-distribution  
 204 is generated (i.e., what shuffle is used), some secondary criteria (e.g., number of active cells, unit  
 205 cluster quality, peak firing rate, trajectory “jumps”, etc.), and the value of the significance threshold  
 206 arbitrarily chosen to be 90%, 95%, or 99% of shuffles in different reports. When we combined  
 207 across datasets, we found that our transition matrix shuffle yielded a null distribution for which a  
 208 99% confidence interval identified slightly fewer PBEs as significant than the column-cycle shuffle  
 209 did for Bayesian decoding (**Figure 4c**). To make a principled comparison of Bayesian- and HMM-  
 210 based replay detection schemes, we fixed the Bayesian-based significance threshold at 99% but  
 211 selected the significance threshold for the HMM-congruence null distribution so that the fraction  
 212 of replay events detected would be the same between the two schemes. Following this approach,  
 213 we found that model-congruent/incongruent PBEs largely overlapped with the replay/non-replay  
 214 events detected using Bayesian decoding of the place cell map (**Figure 4d**). Thus, using only the  
 215 neural activity during PBEs, without access to any place cell activity, we are remarkably able to  
 216 detect the sequential patterns typically described as “replay” based only on their consistency with  
 217 the structure of other PBE activity.

218 There were, however, also differences between the Bayesian and HMM-congruent approaches,  
 219 including events that reached significance in one but not the other formalism. We wanted to  
 220 understand where and why these approaches differed in identifying significant sequences. When  
 221 we examined individual PBEs, we found sequences for which both Bayesian and model-congruence  
 222 replay detection approaches appeared to malfunction (**Figure 5a**). This was not a failure of the  
 223 choice of significance threshold, as for both techniques we found what appeared to be false-  
 224 negatives (patterns which looked like replay but were not significant) as well as false-positives  
 225 (patterns which looked noisy but were identified as replay). Thus, in order to quantitatively compare  
 226 the two approaches, we asked eight humans to visually examine all the PBEs in our database. They  
 227 were instructed to label as replay PBEs in which the animal’s Bayesian decoded position translated  
 228 sequentially without big jumps (*Silva et al., 2015*, see Materials and Methods).

229 We marked each event as a “true” community replay if it was identified by a majority of scorers  
 230 (six individuals scored  $n = 1883$  events, two individuals scored a subset of  $n = 1423$  events, individual  
 231 scores are shown in **Figure 5–Figure Supplement 1**). We calculated an receiver operating charac-  
 232 teristic (ROC) curve which compared the rate of true positive and false positive detections as the  
 233 significance thresholds for Bayesian and model-congruence approaches were varied (**Figure 5b**). A



**Figure 4.** **a.** Example PBEs decoded to position using Bayesian decoding. **b.** (left) Same examples decoded to the latent space using the learned HMM. (right) Examples decoded after shuffling the transition matrix, and (middle) the sequence likelihood using actual and shuffled models. **c.** Effect of significance threshold on the fraction of events identified as replay using Bayesian decoding and model congruent events using the HMM approach. **d.** Comparing Bayesian and model-congruence approaches for all PBEs recorded, we find statistically significant agreement in event identification (60.9% agreement,  $n = 1883$  events from 18 sessions,  $p < 0.001$ , Fisher's exact test two sided).

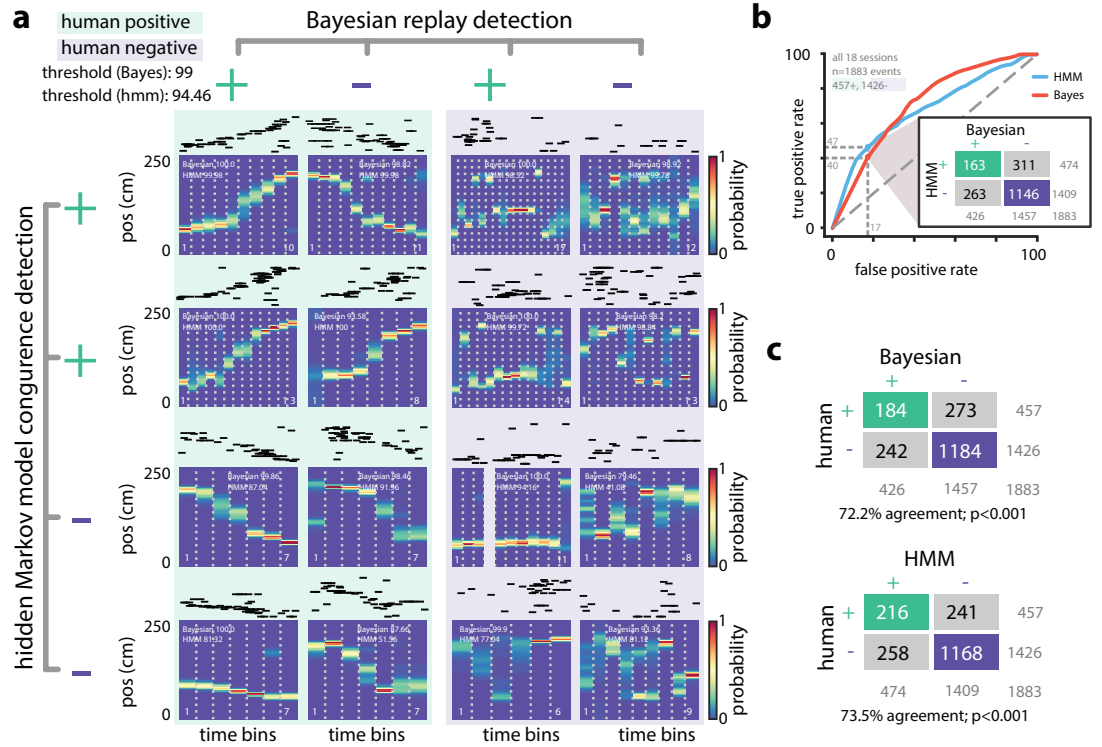
**Figure 4-Figure supplement 1.** Number of significant PBEs.

234 perfect detector would have an area under the curve (AUC) of unity. We did not find a significant  
 235 difference between the AUCs of Bayesian decoding and model-congruence ( $p = 0.14$ , bootstrap,  
 236 see Methods). If we select thresholds such that our algorithms yield a similar fraction of significant  
 237 vs. total events as the 24% denoted by our human scorers, we find that both Bayesian and model-  
 238 congruence yield agreement of  $\approx 70\%$  labeled events with each other and with human scorers  
 239 (**Figure 5c**).

240 Thus, congruence with an HMM trained only on PBEs appears to work as reliably as Bayesian  
 241 decoding in detecting sequential reactivation of linear track behaviors. However, when we examined  
 242 individual sessions, we noticed that performance was quite variable. Given that our models are  
 243 learned only from PBEs, we reasoned that the statistics or structure of the PBEs within each session  
 244 might yield models which vary in quality depending on the number of recorded units, the number  
 245 of PBEs detected, and their self-consistency across events. We created a model quality metric by  
 246 comparing cross-validated learning statistics to models which were learned from shuffled events  
 247 (see Materials and Methods). We found that the performance of model-congruence detection was  
 248 tied to model quality ( $R^2 = 0.17$ ,  $F = 2.9$ ,  $n = 18$  sessions, **Figure 5-Figure Supplement 1**). Model  
 249 quality, in turn, was highly correlated with the number of PBEs during the session ( $R^2 = 0.96$ ,  
 250  $F = 392.6$ ,  $n = 18$  sessions, **Figure 5-Figure Supplement 1**). Not surprisingly, the performance of  
 251 Bayesian decoding relative to human scorers was independent of the quality of the HMM, or the  
 252 number of PBEs, as the place field model is learned from ensemble neural activity during running.  
 253 Thus, we find an intriguing contrast—when there is an abundance of PBEs (indicating novelty,  
 254 learning, hippocampus-dependent planning, etc. (**Buzsáki, 2015**)), even in the absence of repeated  
 255 experience, replay detection based on PBE activity is highly effective. Conversely, when there are  
 256 few PBEs (i.e., scenarios in which PBEs are uncorrelated with cognitive function), but an abundance  
 257 of repeated behavioral trials, Bayesian decoding of these limited events proves more effective.

## 258 Modeling internally generated activity during open field behavior

259 The linear track environment represents a highly-constrained behavior. We therefore asked whether  
 260 the hidden Markov model approach could generalize to more complex environments and behavioral  
 261 tasks. **Pfeiffer and Foster (2013, 2015)** had previously recorded activity of CA1 neurons in rats as  
 262 they explored in a 2 m  $\times$  2 m open field arena for liquid reward. Briefly, animals were trained to  
 263 discover which one of 36 liquid reward wells would be the “home” well on a given day. They then



**Figure 5. a.** Eight examples from one session show that Bayesian decoding and HMM model-congruence can differ in labeling of significant replay events. For each event, spike rasters (ordered by the location of each neuron's place field) and the Bayesian decoded trajectory are shown. "+" ("-") label corresponds to significant (insignificant) events. (left) Both methods can fail to label events that appear to be sequential as replay and (right) label events replay that appear non-sequential. **b.** We recruited human scorers to visually inspect Bayesian decoded spike trains and identify putative sequential replay events. Using their identifications as labels, we can define an ROC curve for both Bayesian and HMM model-congruence which shows how detection performance changes as the significance threshold is varied. (inset) Human scorers identify 24% of PBEs as replay. Setting thresholds to match this value results in agreement of 70% between Bayesian and HMM model-congruence. **c.** Using the same thresholds, we find  $\approx 70\%$  agreement between algorithmic and human replay identification. (All comparison matrices,  $p < 0.001$ , Fisher's exact test two-tailed.)

**Figure 5-Figure supplement 1.** Human scoring of PBEs and session quality.



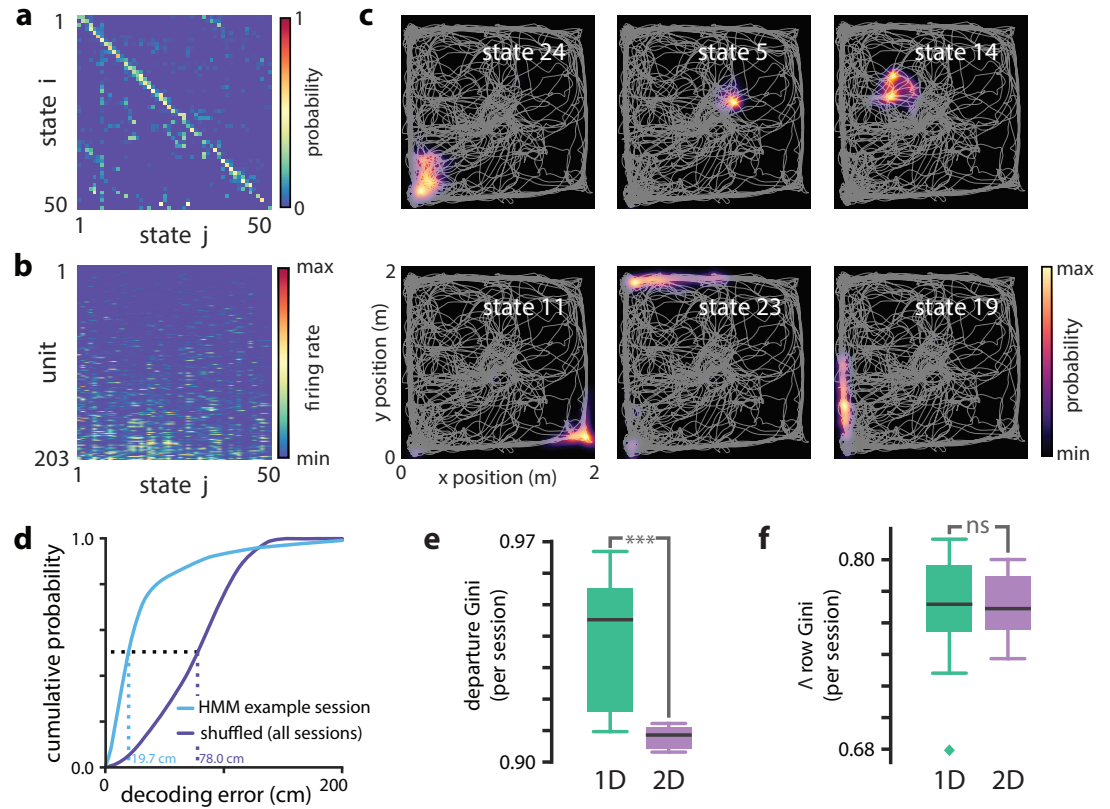
264 were required to alternate between searching for a randomly rewarded well and returning to the  
 265 home well. Using the place cell map in this task and Bayesian decoding, many PBEs were decoded  
 266 to trajectories through two-dimensional space that were predictive of behavior and shaped by  
 267 reward. Using this same dataset, we trained a HMMs on neural activity during PBEs in the open field.  
 268 Here, we used the same PBEs detected previously (*Pfeiffer and Foster, 2013, 2015*) which occupied  
 269 an average of  $2.53 \pm 0.42\%$  of the periods during which animals were behaving ( $77.91 \pm 21.16$  s out  
 270 of  $3064.86 \pm 540.26$  s). Given the large number of units available in this dataset and the increased  
 271 behavioral variability in the open field environment compared to the linear track, we chose to  
 272 estimate HMMs with 50 latent states. The transition matrix and observation model from a sample  
 273 session are shown in *Figure 6a,b*. Despite the complex and varied trajectories displayed by animals,  
 274 the HMM captured sequential dynamics in PBE activity, as in the 1D case, when we compared  
 275 learned models against both actual and surrogate test data, we found that the model likelihood  
 276 was significantly greater for real data ( $p < 0.001$ , Wilcoxon signed-rank test).

277 In the case of the linear track, we linked sparsity of the transition matrix to the sequential nature  
 278 of behaviors in that environment. An unconstrained, two-dimensional environment permits a much  
 279 richer repertoire of behavioral trajectories. However, behavior is still constrained by the structure of  
 280 space—arbitrary teleportation from one location to another is impossible. We found that learning  
 281 from PBEs in the open field yielded transition matrices (*Figure 6a*) that were significantly sparser  
 282 than models learned from shuffled data ( $p < 0.05$ , Wilcoxon signed-rank test,  $n = 8$  sessions,  
 283 *Figure 5–Figure Supplement 1*). However, consistent with increased freedom of potential behaviors,  
 284 when we compared the sparsity of models learned from open field PBEs with 50-state models  
 285 learned from PBEs in linear tracks, the open field transition matrices were less sparse ( $p < 0.001$ ,  
 286 Mann–Whitney  $U$  test comparing 8 and 18 sessions, *Figure 4–Figure Supplement 1*). Likewise,  
 287 when we examined the observation model for the open field, we found that the activity across  
 288 states for individual neurons was significantly more sparse than in models learned from shuffled  
 289 data ( $p < 0.05$ , Wilcoxon signed-rank test,  $n = 8$  sessions, *Figure 6–Figure Supplement 1*). The  
 290 sparsity of linear track and open field observation models were not significantly different ( $p = 0.44$ ,  
 291 Mann–Whitney  $U$  test).

292 Do the latent states learned from PBEs capture spatial information in a 2D environment? We  
 293 used the PBE-trained model to decode run data, as in the linear track case. We found that the latent  
 294 states corresponded with specific locations in the open field, as we expected (*Figure 6c*). Moreover,  
 295 we were able to decode animals' movements with significantly greater than chance accuracy by  
 296 converting decoded latent states to positions using the IsPF ( $p < 0.001$ , *Figure 6d*). Finally, we  
 297 examined model-congruency for PBEs detected in the open field. Previously, it was reported that  
 298 27.3% (815 of 2980,  $n = 8$  sessions) were identified as “trajectory events” (*Pfeiffer and Foster,*  
 299 *2015*). We chose a significance threshold to match this fraction (*Figure 6–Figure Supplement 3*)  
 300 and found that there was significant overlap between the events detected through Bayesian and  
 301 model-congruence techniques ( $p < 0.01$ , Fisher's exact test). These events overlapped significantly  
 302 with replay events detected using traditional Bayesian decoding (*Figure 6–Figure Supplement 3*).  
 303 Thus, an HMM of the activity during population bursts captures the structure of neural activity in  
 304 two dimensional environments during complex tasks and can be used to decode events consistent  
 305 with trajectories through that environment.

### 306 **Extra-spatial Information**

307 As described earlier, while we observed a similar fraction of events to be similar by HMM-congruence  
 308 and Bayesian decoding, there was not an exact event-to-event correspondence. An intriguing po-  
 309 tential explanation is that the latent space represented in PBE sequential firing and captured by the  
 310 HMM is richer than simply the spatial structure of the present environment. In most hippocampal  
 311 ensemble recording experiments, maze or open field tasks are structured to intentionally map  
 312 memory elements to spatial behavior, and thus this potential richness is difficult to test. We used  
 313 two sample datasets to explore the potential of the HMM to capture extra-spatial richness in the



**Figure 6.** Modeling PBEs in open field. **a.** The transition matrix estimated from activity detected during PBEs in an example session in the open field. **b.** The corresponding observation model (203 neurons) shows sparsity similar to the linear track. **c.** Example latent state place fields show spatially-limited elevated activity in two dimensions. **d.** For an example session, position decoding through the latent space gives significantly better accuracy than decoding using the shuffled latent state place fields. **e.** Comparing the sparsity of the transition matrices (mean Gini coefficient of the departure probabilities) between the linear track and open field reveals that, as expected, over the sessions we observed, the open field is significantly *less sparse* ( $p < 0.001$ ), since the environment is less constrained. **f.** In contrast, there is not a significant difference between the sparsity of the observation model (mean Gini coefficient of the rows) between the linear track and the open field. Note that the linear track models are sparser than in *Figure 2* due to using 50 states rather than 30 to match the open field.

**Figure 6-Figure supplement 1.** Open field PBE model states typically only transition to a few other states.

**Figure 6-Figure supplement 2.** Each neuron is active in only a few model states in the open field.

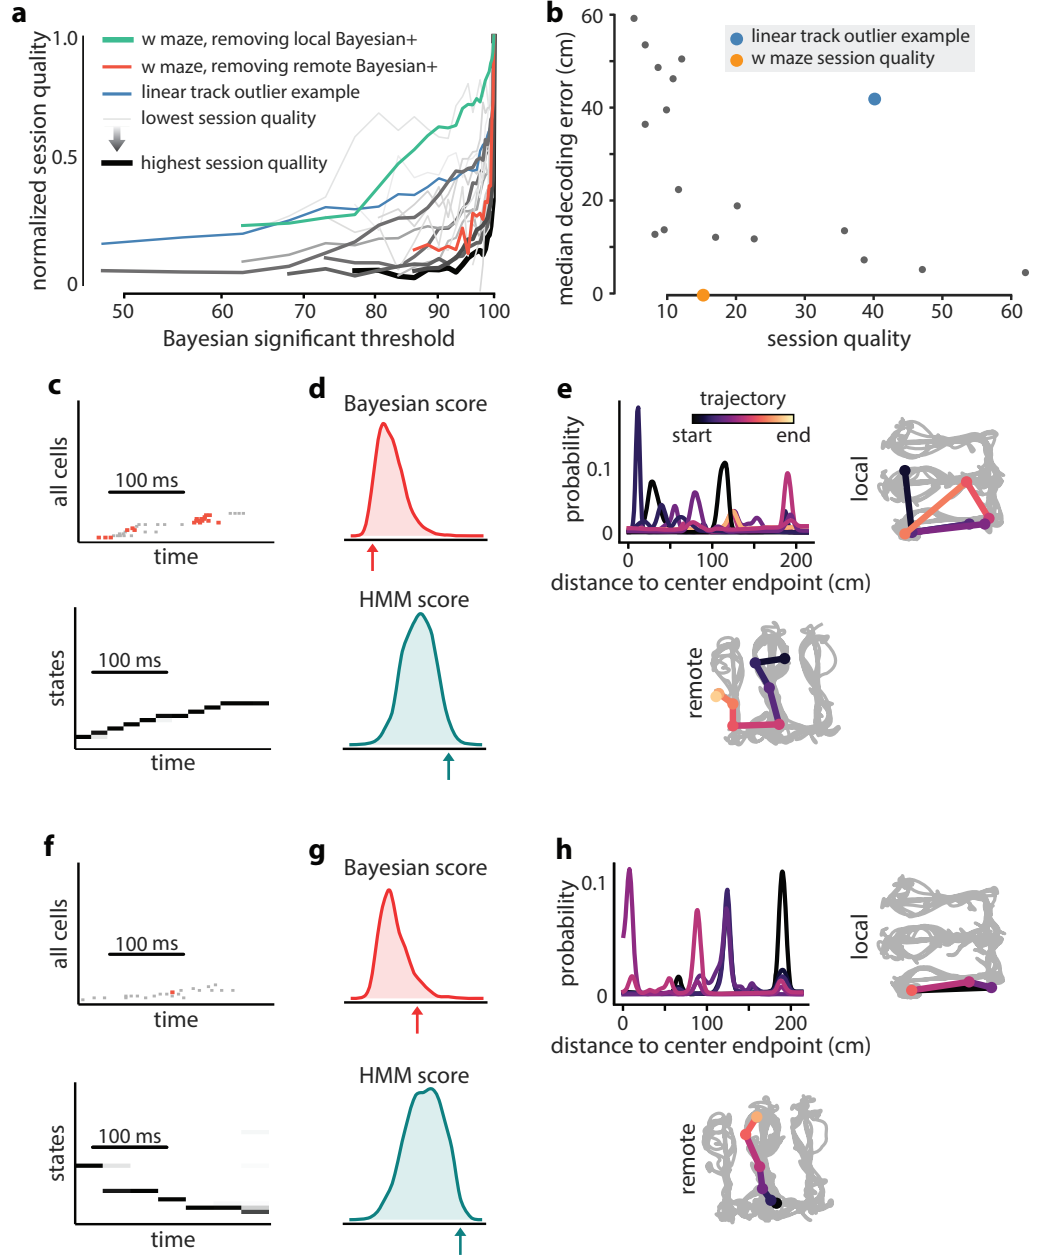
**Figure 6-Figure supplement 3.** IsPFs and position decoding in an open field.

**Figure 6-Figure supplement 4.** Examples of open field PBEs.

314 PBE sequences.

315 First, we considered the possibility that in the awake behaving animal, PBE activity might be  
316 sequential reactivation of environments other than the one being explored (“remote replay”). We  
317 reasoned that we could enhance the model’s representation of remote environments by filtering  
318 out local replay from the training data. We evaluated how the model-quality of our HMM changed  
319 as progressively more sequences labeled as replay by Bayesian decoding were removed from the  
320 training data. In the linear track sessions we considered, we found that refining the training data  
321 resulted in models that lowered in quality at different rates as the threshold for Bayesian replay  
322 was decreased (**Figure 7**). Most, but not all, models dropped precipitously in quality: > 50% when  
323 we removed events detected as Bayesian replay at a 95% threshold, as would be expected if the  
324 HMM represented only the local environment. In many outlier sessions in which model quality  
325 decreased more slowly, the initial (baseline) model quality was low. Intriguingly, however, in at least  
326 one outlier session where model quality decreased slowly with refinement (blue line, **Figure 7a**), the  
327 initial model quality was still high, and we further noted that position decoding using ISPFs yielded  
328 relatively high error (blue dot, **Figure 7b**). Thus, we wondered whether this and similar sessions  
329 might have contained non-local or extra-spatial PBEs that were captured by the HMM.

330 In order to validate the concept of model-training refinement, we considered a dataset in which  
331 multiple environments were explored on the same day and remote replay was previously observed  
332 (**Karlsson et al., 2015**). These data consisted of a series of short exploratory sessions in which an  
333 animal first explored a novel maze (E1) and then was placed in a familiar one (E2). We identified  
334 awake PBEs during the familiar E2 session and used them to train an HMM. When we refined this  
335 model by removing Bayesian-significant local replay events from the training data, we found that  
336 the model quality decreased comparatively slowly (**Figure 7a**, green line), indicating that the HMM  
337 was capturing more than the local spatial structure. In contrast, when we used place fields from E1  
338 to identify Bayesian-significant remote replay events and removed these from the training data, we  
339 found that the model quality decreased rapidly as with the general linear track cases (**Figure 7a**,  
340 red line). When we examined individual events in detail in this data, we found many examples  
341 in which HMM-significant, Bayesian non-significant PBEs decoded to extended state sequences  
342 which turned out to correspond to reactivation of the remote track (two are shown in **Figure 7c-l**).  
343 If we imagine that in this experiment data were only recorded during exploration of the familiar  
344 environment, classical Bayesian decoding would treat these events as noise, as shown in the bottom  
345 half of the two examples. In contrast, our HMM-based analysis finds these events to be significant,  
346 as shown in the top half of the two examples. Thus, by combining classical Bayesian decoding and  
347 HMM-congruence, we are able to identify a signature of when a HMM trained on PBEs captures  
348 sequential structure distinct from that dictated by the local environment. Additionally, in these  
349 cases, we show that specific non-local reactivation events can be identified.




**Figure 7.** Examples of remote replay events identified with HMM-congruence. We trained and evaluated HMMs on the events that were not Bayesian significant (residual events) to identify potential extra-spatial structure. **a.** The normalized session quality drops as local-replay events above the Bayesian significance threshold are removed from the data. Each trace corresponds to one of the 18 linear track sessions, with the stroke width and the stroke intensity proportional to the baseline (all-events) session quality. The blue line identifies a session in which model quality drops more slowly, indicating the potential presence of extra-spatial information. The reduction in session quality for a W maze experiment with known extra-spatial information is even slower (green). When, instead, Bayesian-significant *remote* events are removed, rapid reduction in session quality is again revealed (red). **b.** The IsPF-based median decoding errors are shown as a function of baseline session quality for all 18 linear track sessions. The blue dot indicates the outlier session from panel **a** with potential extra-spatial information: this session shows high decoding error combined with high session quality. Session quality of the W maze session is also indicated on the x-axis (decoding error is not directly comparable). **c-n.** Two example HMM-congruent but not Bayesian-significant events from the W maze session are depicted to highlight the fact that congruence can correspond to remote replay. **c.** Spikes during ripple with local place cells highlighted (top panel) and the corresponding latent state probabilities (bottom panel) decoded using the HMM show sequential structure (grayscale intensity corresponds to probability). **d.** In this event, the Bayesian score relative to the shuffle distribution (top panel) indicates that the event is not-significant, whereas the HMM score relative to shuffles indicates (bottom panel) the ripple event is HMM-congruent. **e.** Estimates of position using local place fields show jumpy, multi-modal *a posteriori* distributions over space in 1D (top left panel) and 2D (top right panel; distribution modes and time is denoted in color). Bayesian decoding using the remote environment place fields (bottom panel) indicates that the sample event is a remote replay. Note that in a typical experiment, only the local place fields would be available. **f-h.** Same as **c-e**, but for a different ripple event.

350

351





**Figure 8.** Temporal structure during a sleep period following object-location memory task. Using cross validation, we calculate the HMM-congruence score (which ranges from 0 to 1) for test PBEs. For each event, we also calculate the score of a surrogate chosen using a pooled time-swap shuffle across all test events. The distribution of scores of actual events is significantly higher than that of the surrogate data ( $p < 0.001$ , Mann-Whitney  $U$  test).

352 Finally, we considered the potential of our methodology for uncovering temporal patterns  
 353 in PBE activity under scenarios where complex behavior does not permit identification of well-  
 354 defined place-fields or in the absence of behavior, such as during sleep. As we have emphasized,  
 355 a remarkable aspect of learning HMMs from PBE activity is that the model can be built entirely  
 356 without behavioral data, so can our model capture significant sequential information outside  
 357 of immobility periods during quiet waking? To demonstrate this potential, we examined HMMs  
 358 trained on PBEs in sleep following the learning phase of an object-location memory task when  
 359 animals explored three objects in an open field (see Material and Methods). Previous studies have  
 360 demonstrated that subsequent recall of this memory is hippocampus-dependent, and requires  
 361 consolidation in post-task sleep *Prince et al. (2014); Inostroza et al. (2013)*. However, while this  
 362 task involves spatial exploration of objects in an arena, whether the subsequent post-task sleep  
 363 contains sequential structure and whether object memory is contained in this code has remained  
 364 elusive (*Larkin et al., 2014*). In order to assess the presence of sequential structure in the PBEs,  
 365 we first used cross validation to generate a distribution of sequence HMM-congruence scores. For  
 366 each set of test PBEs, we also generated surrogates by shuffling time bins across events (pooled  
 367 time-swap). Using our HMM-congruence score which explicitly tests for sequences through state  
 368 space, the large difference between actual and shuffled score distributions indicates evidence for  
 369 significant sequential structure in the PBEs ( $p < 0.001$ , Mann-Whitney  $U$  test, **Figure 8**). While more  
 370 work is needed to evaluate the mnemonic relevance of these HMM-congruent sequences, these  
 371 data support the notion that the HMM can uncover sequential activity in sleep away from the task  
 372 environment. This approach further demonstrates the utility of the HMM approach as an initial  
 373 analysis of a novel dataset, or as a way of comparing the sequential content encoded in PBEs during  
 374 different periods.

### 375 Discussion

376 Increasing lines of evidence point to the importance of hippocampal ensemble activity during PBEs  
 377 in guiding on-going behavior and active learning. Despite being the strongest output patterns of  
 378 the hippocampus, however, this activity has been assumed to be interpretable only in the context  
 379 of other theta-associated place cell activity expressed during behavior. Our findings demonstrate  
 380 that over the course of a behavioral session, ensemble activity during PBEs alone is sufficient to  
 381 form a model which captures the spatial relationships within an environment. This suggests that  
 382 areas downstream of the hippocampus might be able to make use solely of PBE activity to form

383 models of external space. In an extreme view, place cell activity might merely subserve the internal  
384 mechanisms in the hippocampus which generate PBE sequences. To the extent that animals might  
385 wish to use the spatial code obtained from PBEs to identify their current location, we show that  
386 this can be done after translating ensemble activity into the latent states of the model. Do the  
387 PBEs contain “full information” about the environment? Bayesian decoding of location from place  
388 cell activity results in lower error than location estimates generated using the latent states and  
389 IsPFs. This suggests that the manifold defined by the HMM may not capture all the dimensions of  
390 information represented during exploration. However, it is possible that with more PBE data, we  
391 would learn a more refined state space. Thus, the difference between the latent space represented  
392 during behavior and within PBEs may be an interesting focus of future study.

393 When we examined the transition matrices we learned from PBEs, we found that they were  
394 marked by significant sparsity. This sparsity results from the sequential patterns generated during  
395 PBEs. Latent variable models have previously been used to analyze the structure of hippocampal  
396 place cell activity (*Chen et al., 2012, 2014; Dabaghian et al., 2014*). In these studies, the learned  
397 transition matrices were mapped to undirected graphs which could be analyzed using topological  
398 measures. It is intriguing that similar structure is apparent in PBE activity. For example, we observed  
399 that transition matrices learned from PBEs associated with linear track behavior were significantly  
400 sparser than those learned from the open field, which we hypothesize is a consequence of the  
401 greater freedom of behavior in the latter (a topological difference). Whether hippocampal PBE  
402 activity must always be sequential, i.e., evolve through a sparsely-connected latent space, is an  
403 open and interesting question, as are differences between the latent state space dynamics learned  
404 during PBEs and those learned from place cell activity.

### 405 **Graded, Non-binary Replay Detection**

406 Remarkably, evaluating the congruence or likelihood of test data against our HMM provided a highly  
407 novel method to detect events that are consistent with replay, without a need to access the “play”  
408 itself. In the process of evaluating the potential of HMMs for detecting replay, we developed an  
409 approach to compare different replay-detection strategies. Our results highlight how the data does  
410 not readily admit to a strict separation between “replay” and “non-replay” events. While it is possible  
411 that with additional shuffles or other restrictions (*Silva et al., 2015*), automated performance  
412 might be rendered closer to human-labeling, even human scorers had variation in their opinions.  
413 This calls into doubt judgments of memory-related functions which build on a binary distinction  
414 between replay and non-replay sequences. Model congruence, either as a raw statistical likelihood  
415 or weighted against a shuffle distribution, seems to be a very reasonable metric to associate  
416 with individual PBEs. Moreover, evaluating congruence with an HMM does not require access to  
417 repeated behavioral sequences, which may be infeasible under widely-used single- or few-trial  
418 learning paradigms or when the events involve replay of a remote internalized environment. Given  
419 these benefits, along with computational efficiency, we would suggest that future analyses of the  
420 downstream impact of hippocampal reactivation regress effects against this measure rather than  
421 assuming a binary distinction.

### 422 **Learning, Model Congruence and Replay Quality**

423 Not surprisingly, the rate of PBEs had a large effect on our ability to measure model congruence.  
424 Interestingly, it has been noted that the density of PBEs is higher during early exposure to a  
425 novel environment (*Cheng and Frank, 2011; Frank et al., 2004; Kemere et al., 2013; Kudrimoti*  
426 *et al., 1999*). This might suggest that for the animal, PBE activity could be an important source  
427 for generating models of the world when the animal is actively learning about the environment.  
428 If as hypothesized, replay is a form of rehearsal signal generated by the hippocampus to train  
429 neocortical modules (*McClelland et al., 1995; Buzsáki, 1989*), then indeed the brain’s internal  
430 machinery may also be evaluating whether a given sequential PBE pattern is congruent and  
431 consistent with previously observed PBEs. In later sessions, as animals have been repeatedly

432 exposed to the same environments, downstream regions will have already witnessed many PBEs  
 433 from which to estimate the structure of the world. Overall, our approach provides a novel viewpoint  
 434 from the perspective of hippocampal PBEs. An interesting future line of inquiry would be to assess  
 435 the extent to which a model built on PBEs during first experience of a novel environment is slower  
 436 or faster to converge to the final spatial map than models built on theta-associated place activity.

### 437 **Application to Extra-spatial Behaviors**

438 We have analyzed data gathered in experiments in which rats carried out simple spatial navigation  
 439 tasks. Thus, to some extent it is not surprising that when we decoded ensemble activity during  
 440 behavior we found that spatial positions the animal is exploring are strongly associated with the  
 441 latent states.

442 We anticipate that our approach for calculating IsPFs would be equally useful in tasks in which  
 443 the hippocampal map is organized around time (*Eichenbaum, 2014; Rodriguez and Levy, 2001*) or  
 444 other continuous variables (e.g. sound frequency (*Aronov et al., 2017*)). Our two proof-of-concept  
 445 analyses, however, suggest that it should be possible to use HMMs to infer the presence of extra-  
 446 spatial sequential reactivation in PBEs. For example, we showed that there is significant sequential  
 447 structure during sleep after an animal explores novel objects in an environment. We anticipate that  
 448 careful experimental design and further algorithmic development would allow for the conjunctive  
 449 coding of object identity and spatial locations to be detected in the latent states we learn from PBEs,  
 450 with model-congruence providing a tool to study sequential hippocampal reactivation in these  
 451 types of tasks.

452 Conjunctive, non-spatial information might be one source of the apparent variability that results  
 453 in many PBEs not being detected as replay using traditional Bayesian decoding. Another proposed  
 454 source of this variability is reactivation of other environments. Our second proof-of-concept analysis  
 455 suggests that HMMs learned from PBEs can, in fact, capture the spatial structure of environments  
 456 beyond the one the animal is currently exploring. It appears that it should be possible to use only  
 457 the PBEs and information about the place-cell map of the local environment to refine the training  
 458 set for remote replay activity and learn the structure of a remote environment. While we used  
 459 Bayesian decoding to detect putative local replays, we anticipate related approaches might use an  
 460 HMM or other approaches to model local place cell activity.

### 461 **Future possibilities**

462 It has been previously observed that the rate of hippocampal reactivations in PBEs during awake  
 463 behavior is much higher than during sleep (*Grosmark and Buzsáki, 2016; Karlsson and Frank,*  
 464 *2008*), but the reasons for this are not well understood. One hypothesis is that many sleep PBEs  
 465 contain the reactivation of contexts other than those measured during a behavioral experiment.  
 466 Another hypothesis is that sleep activity involves remodeling of dynamic network architectures  
 467 (*Buhry et al., 2011; Tononi and Cirelli, 2014*). Our approach has the potential to illuminate some  
 468 sources of variability during sleep. While we have given preliminary evidence that information  
 469 about a remote context can be present in PBEs along with the local context, further work is required  
 470 to understand how our model's ability to capture this structure scales with the number of different  
 471 contexts. With sufficient data, our HMM approach should be able to learn disjoint sets of latent  
 472 states (or "sub-models") which would capture these separate contexts and allow us to test this  
 473 possibility. Alternatively, sleep PBEs could yield models which represent a known behavioral context  
 474 but are markedly different (e.g., less sparse) than those learned from awake PBEs. This might  
 475 support the network remodeling function of sleep. In the latter case, we might imagine that only a  
 476 small subset of sleep PBEs—corresponding to learning-related replay—would be congruent with a  
 477 model learned from awake PBE data.

## 478 **Conclusions**

479 We have demonstrated a new analytical framework for studying hippocampal ensemble activity  
 480 which enables primacy of PBEs in model formation. We use an unsupervised learning technique  
 481 commonly used in the machine learning field to study sequential patterns, the hidden Markov  
 482 model. This contrasts with existing approaches in which the model—estimated place fields for  
 483 the ensemble—is formed using the theta-associated place cell activity. We find that our PBE-first  
 484 approach results in a model which still captures the spatial structure of the behavioral tasks we  
 485 studied. Additionally, we demonstrate that we can use model-congruence as a tool for assessing  
 486 whether individual PBEs contain hippocampal replay. Finally, we present proofs-of-concept that this  
 487 analytical approach can detect the presence of sequential reactivation in experimental scenarios in  
 488 which existing approaches are insufficient. Thus, the use of unsupervised learning of latent variable  
 489 models—specifically HMMs and statistical congruence as a marker of individual events—bears  
 490 much promise for expanding our ability to understand how PBEs enable the cognitive functions of  
 491 the hippocampus.

## 492 **Materials and Methods**

### 493 **Experiment paradigm/Neural data recording**

494 We analyzed neural activity recorded from the hippocampus of rats during periods in which they  
 495 performed behavioral tasks in different paradigms. First, we considered data from animals running  
 496 back and forth in a linear track 150 or 200 cm long. As previously reported using these same data  
 497 (*Diba and Buzsáki, 2007*), we recorded neural activity using chronically-implanted silicon probes to  
 498 acquire the activity of hippocampal CA1/CA3 neurons. From these experiments, we chose sessions  
 499 during which we observed at least 20 place cells during active place-field exploration, and at least  
 500 30 PBEs (see below). Place cells were identified as pyramidal cells which had (i) a minimum peak  
 501 firing rate of 2 Hz, (ii) a maximum mean firing rate of 5 Hz, and (iii) a peak-to-mean firing rate ratio  
 502 of at least 3, all estimated exclusively during periods of run (as defined before, that is, when the  
 503 animal was running > 10 cm/s). This selection yielded  $n = 18$  session with 41–203 neurons (36–186  
 504 pyramidal cells). All procedures were approved by the Institutional Animal Care and Use Committee  
 505 of Rutgers University and followed US National Institutes of Health animal use guidelines (protocol  
 506 90-042).

507 A second dataset used tetrodes to record a large number (101–242) of putative pyramidal  
 508 neurons in area CA1 during two sessions each in four rats. Briefly, as was previously reported  
 509 using these data (*Pfeiffer and Foster, 2013, 2015*), rats explored an arena in which there were 36  
 510 reward sites. In each session, one site was designated as “home”. During a session, rats would  
 511 repeatedly alternate between retrieving a random reward site in one of the remaining 35 locations  
 512 and retrieving a reward at the home location. All procedures were approved by the Johns Hopkins  
 513 University Animal Care and Use Committee and followed US National Institutes of Health animal  
 514 use guidelines (protocols RA08M138, RA11M16, and RA14M48).

515 In order to investigate remote replay, we used data from an experiment in which this phe-  
 516 nomenon has been previously reported (*Karlsson and Frank, 2009*). Briefly, rats were implanted  
 517 with multi-electrode microdrives with tetrodes targeting CA1 and CA3. They were trained to carry  
 518 out a continuous-alternation task in an initially novel “w”-shaped maze (E2) for liquid reward for  
 519 multiple daily run sessions interspersed by rest-periods in an enclosed box. After they learned the  
 520 task, they were introduced to a novel w-maze (E1) in a different orientation in which they had two  
 521 run sessions followed by a run in the now-familiar E2. For our proof-of-concept analysis (*Figure 7*),  
 522 we used data from the second day of the novel maze (i.e., third and fourth exposures) in animal  
 523 ‘Bon’.

524 Finally, we recorded neural activity during an object-location memory task using a 32-channel  
 525 silicon probes (Buzsaki32, Neuronexus, MI) equipped with light fibers lowered to area CA1 of the  
 526 dorsal hippocampus. The animal was previously infused with AAV-CamKIIa-ArchT-GFP for the

527 purpose of another experiment. Putative pyramidal cells and interneurons were distinguished  
 528 based on their spike waveforms and spike auto-correlograms. On the day before the recordings,  
 529 the animal was repeatedly exposed to an empty test chamber on four successive six minute blocks,  
 530 interleaved by three minute rest periods in the home cage. On the recording day, the first of these  
 531 six-minute blocks was again the empty test chamber, but on the remaining blocks, the animal was  
 532 exposed to a fixed configuration of three different novel objects placed in the northeast, center  
 533 and southeast corners of the box. These blocks were again interleaved with three minute rest  
 534 periods in the home cage. The test chamber was a 60 × 60 cm<sup>2</sup> box with a local cue (8.5 in. ×  
 535 11 in. sheet printout) placed on one test wall. Following the last acquisition exposure, the animal  
 536 was returned to its home cage for a four hour extended sleep period. The subsequent day, one  
 537 of the objects in the box was displaced and the animal was reintroduced into the box to test for  
 538 interactions with the displaced versus non-displaced objects. All procedures were approved by  
 539 the Institutional Animal Care and Use Committee of the University of Wisconsin-Milwaukee and  
 540 followed US National Institutes of Health animal use guidelines (protocol 13-14 #28)

### 541 Population burst events

542 To identify PBEs in the linear track data, a spike density function (SDF) was calculated by counting  
 543 the total number of spikes across all recorded single and multi-units in non-overlapping 1 ms  
 544 time bins. The SDF was then smoothed using a Gaussian kernel (20 ms standard deviation, 60 ms  
 545 half-width). Candidate events were identified as time windows with a peak SDF of at least three  
 546 standard deviations above the mean calculated over all the session. The boundaries of each event  
 547 were set to time points of crossing the mean, preceding and following the peak. Events during  
 548 which animals were moving (average movement speed of > 5 cm/s) were excluded from all further  
 549 analyses to prevent possible theta sequences from biasing our results. For analysis, we then binned  
 550 each PBE into 20 ms (non-sliding) time bins. Spikes from putative interneurons (mean firing rate  
 551 when moving > 10 Hz) were excluded, as were events with duration less than four time bins or with  
 552 fewer than four active pyramidal cells. For the open field data, we used the previously reported  
 553 criteria (*Pfeiffer and Foster, 2013*) for identifying PBEs prior to binning (10 ms standard deviation  
 554 kernel, minimum of 10% of units active, duration between 50 ms and 2000 ms).

### 555 Hidden Markov model of PBE activity

556 We trained HMMs on the PBEs. In an HMM, an unobserved discrete latent state  $q_t$  evolves through  
 557 time according to a first order Markov process. The temporal evolution of the latent state is  
 558 described by the  $M \times M$  matrix  $\mathbf{A}$ , whose elements  $\{a_{ij}\}$  signify the probability after each time  
 559 bin of transitioning from state  $i$  to state  $j$ ,  $a_{ij} = \Pr(q_{t+1} = j | q_t = i)$ . The number of states,  $M$ , is a  
 560 specified hyperparameter. We found that our results were insensitive to the value of  $M$  through a  
 561 wide range of values from 20 to 100 (*Figure 3–Figure Supplement 1*). During each time bin of an  
 562 event, the identity of the latent state influences what is observed via a state-dependent probability  
 563 distribution. We modeled the  $N$ -dimensional vector of binned spiking from our ensemble of  $N$   
 564 neurons at time  $t$ ,  $O_t$ , as a Poisson process. Specifically, for each state,  $i$ , we model neuron  $n$  as  
 565 independently firing according to a Poisson process with rate  $\lambda_{ni}$ .

$$\Pr(O_t | q_t = i) = \prod_{n=1}^N \Pr(o_{nt} | q_t = i) \propto \prod_{n=1}^N (\lambda_{ni})^{o_{nt}} \exp(-\lambda_{ni})$$

566 where  $o_{nt}$  is the number of spikes observed from neuron  $n$  at time  $t$ . The final parameter which  
 567 specifies our model is the probability distribution of the initial state for a given event,  $\pi_i = \Pr(q_1 = i)$ .  
 568 Thus, our model is specified by parameters  $\theta = \{\mathbf{A}, \mathbf{\Lambda}, \boldsymbol{\pi}\}$ , where  $\mathbf{\Lambda} = \{\lambda_{ni}\}$  is an  $N \times M$  matrix and  
 569  $\boldsymbol{\pi} = \{\pi_i\}$  is an  $M$ -dimensional vector.

570 To learn model parameters, we follow the well-known iterative EM procedure (*Rabiner, 1989*),  
 571 treating each training PBE as an observation sequence. In order to regularize the model, we impose  
 572 a minimum firing rate for each neuron of 0.001 (0.05 Hz) during the M-step of EM. For a given PBE



573 (i.e., observation sequence) with  $K$  bins, we use the “forward-backward algorithm” (*Rabiner, 1989*)  
 574 to calculate the probability distribution of the latent state for each time bin,  $\Pr(q_t|O_1, \dots, O_t, \dots, O_K)$ .  
 575 For a particular time bin,  $t$ , in a given sequence, the forward-backward algorithm allows information  
 576 from all observation bins, previous and subsequent, to affect this state probability distribution  
 577 (as well the observation bin at time  $t$ ). The forward-backward algorithm also efficiently calculates  
 578 the “score”, or likelihood of the complete sequence,  $\Pr(O_1, \dots, O_K)$ . All HMMs learned in this work  
 579 used five-fold cross validation, i.e., the PBEs were divided into five randomly selected fifths (“folds”),  
 580 and then each fold was evaluated as a test set, with the model trained using the remaining four  
 581 folds. We define the model likelihood of an HMM as the product of the scores of each event using  
 582 this five-fold cross validation. To initially evaluate model learning, we compared model likelihoods  
 583 calculated using real and shuffled test data. Models which have learned to properly represent the  
 584 data should show significant increases. To quantify the presence of PBE sequences in a model we  
 585 used a model quality metric as described below.

### 586 **Ordering states for visualization**

587 For visualization, we wanted to order the states to maximize the super diagonal of the transition  
 588 matrix. We used a greedy approach which typically yields this solution. We started by assigning  
 589 the first index to the state with the highest initial probability and added states based on the most  
 590 probable state transitions. The undirected connectivity graphs were then generated from this  
 591 transition matrix, averaging the strength of reciprocal connections,  $a_{ij}$  and  $a_{ji}$ .

### 592 **Surrogate datasets and shuffle methods**

593 In order to analyze the HMMs we learned, we compared them against different types of surrogate  
 594 datasets obtained by shuffling the neural activity during PBEs in distinct ways. 1) **Temporal shuffle**:  
 595 within each event, the binned spiking activity was circularly permuted across time for each unit,  
 596 independently of other units. This goal of this shuffle is to disrupt unit co-activation, while maintain-  
 597 ing the temporal dynamics for each unit. 2) **Time-swap shuffle**: within each event, the order of the  
 598 binned columns of neural activity was randomly permuted across time, coherently across units. The  
 599 goal of this shuffle is to change the temporal dynamics of ensemble activity, while maintaining unit  
 600 co-activation. 3) **Poisson surrogate “shuffle”**: we estimated each unit’s mean firing rate across all  
 601 PBEs, and then produced surrogate PBEs from independent Poisson simulations according to each  
 602 unit’s mean firing rate. 4) **Pooled time-swap shuffle**: the order of the binned columns of neural  
 603 activity was randomly permuted across all pooled events, coherently across units. This shuffle has  
 604 been previously used in Bayesian replay detection (*Davidson et al., 2009*).

### 605 **Calculating sparsity and connectivity of the model parameters**

606 Sparsity of the transitions from individual states (departure sparsity) was measured by calculating  
 607 the Gini coefficient of corresponding rows of the transition matrix (*Hurley and Rickard, 2009*). The  
 608 Gini coefficient is a measure of how variable the values of this probability distribution are, with  
 609 equality across states corresponding to a coefficient of zero (minimal sparsity), and a singular  
 610 distribution with a probability-one transition to a single other state corresponding to a coefficient  
 611 of one (maximal sparsity). The sparsity of the full transition matrix was calculated by averaging  
 612 the Gini coefficient across rows. For analyses of PBE models from linear tracks, we computed the  
 613 mean sparsity across states for each of the 250 surrogate datasets, and these means were used to  
 614 generate the box plots of *Figure 2c*. Note that for the actual data, we generate a distribution by  
 615 randomly initializing the model 250 times and calculating the mean sparsity over all initializations.  
 616 For analyses of models learned from PBEs in open fields (and the linear track comparison with  
 617 50 states), we created 50 surrogates/random initializations (*Figure 6–Figure Supplement 1*). To  
 618 compare across sessions, we calculated the mean sparsity by averaging over all 250 surrogate  
 619 datasets to obtain a single mean sparsity per session, so that  $n = 18$  per-session means were used  
 620 to create the box-plots of *Figure 2e*.

621 Firing rates can be highly variable for different units. Thus, when evaluating the sparsity of  
 622 the observation matrix, we measured the extent to which individual units were specifically active  
 623 in a few states by calculating the Gini coefficients of the rows of the observation matrix. As with  
 624 transitions, we calculated mean sparsity across units for each surrogate dataset (e.g., linear track,  
 625 *Figure 2d*; open field, *Figure 6–Figure Supplement 2*), and we then averaged over all surrogate  
 626 datasets to obtain a per-session average, used in *Figure 2f*.

### 627 **Model connectivity and sequences**

628 To measure the degree of sequential connectivity within the graph corresponding to the transition  
 629 matrix—with nodes and edges representing the states and transitions, respectively—we developed  
 630 an algorithm for measuring the length of the longest path that can be taken through the graph. This  
 631 method is analogous to the “depth-first search” algorithm for traversing the graph’s tree structure  
 632 without backtracking. First, we made an adjacency matrix for a corresponding unweighted directed  
 633 graph by binarizing the transition matrix using a threshold of 0.2 on the transition probabilities.  
 634 Starting from each node, we then found the longest path that ended at either a previously visited  
 635 node or a terminal node (a node without any outgoing edges). To compare models trained on actual  
 636 versus surrogate datasets, we adjusted the thresholds to match the average degree (defined as the  
 637 average number of edges per node) between the models, thus ruling out possible effects due to  
 638 differences in the number of graph edges. We carried out this analysis on the same set of models  
 639 that were generated for analyzing sparsity. To compare across sessions, we calculated the median  
 640 maximum path length for each session ( $n = 18$ ) and used the per-session medians to generate box  
 641 plots of *Figure 2–Figure Supplement 3c*.

### 642 **Latent state place fields**

643 To calculate the latent state place fields, we first identified bouts of running by identifying periods  
 644 when animals were running (speed > 10 cm/s). We then binned the spiking during each of these  
 645 bouts in 100 ms bins. Using the forward-backward algorithm *Rabiner (1989)* and the HMM model  
 646 parameters learned from PBEs, we decoded each bout into a sequence of latent state probability  
 647 distributions,  $\Pr(q_t|O_t)$ . Using the track positions corresponding to each time bin, we then found the  
 648 average state distribution for each position bin,  $x_p$ , and normalized to yield a distribution for each  
 649 state,  $\Pr(x_p|q_t = i)$ .

### 650 **Decoding position from latent state sequences**

651 We used the IsPFs to decode the animal’s position after determining the probability of different  
 652 latent state trajectories during bouts of running. With five-fold cross validation, we estimated IsPFs  
 653 in a training dataset, then used the HMM model to decode latent state trajectory distributions from  
 654 ensemble neural activity in the test data. The product of IsPFs and decoded latent state distribution  
 655 at time  $t$  is the joint distribution  $\Pr(x_p, q_t|O_t)$ . We decode position as the mean of the marginal  
 656 distribution  $\Pr(x_p|O_t)$ .

### 657 **Bayesian Replay Detection**

658 We followed a frequently used Bayesian decoding approach to detect replay in our 1D data (*Kloost-*  
 659 *erman, 2012*). For each 20 ms time bin  $t$  within a PBE, given a vector comprised of spike counts  
 660 from  $N$  units,  $O_t = (o_{1t} o_{2t} \dots o_{Nt})$  in that bin, the posterior probability distribution over the binned  
 661 track positions was calculated using Bayes’ rule:

$$\Pr(x_p|O_t) = \frac{\Pr(O_t|x_p)\Pr(x_p)}{\sum_{q=1}^P \Pr(O_t|x_q)\Pr(x_q)},$$

662 where  $x_p$  is the center of  $p$ -th linearized position bin (of  $P$  total bins). We assumed Poisson firing  
 663 statistics, thus the prior probability,  $\Pr(O_i|x_p)$ , for the firing of each unit  $n$  is equal to

$$\Pr(O_i|x_p) = \prod_{n=1}^N \Pr(o_{ni}|x_p) \propto \prod_{n=1}^N (\tau \lambda_{np})^{o_{ni}} \exp(-\tau \lambda_{(n,p)})$$

664 where  $\tau$  is the duration of time bin (100 ms during estimation, 20 ms during decoding), and  $\lambda_{np}$   
 665 characterizes the mean firing rate of the  $n$ -th unit in the  $p$ -th position bin. We assumed a uniform  
 666 prior distribution  $\Pr(x_p)$  over the position bins.

667 For each PBE, the estimated posterior probability distribution was used to detect replay as  
 668 follows. Many (35,000) lines with different slopes and intercepts were sampled randomly following  
 669 the approach in (Kloosterman, 2012). The Bayesian replay score for a given event was the maximum  
 670 score obtained from all candidate lines, where the score for a particular line was defined as the  
 671 mean probability mass under the line, within a bandwidth (of 3 cm). For time bins during which the  
 672 sampled line fell outside of the extent of the track, the median probability mass of the corresponding  
 673 time bin was used, and for time bins during which no spikes were observed, we used the median  
 674 probability mass across all on-track time bins. To evaluate the significance of this score, for each  
 675 event we generated 5,000 surrogates of the posterior probability distribution by cycling the columns  
 676 (i.e., for each time bin, circularly permuting the distribution over positions by a random amount)  
 677 and calculated the replay score for each surrogate. The Monte Carlo  $p$ -value for each event was  
 678 obtained from the number of shuffled events with replay scores higher than the raw data. The  
 679 threshold for significance was varied as described in the text. For the open field, we used previously  
 680 reported criteria (Pfeiffer and Foster, 2013) to identify replay events from PBEs.

### 681 **Replay detection via PBE model congruence**

682 To identify replay as model congruence, for each PBE, we used the forward-backward algorithm to  
 683 calculate the sequence likelihood  $\Pr(O_1, \dots, O_K)$ , as defined earlier. Using five-fold cross validation,  
 684 the parameters of a HMM were learned from training PBEs. The sequence score was then calculated  
 685 for each event in the test data. To evaluate the significance of this score, for each event we generated  
 686 5,000 surrogate scores using a computationally-efficient scheme. Specifically, for each surrogate,  
 687 we randomly shuffle the rows of the transition matrix, excepting the diagonal. By maintaining the  
 688 diagonal (i.e., transitions that begin and end in the same state) and leaving the observation model  
 689 unchanged, this shuffle specifically selects against PBEs in which the latent states do not evolve  
 690 in temporal sequences. The Monte Carlo  $p$ -value for each event was calculated as the fraction of  
 691 shuffled events with HMM sequence scores higher than the raw data. The threshold for significance  
 692 was varied as described in the text. Note that while we describe this as HMM-congruence, we  
 693 have maintained the diagonal of the transition matrix, which specifically selects against PBEs  
 694 which might be model-congruent by maintaining a single state over many time bins. In reality  
 695 there are other dimensions of the HMM that we could assess congruence against, for example  
 696 the observation model, the initial state distribution, or combinations of these and the transition  
 697 matrix. In comparing against Bayesian decoding, our current definition seemed most appropriate  
 698 for sequence detection, but we can imagine future studies expanding on our approach.

### 699 **Human scoring and detection comparison**

700 We organized a group of human scorers to visually evaluate whether individual PBEs should  
 701 be described as replay. More specifically, scorers were only presented with Bayesian decoded  
 702 probability distributions such as those in **Figure 4a**, but without access to the spike raster or any  
 703 additional information. The scorers included six graduate students (including one of the authors)  
 704 and two undergraduates, all of whom were generally familiar with the concept of hippocampal  
 705 replay. We built an automatic presentation system which would display each event in random order,  
 706 and record one of six possible scores: “excellent” (highly sequential with no jumps and covering  
 707 most of the track), “good” (highly sequential with few or no jumps), “flat” (decoded position stayed

708 mostly in the same place, i.e. no temporal dynamics), “uncertain” (some semblance of structure,  
 709 but not enough to fall into any of the previous categories) or “noise” (no apparent structure, or  
 710 nonsensical trajectories such as teleportation). An event was then designated as replay if it was  
 711 labeled as “excellent” or “good” by a majority of scorers (ties were labeled as non-replay).

712 To calculate an ROC curve for replay detection algorithms, we used our shuffle statistics for each  
 713 event to create a vector which related the significance threshold (e.g., 99%) to the label supplied by  
 714 the algorithm (i.e., significant replay or not). Then, as a function of threshold, the sensitivity (fraction  
 715 of true positives identified) and selectivity (fraction of true negatives identified) were averaged  
 716 over events to yield an ROC curve. To evaluate whether the AUC differed between Bayesian and  
 717 model-congruence techniques we used a bootstrap approach. To generate a null hypothesis, we  
 718 combined the event/threshold vectors from both groups, and then sampled two random groups (A  
 719 and B) with replacement from the pooled data. The AUCs for these two random groups of events  
 720 were measured, and a distribution for the difference between the randomly chosen AUCs was  
 721 calculated. The two-sided  $p$ -value we report is the fraction of differences in random AUCs which are  
 722 more extreme than the actual difference.

### 723 HMM model quality across sessions

724 In order to understand the extent to which an HMM trained on PBEs from a given session contained  
 725 sequentially-structured temporal dynamics, we calculated the “session quality” (equivalently model  
 726 quality) as follows. Again using five-fold cross validation, we learn an HMM on the training subset of  
 727 PBEs, and score (using the forward-backward algorithm, as before), the remaining subset of test  
 728 PBEs. Then, we also score a pooled time-swap surrogate of the test PBEs and we repeat this pooled  
 729 time-swap scoring  $n = 2500$  times. Finally, we obtain a z score for each PBE by comparing the score  
 730 from the actual test PBE, to the distribution of pooled time-swap scores of the corresponding PBE.  
 731 The session quality is then defined as the average of these z scores, over all events in a session.  
 732 This measure of session quality was then used to detect the presence of putative remote replay  
 733 events or other extra-spatial structure in PBEs, since a high session quality after removing local  
 734 Bayesian significant events is highly suggestive of remaining (unexplained) sequential structure.

### 735 Software and data analysis

736 Data analyses were performed using MATLAB and Python. Jupyter notebooks (using Python) are avail-  
 737 able at <https://github.com/kemerelab/UncoveringTemporalStructureHippocampus>, where most of  
 738 the results presented here are reproduced. We have also developed and open-sourced a Python  
 739 package (namely `nelpy`) to support the analyses of electrophysiology data with HMMs, which is  
 740 available from <https://github.com/nelpy> (Ackermann *et al.*, 2018).

### 741 Acknowledgments

742 The authors would like to thank the reviewers for their valuable feedback, as well as MP Karlsson,  
 743 MF Carr, and LM Frank for making their w-maze data available on [crcns.org](http://crcns.org).

### 744 References


- 745 Ackermann E, Chu J, Dutta S, Kemere C, Nelpy: neuroelectrophysiology object model and data analysis in  
 746 Python; 2018. doi: 10.5281/zenodo.1219790. <https://github.com/nelpy>.
- 747 Aronov D, Nevers R, Tank DW. Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit.  
 748 Nature. 2017; 543(7647):719.
- 749 Buhry L, Azizi AH, Cheng S. Reactivation, replay, and preplay: how it might all fit together. Neural plasticity.  
 750 2011; 2011.
- 751 Buzsáki G. Hippocampal sharp waves: their origin and significance. Brain research. 1986; 398(2):242–252.
- 752 Buzsáki G. Two-stage model of memory trace formation: a role for “noisy” brain states. Neuroscience. 1989;  
 753 31(3):551–570.

- 754 **Buzsáki G.** Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hip-*  
755 *pocampus*. 2015; 25(10):1073–1188.
- 756 **Buzsáki G, Vanderwolf CH, et al.** Cellular bases of hippocampal EEG in the behaving rat. *Brain Research Reviews*.  
757 1983; 6(2):139–171.
- 758 **Carr MF, Jadhav SP, Frank LM.** Hippocampal replay in the awake state: a potential substrate for memory  
759 consolidation and retrieval. *Nature neuroscience*. 2011; 14(2):147–153.
- 760 **Chen Z, Gomperts SN, Yamamoto J, Wilson MA.** Neural representation of spatial topology in the rodent  
761 hippocampus. *Neural computation*. 2014; 26(1):1–39.
- 762 **Chen Z, Kloosterman F, Brown EN, Wilson MA.** Uncovering spatial topology represented by rat hippocampal  
763 population neuronal codes. *Journal of computational neuroscience*. 2012; 33(2):227–255.
- 764 **Chen Z, Wilson MA.** Deciphering Neural Codes of Memory during Sleep. *Trends in Neurosciences*. 2017; .
- 765 **Cheng S, Frank LM.** New experiences enhance coordinated neural activity in the hippocampus. *Neuron*. 2008;  
766 57(2):303–313.
- 767 **Cheng S, Frank LM.** The structure of networks that produce the transformation from grid cells to place cells.  
768 *Neuroscience*. 2011; 197:293–306.
- 769 **Chrobak JJ, Buzsáki G.** High-frequency oscillations in the output networks of the hippocampal–entorhinal axis  
770 of the freely behaving rat. *Journal of neuroscience*. 1996; 16(9):3056–3066.
- 771 **Dabaghian Y, Brandt VL, Frank LM.** Reconceiving the hippocampal map as a topological template. *Elife*. 2014;  
772 3:e03476.
- 773 **Davidson TJ, Kloosterman F, Wilson MA.** Hippocampal replay of extended experience. *Neuron*. 2009; 63(4):497–  
774 507.
- 775 **Deppisch J, Pawelzik K, Geisel T.** Uncovering the synchronization dynamics from correlated neuronal activity  
776 quantifies assembly formation. *Biological cybernetics*. 1994; 71(5):387–399.
- 777 **Diba K, Buzsáki G.** Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*.  
778 2007; 10(10):1241–1242.
- 779 **Diekelmann S, Born J.** The memory function of sleep. *Nature Reviews Neuroscience*. 2010; 11(2):114–126.
- 780 **Dupret D, O’neill J, Pleydell-Bouverie B, Csicsvari J.** The reorganization and reactivation of hippocampal maps  
781 predict spatial memory performance. *Nature neuroscience*. 2010; 13(8):995–1002.
- 782 **Ego-Stengel V, Wilson MA.** Disruption of ripple-associated hippocampal activity during rest impairs spatial  
783 learning in the rat. *Hippocampus*. 2010; 20(1):1–10.
- 784 **Eichenbaum H.** Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews*  
785 *Neuroscience*. 2014; 15(11):732.
- 786 **Foster DJ, Wilson MA.** Reverse replay of behavioural sequences in hippocampal place cells during the awake  
787 state. *Nature*. 2006; 440(7084):680–683.
- 788 **Frank LM, Stanley GB, Brown EN.** Hippocampal plasticity across multiple days of exposure to novel environments.  
789 *Journal of Neuroscience*. 2004; 24(35):7681–7689.
- 790 **Fruchterman TMJ, Reingold EM.** Graph Drawing by Force-directed Placement. *Software – Practice & Experience*.  
791 1991; 21(11):1129–1164.
- 792 **Girardeau G, Benchenane K, Wiener SI, Buzsáki G, Zugaro MB.** Selective suppression of hippocampal ripples  
793 impairs spatial memory. *Nature neuroscience*. 2009; 12(10):1222–1223.
- 794 **Grosmark AD, Buzsáki G.** Diversity in neural firing dynamics supports both rigid and learned hippocampal  
795 sequences. *Science*. 2016; 351(6280):1440–1443.
- 796 **Gupta AS, van der Meer MA, Touretzky DS, Redish AD.** Hippocampal replay is not a simple function of experience.  
797 *Neuron*. 2010; 65(5):695–705.

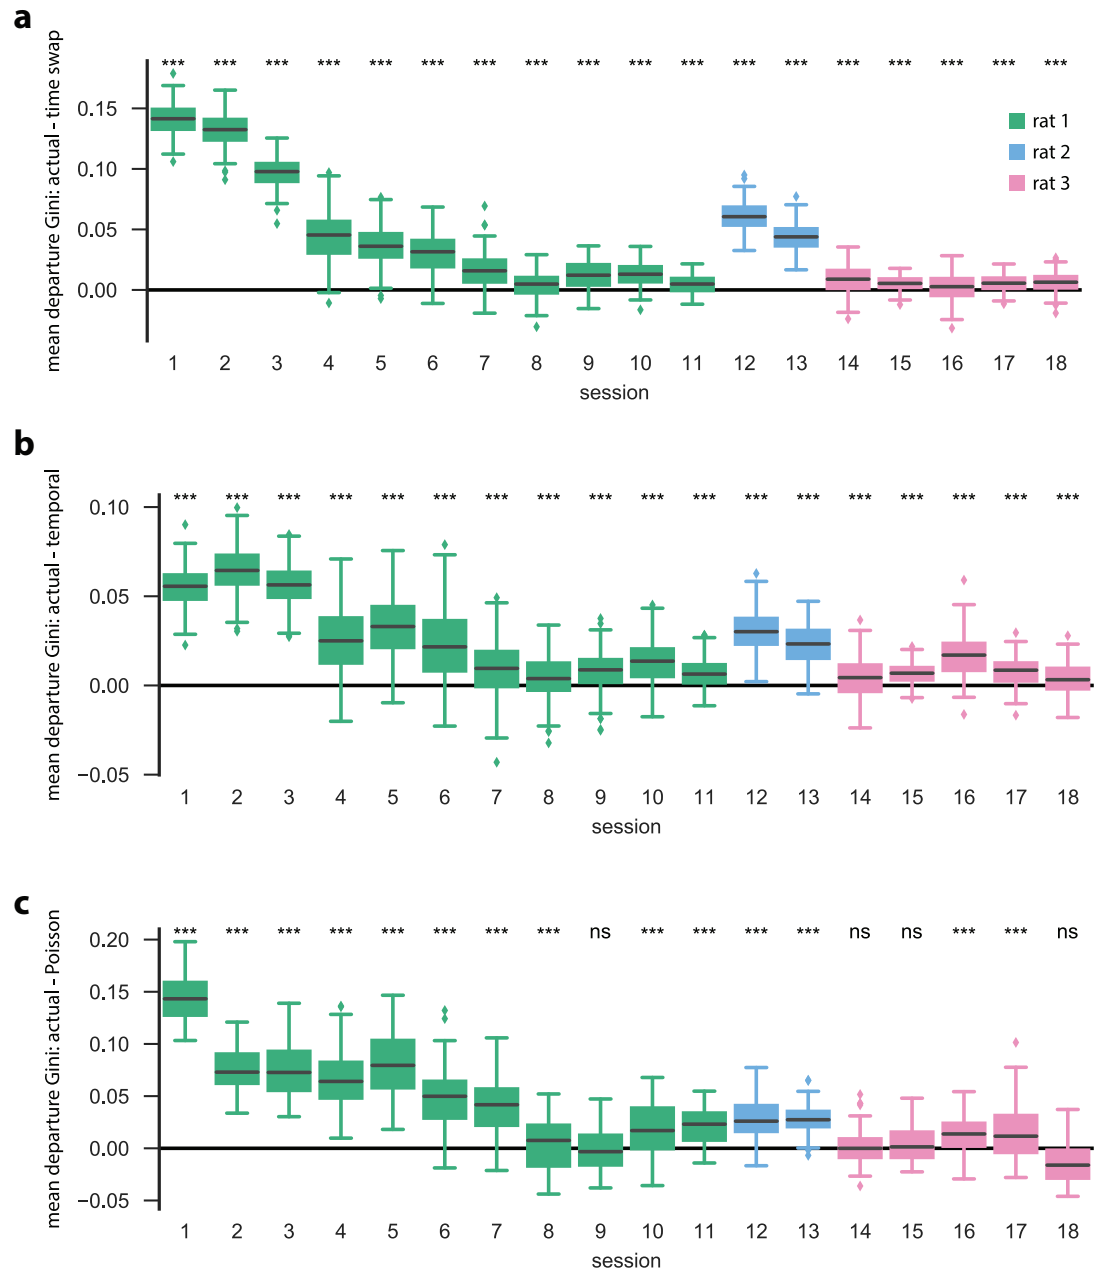


- 798 **Hurley N**, Rickard S. Comparing measures of sparsity. *IEEE Transactions on Information Theory*. 2009;  
799 55(10):4723–4741.
- 800 **Inostroza M**, Binder S, Born J. Sleep-dependency of episodic-like memory consolidation in rats. *Behavioural*  
801 *brain research*. 2013; 237:15–22.
- 802 **Jadhav SP**, Kemere C, German PW, Frank LM. Awake hippocampal sharp-wave ripples support spatial memory.  
803 *Science*. 2012; 336(6087):1454–1458.
- 804 **Karlsson MP**, Carr MF, Frank LM, Simultaneous extracellular recordings from hippocampal areas CA1 and CA3  
805 (or MEC and CA1) from rats performing an alternation task in two W-shaped tracks that are geometrically  
806 identically but visually distinct. *CRCNS.org*; 2015. doi: 10.6080/KONK3BZJ. <http://dx.doi.org/10.6080/KONK3BZJ>.
- 807 **Karlsson MP**, Frank LM. Network dynamics underlying the formation of sparse, informative representations in  
808 the hippocampus. *Journal of Neuroscience*. 2008; 28(52):14271–14281.
- 809 **Karlsson MP**, Frank LM. Awake replay of remote experiences in the hippocampus. *Nature neuroscience*. 2009;  
810 12(7):913–918.
- 811 **Kemere C**, Carr MF, Karlsson MP, Frank LM. Rapid and continuous modulation of hippocampal network state  
812 during exploration of new places. *PloS one*. 2013; 8(9):e73114.
- 813 **Kemere C**, Santhanam G, Byron MY, Afshar A, Ryu SI, Meng TH, Shenoy KV. Detecting neural-state transitions  
814 using hidden Markov models for motor cortical prostheses. *Journal of neurophysiology*. 2008; .
- 815 **Kloosterman F**. Analysis of hippocampal memory replay using neural population decoding. *Neuronal Network*  
816 *Analysis: Concepts and Experimental Approaches*. 2012; p. 259–282.
- 817 **Kudrimoti HS**, Barnes CA, McNaughton BL. Reactivation of hippocampal cell assemblies: effects of behavioral  
818 state, experience, and EEG dynamics. *Journal of Neuroscience*. 1999; 19(10):4090–4101.
- 819 **Larkin MC**, Lykken C, Tye LD, Wickelgren JG, Frank LM. Hippocampal output area CA1 broadcasts a generalized  
820 novelty signal during an object-place recognition task. *Hippocampus*. 2014; 24(7):773–783.
- 821 **Lee AK**, Wilson MA. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*.  
822 2002; 36(6):1183–1194.
- 823 **Logothetis NK**, Eschenko O, Murayama Y, Augath M, Steudel T, Evrard H, Besserve M, Oeltermann A.  
824 Hippocampal-cortical interaction during periods of subcortical silence. *Nature*. 2012; 491(7425):547–553.
- 825 **McClelland JL**, McNaughton BL, O'reilly RC. Why there are complementary learning systems in the hippocampus  
826 and neocortex: insights from the successes and failures of connectionist models of learning and memory.  
827 *Psychological review*. 1995; 102(3):419.
- 828 **Nádasy Z**, Hirase H, Czurkó A, Csicsvari J, Buzsáki G. Replay and time compression of recurring spike sequences  
829 in the hippocampus. *Journal of Neuroscience*. 1999; 19(21):9497–9507.
- 830 **O'Keefe J**. Place units in the hippocampus of the freely moving rat. *Experimental neurology*. 1976; 51(1):78–109.
- 831 **Pfeiffer BE**, Foster DJ. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*.  
832 2013; 497(7447):74–79.
- 833 **Pfeiffer BE**, Foster DJ. Autoassociative dynamics in the generation of sequences of hippocampal place cells.  
834 *Science*. 2015; 349(6244):180–183.
- 835 **Prince TM**, Wimmer M, Choi J, Havekes R, Aton S, Abel T. Sleep deprivation during a specific 3-hour time window  
836 post-training impairs hippocampal synaptic plasticity and memory. *Neurobiology of learning and memory*.  
837 2014; 109:122–130.
- 838 **Rabiner LR**. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings*  
839 *of the IEEE*. 1989; 77(2):257–286.
- 840 **Radons G**, Becker J, Dülfer B, Krüger J. Analysis, classification, and coding of multielectrode spike trains with  
841 hidden Markov models. *Biological cybernetics*. 1994; 71(4):359–373.
- 842 **Rodriguez P**, Levy WB. A model of hippocampal activity in trace conditioning: Where's the trace? *Behavioral*  
843 *Neuroscience*. 2001; 115(6):1224.

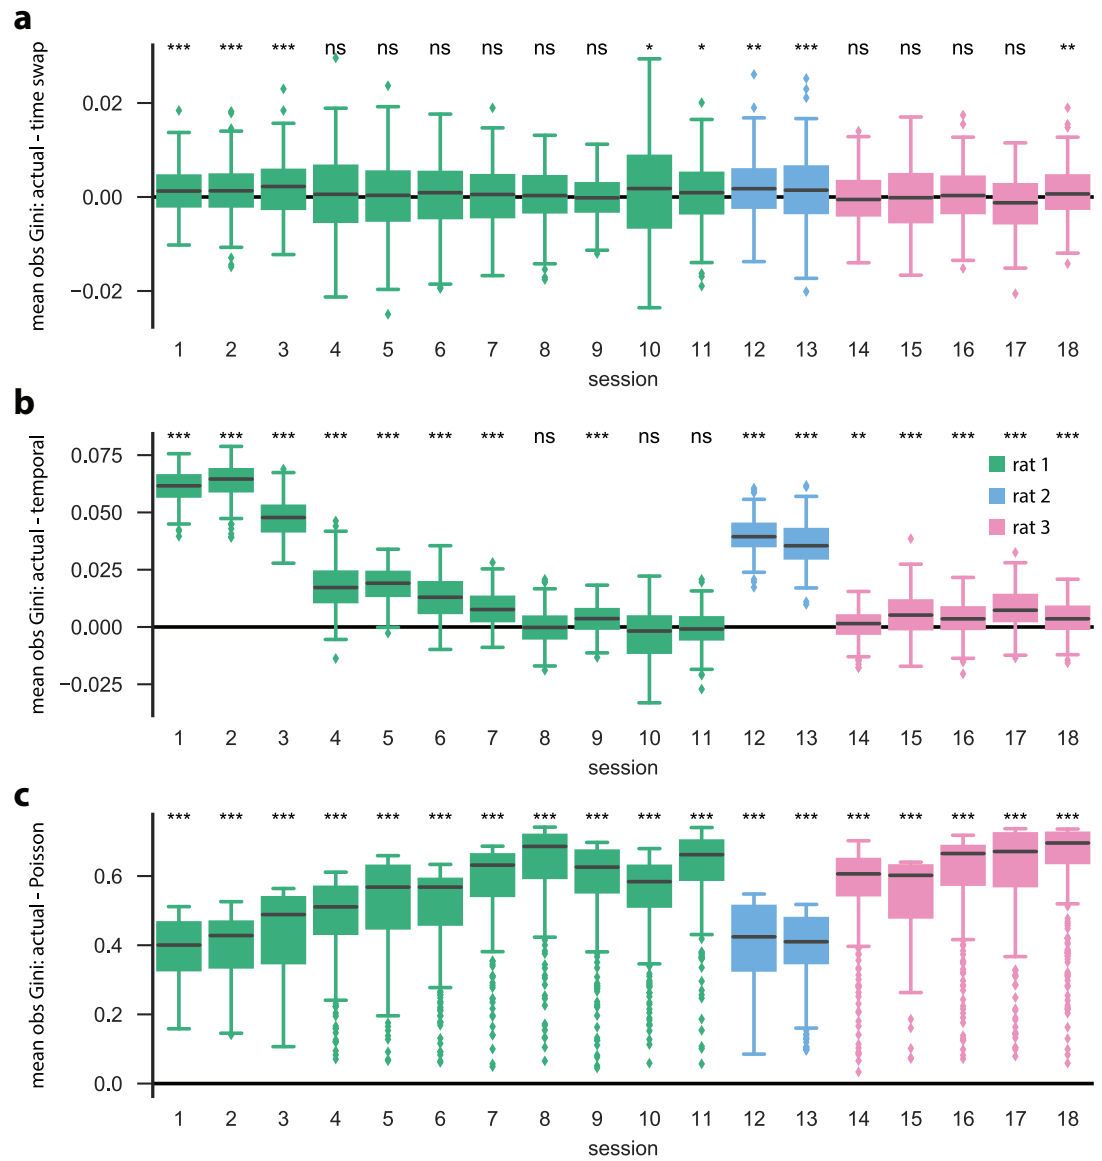
- 844 **Siapas AG**, Wilson MA. Coordinated interactions between hippocampal ripples and cortical spindles during  
845 slow-wave sleep. *Neuron*. 1998; 21(5):1123–1128.
- 846 **Silva D**, Feng T, Foster DJ. Trajectory events across hippocampal place cells require previous experience. *Nature*  
847 *neuroscience*. 2015; 18(12):1772–1779.
- 848 **Tononi G**, Cirelli C. Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory  
849 consolidation and integration. *Neuron*. 2014; 81(1):12–34.
- 850 **Yamamoto J**, Tonegawa S. Direct Medial Entorhinal Cortex Input to Hippocampal CA1 Is Crucial for Extended  
851 Quiet Awake Replay. *Neuron*. 2017; 96(1):217–227.



**Figure 1-Figure supplement 1.** Actual cross-validated test data and surrogate test data evaluated in actual-data-optimized HMMs for all 18 linear track sessions. For each session, we performed five-fold cross validation to score the validation (=test) set in an HMM that was learned on the corresponding training set. In addition, two surrogate datasets of the validation data (obtained by either temporal shuffle or time-swap shuffle) were scored in the same HMM as the actual validation data.  $k = 50$  shuffles of each event and of each type were performed. **a.** Difference between the data log likelihoods of actual and time-swap surrogate test events, evaluated in the actual train-data-optimized models. **b.** Same as in **a.**, except that the differences between the actual data and the temporal surrogates are shown. For each of the  $n = 18$  sessions, the actual test data had a significantly higher likelihood than either of the shuffled counterparts ( $p < 0.001$ , Wilcoxon signed-rank test). Sessions are arranged first by animal, and then by number of PBEs, in decreasing order.

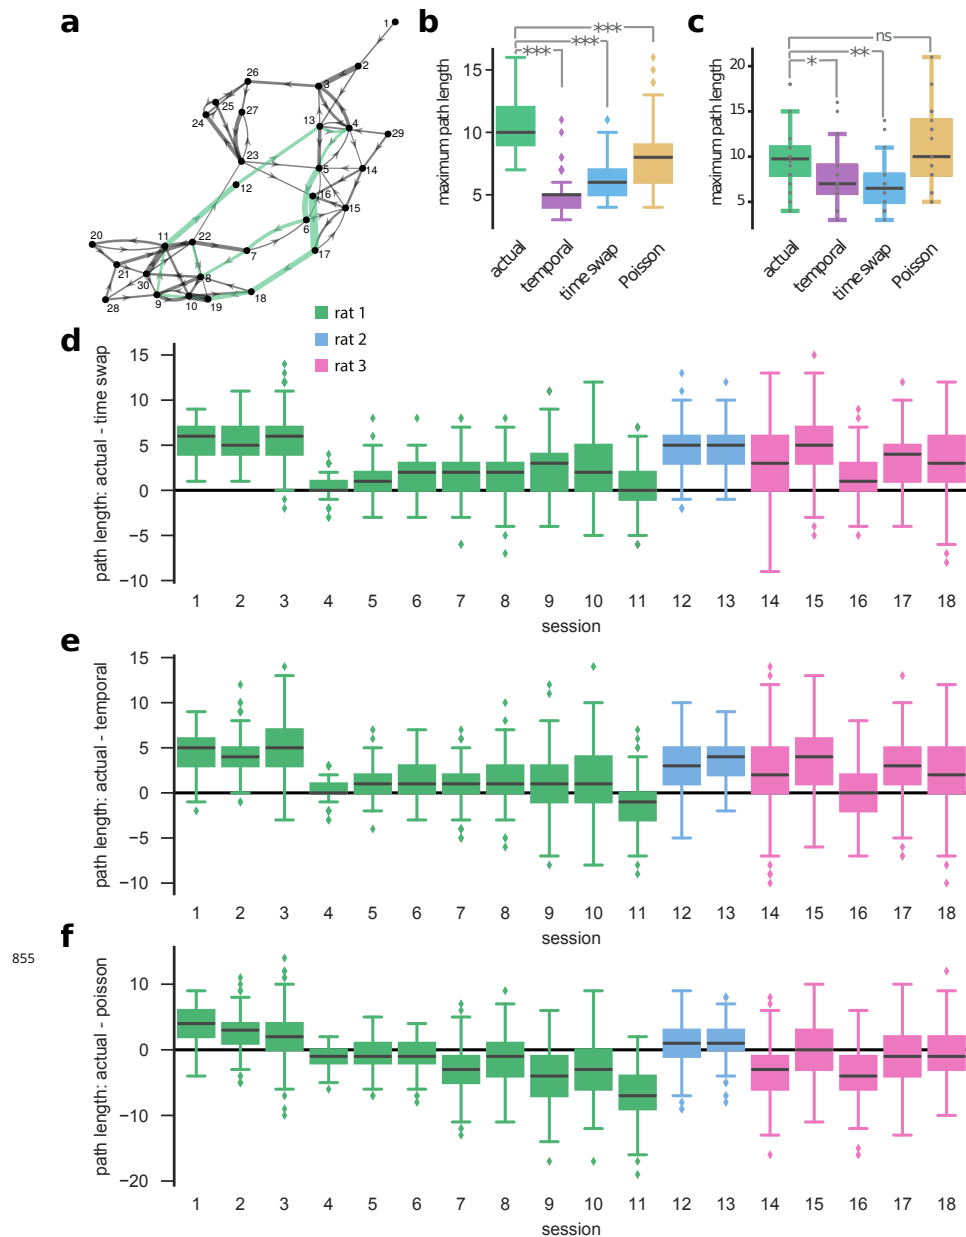


**Figure 2-Figure supplement 1.** We trained HMMs on neural activity during PBEs (in 20 ms bins), and asked how sparse the resulting state transitions were. In particular, we calculated the Gini coefficient for each row of our state transition matrix, so that the Gini coefficient for a particular row reflects the sparsity of state transitions from that state (row) to all other states (so-called “departure sparsity”). A high (close to one) Gini coefficient implies that the state is likely to only transition to a few other states, whereas a low (close to zero) Gini coefficient implies that the state is likely to transition to many other states. For each transition matrix, we computed the mean departure sparsity for  $n = 250$  initializations, and for  $n = 250$  shuffled counterparts for each of the surrogate datasets (**a.** time-swap shuffle, **b.** temporal shuffle, **c.** Poisson surrogate), and in each case we show the difference between the actual test data, and the surrogate test data. The actual data are significantly more sparse than both the temporal and time-swap surrogates for all sessions ( $p < 0.001$ , Mann-Whitney  $U$  test) and significantly more sparse than the Poisson surrogate for 14 of the 18 sessions ( $p < 0.001$ , Mann-Whitney  $U$  test).



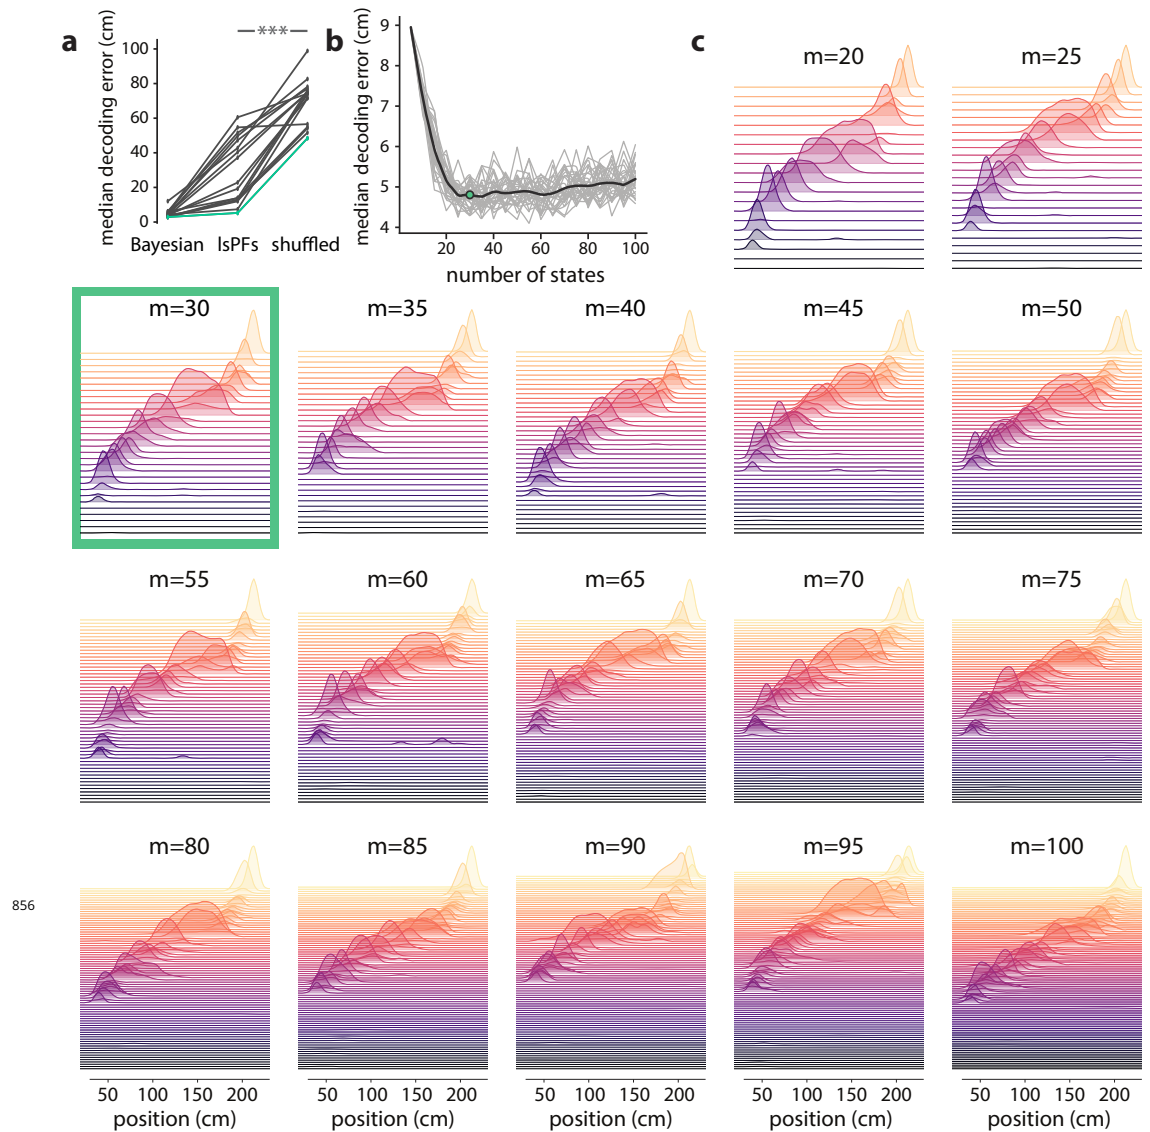
854

**Figure 2-Figure supplement 2.** Using the same PBE models and surrogate datasets ( $n = 250$  shuffles each) as in *Figure 2-Figure Supplement 1*, we investigated the sparse participation of neurons/units in our models by calculating the Gini coefficient of each row (that is, for each unit) of the observation matrix. A high Gini coefficient implies that the unit is active in only a small number of states, whereas a low Gini coefficient implies that the unit is active in many states. For each initialization / shuffle, we calculate the mean Gini coefficient over all units, and the differences between those obtained using actual data and those obtained using surrogate data are shown: differences between actual and **a.** time-swap, **b.** temporal, and **c.** Poisson surrogates. We find that the actual data are significantly more sparse than the temporal and Poisson surrogates for most of the sessions ( $p < 0.001$ , Mann-Whitney  $U$  test), but that for many (10 out of 18) sessions, there is no significant difference between the mean row-wise observation sparsity of the actual data compared to the time-swap surrogate. This is an expected result, since the time-swap shuffle leaves the observation matrix largely unchanged.



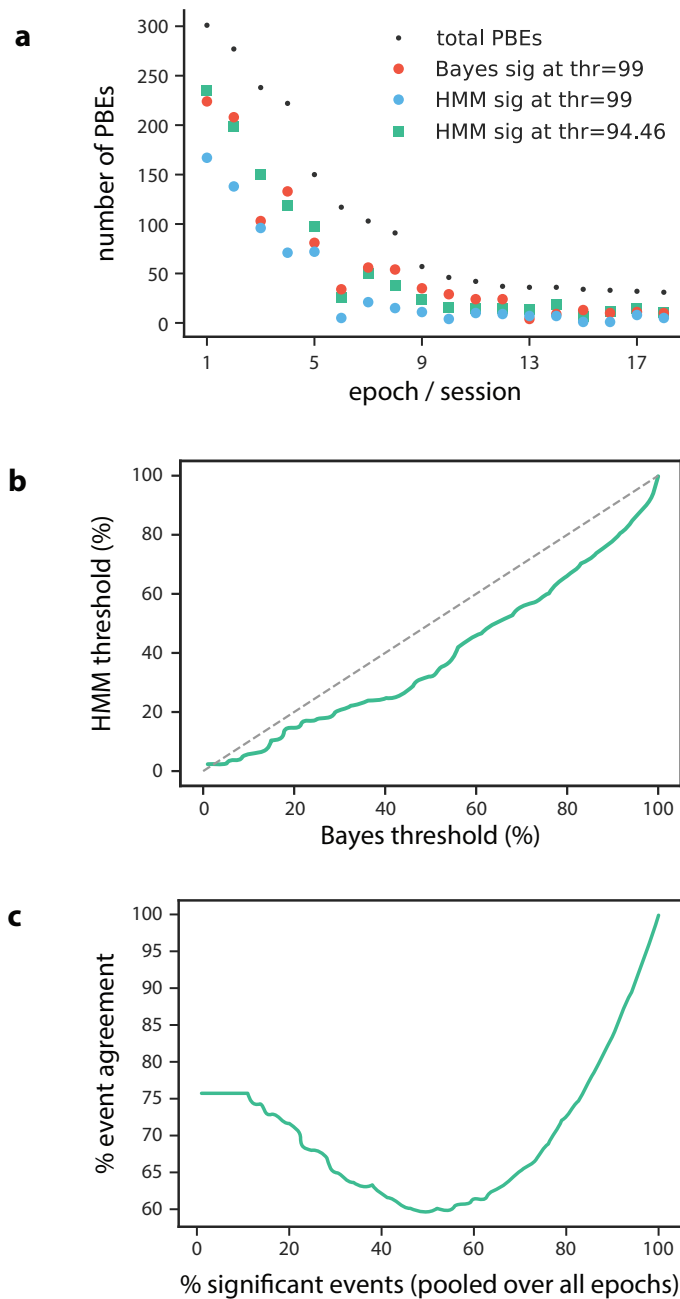
**Figure 2-Figure supplement 3.** We calculated the longest path within an unweighted directed graph corresponding to the transition matrices of HMMs, with nodes representing states and edges reflecting the transition probabilities (see Materials and Methods). **a.** The graph—displayed using the “force-directed layout” (*Fruchterman and Reingold, 1991*)—represents a model trained on actual data. For illustration purposes, we ignored transition probabilities below 0.1. The green path shows the longest path in the example. **b.** For this example session, we computed the maximum path length (the number of nodes in the longest path) for actual and corresponding shuffle datasets (temporal, time-swap, and Poisson) ( $n = 250$  initializations / shuffles). **c.** The panel shows aggregate results built of median maximum path lengths from all sessions. We find that the actual data results in longer paths compared to time-swap ( $p = 0.008$ , Mann-Whitney  $U$  test) and temporal surrogate datasets ( $p = 0.04$ , Mann-Whitney  $U$  test). On the contrary, no significant difference is found in comparison with the Poisson datasets ( $p = 0.57$ , Mann-Whitney  $U$  test). Nevertheless, due to non-sparseness of the observation matrix for a Poisson model (*Figure 2-Figure Supplement 2*), in most instances these paths correspond to highly overlapping ensemble sequences. In panels **d-f**, difference between maximum path lengths obtained from actual data and surrogate datasets are shown separately for all sessions : actual versus **d.** time-swap, **e.** temporal, and **f.** Poisson. The data results in longer paths compared to time-swap and temporal shuffle datasets in most sessions (15 out of 18) ( $p < 0.001$ , Mann-Whitney  $U$  test), though in only five sessions compared to Poisson surrogate datasets.





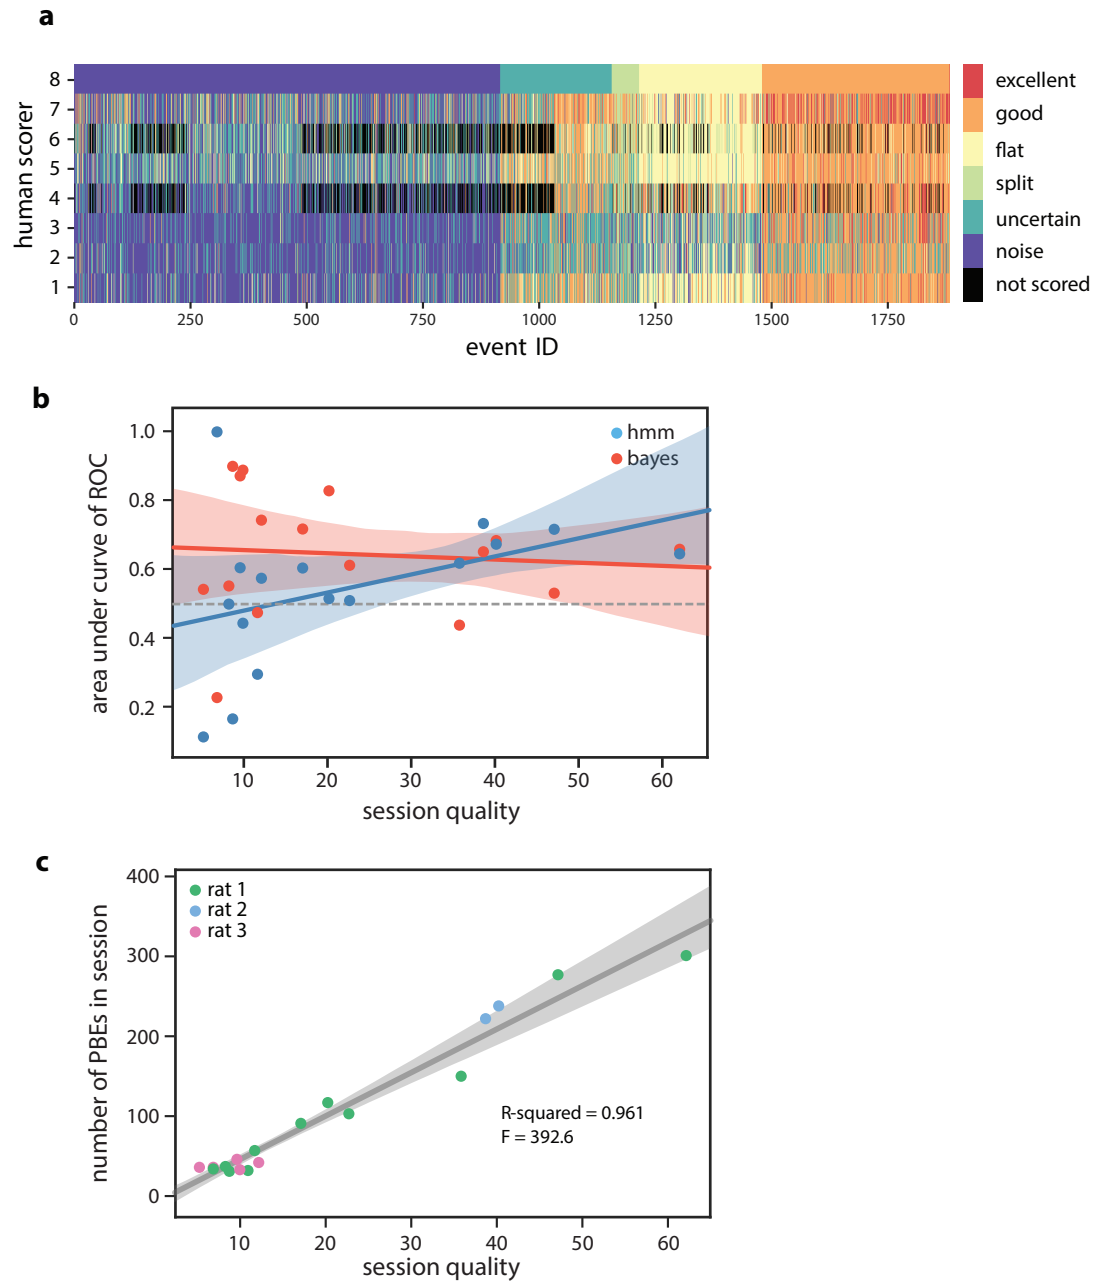
856

**Figure 3-Figure supplement 1.** We investigated to what extent our PBE models encoded information related to the animal's positional code by learning an additional mapping from the latent-state space to the animal's position (resulting in a latent-space place field, IsPF), and then using this mapping, we decoded run epochs to position and assessed the decoding accuracy. **a.** We computed the median position decoding accuracy (via the latent space) for each session on the linear track ( $n = 18$  sessions) using cross validation. In particular, we learned a PBE model for each session, and then using cross validation we learned the latent space to animal position mapping on a training set, and recorded the position decoding accuracy on the corresponding test set by first decoding to the state space using the PBE model, and then mapping the state space to the animal position using the IsPF learned on the training set. The position decoding accuracy was significantly greater than chance for each of the 18 sessions ( $p < 0.001$ , Wilcoxon signed-rank test). **b.** For an example session, we calculated the median decoding accuracy as we varied the number of states in our PBE model ( $n = 30$  realizations per number of states considered). Gray curves show the individual realizations, and the black curve shows the mean decoding accuracy as a function of the number of states. The decoding accuracy is informative over a very wide range of number of states, and we chose  $m = 30$  states for the analysis in the main text. **c.** For the same example session, we show the IsPFs for different numbers of states. The IsPFs are also informative over a wide range of number of states, suggesting that our analyses are largely insensitive to this particular parameter choice (the number of states). The coloration of the IsPFs is only for aesthetic reasons.



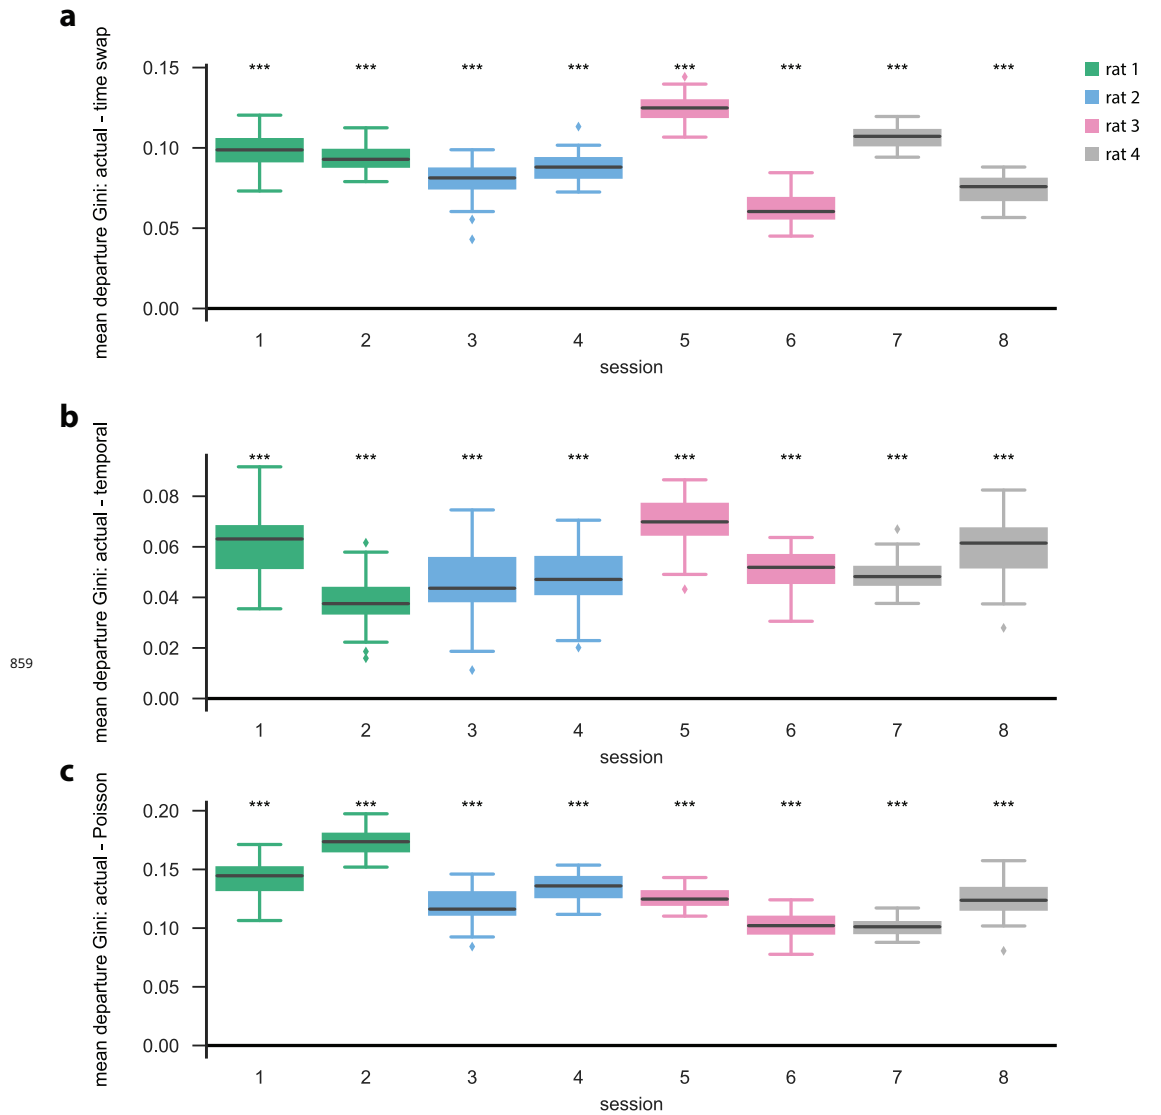
857

**Figure 4-Figure supplement 1. a.** The number of Bayesian significant PBEs, as well as the total number of PBEs are shown for each session ( $n = 18$ ) when using a significance threshold of 99%. We find that 57% of PBEs (1064 of 1883) are Bayesian significant at this threshold. When using this same threshold for the model-congruence (HMM) significance testing, we find that only 35% of PBEs (651 of 1883) are model congruent. In order to compare the Bayesian and model-congruence approaches more directly, we therefore lowered the model-congruence threshold to 94.46%, at which point both methods had the same number of significant events (1064 of 1883). **b.** For each Bayesian significance threshold, we can determine the corresponding model-congruence threshold that would result in the same number of significant PBEs. **c.** Using the thresholds from **b.** such that at each point, both Bayesian and model-congruence approaches have the same number of significant PBEs, we calculate the event agreement between the two approaches. We note that our chosen threshold of 57% significant events has among the worst agreement between the two approaches.

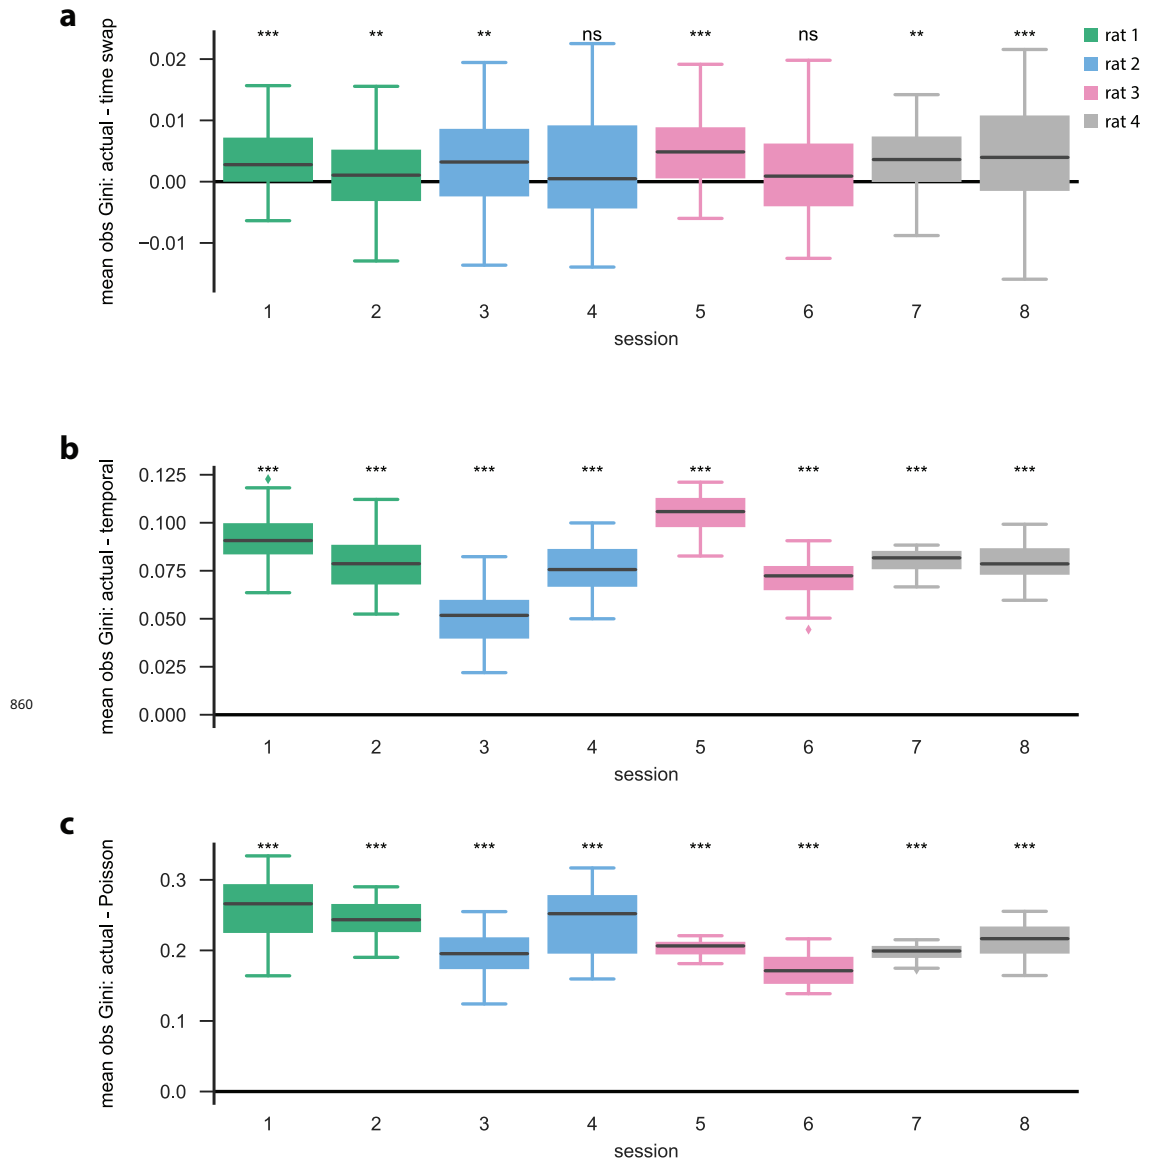


858

**Figure 5-Figure supplement 1. a.** Manual scoring results from 8 human scorers (six individuals scored  $n = 1883$  events, two individuals scored a subset of  $n = 1423$  events). Events were presented to each participant in a randomized order, and individuals were allowed to go back to modify their results before submission. Here, events are ordered according to individual #8's classifications. **b.** The model-congruence (HMM) approach appears to have higher accuracy when the session quality is higher ( $R^2 = 0.17$ ,  $F = 2.9$ ), which is consistent with our expectation that we need many congruent events in the training set in order to learn a consistent and meaningful model. **c.** The session quality is strongly correlated with the number of PBEs recorded within a session ( $R^2 = 0.96$ ,  $F = 392.6$ ).

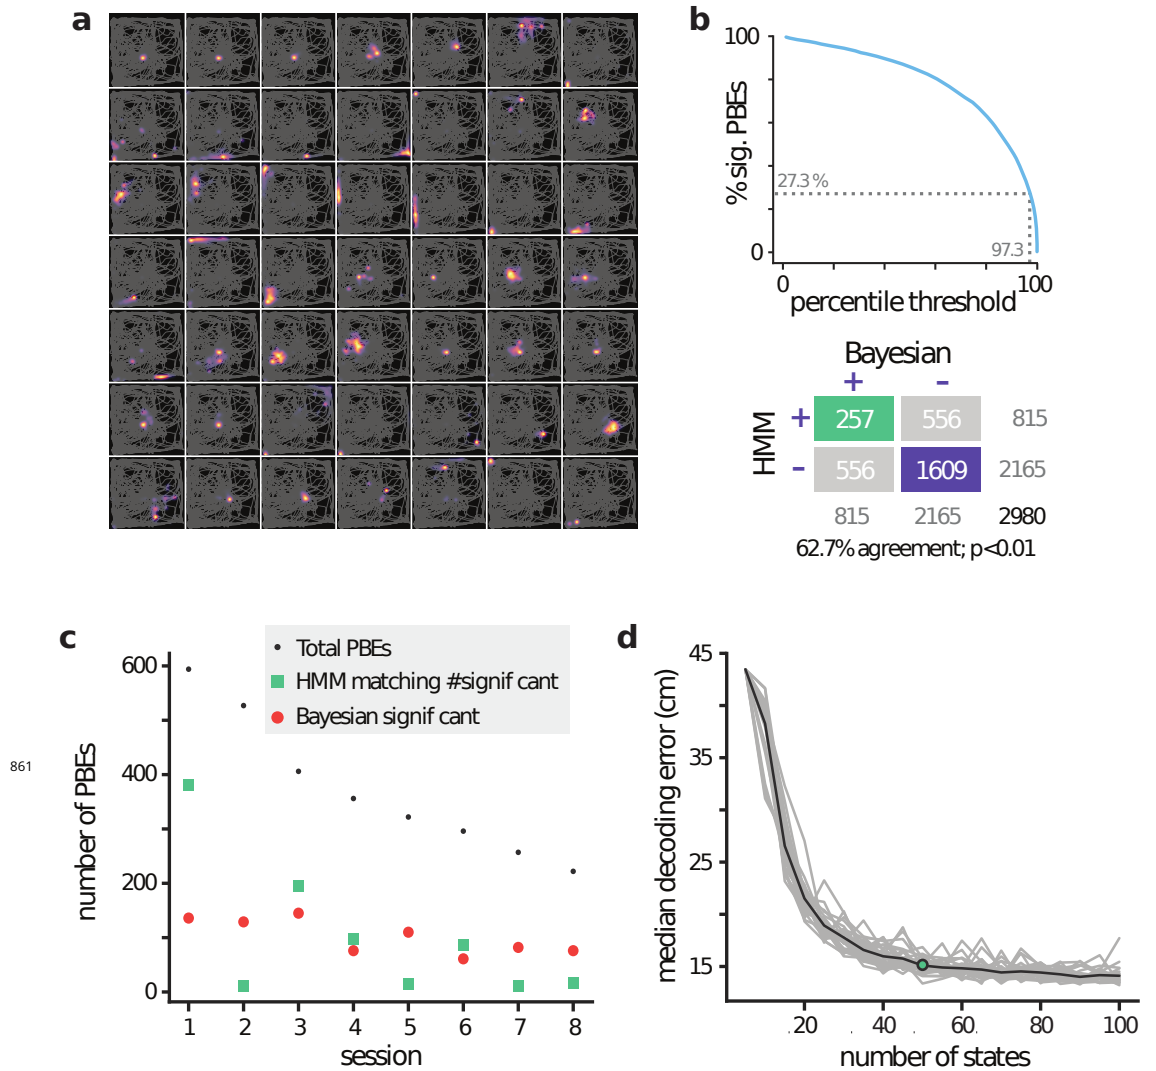


**Figure 6–Figure supplement 1.** Similar to the linear track (one dimensional) case, we find that models learned on actual open field PBE data are significantly more sparse (here showing mean departure sparsity) than their shuffled ( $m = 50$  shuffles) counterparts. This is true for each of the  $n = 8$  open field sessions ( $p < 0.001$ , Mann–Whitney  $U$  test). **a.** Difference [in departure Gini coefficients] between actual and time-swap test data, **b.** between actual and temporal test data, and **c.** between actual and Poisson surrogate data.



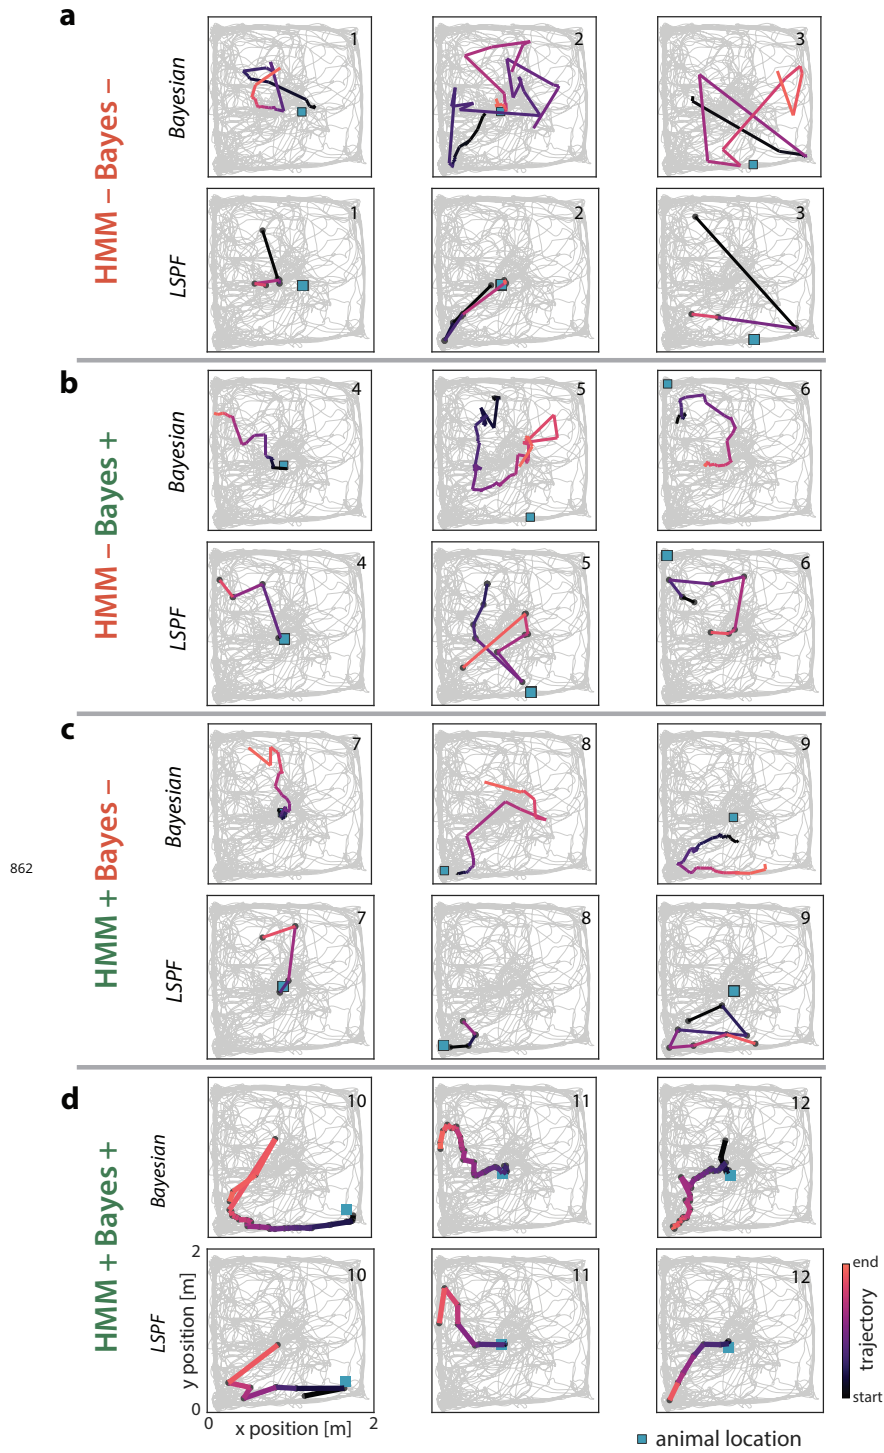
860

**Figure 6-Figure supplement 2. a.** Difference [in observation sparsity Gini coefficients across states] between actual and time-swap test data, **b.** between actual and temporal test data, and **c.** between actual and Poisson surrogate data. Similar to the linear track (one dimensional) case, we find that the observation sparsity across states for actual data are significantly greater than that of both the **b.** temporal and **c.** Poisson surrogates (for each session,  $p < 0.001$ , Mann-Whitney  $U$  test), and that **a.** for some sessions, there are no significant differences between the actual and time-swap surrogates.



**Figure 6-Figure supplement 3. a.** IsPFs for 49 of the 50 latent states from an example session. **b.** (Top) Effect of model-congruence threshold on the number of significant PBEs. (Bottom) Comparison matrix between Bayesian replay detection and our model-congruence approach, where the threshold was chosen to match the total number of significant events pooled over all 8 sessions. **c.** Comparison between number of significant Bayesian events vs number of significant events using our model-congruence approach, when choosing the threshold as in **b.**. Sessions are ordered in decreasing numbers of total PBEs. Note that session 1 is a significant outlier, causing mismatches between many other sessions (2, 5, 7, 8), suggesting that matching on a per-session basis may be more appropriate in this case. **d.** Median position decoding error (via the latent space and IsPFs) was evaluated using cross-validation in an example session ( $n = 30$  realizations for each model considered, shown in gray, mean shown in black), indicating that (i) the PBE-learned latent space encodes underlying spatial information, and (ii) that our PBE models are informative about the underlying position over a wide range of numbers of states.





**Figure 6-Figure supplement 4. a.** Three example PBEs are shown that were classified as non-significant by both the Bayesian and model-congruence approaches. The top row shows the PBEs decoded with place fields using a Bayesian decoder in 20 ms bins, with a 5 ms stride. The bottom row shows the same events, but decoded in 20 ms non-overlapping time bins using the LSPFs. **b.** Three example PBEs are shown that were classified as significant replay by the Bayesian approach, but not by the model-congruence approach. **c.** Three example PBEs are shown that were classified as significant replay by the model-congruence approach, but not by the Bayesian approach. **d.** Three example PBEs are shown that were classified as significant by both approaches.