

Extended Poisson Gaussian-Process Latent Variable Model for Unsupervised Neural Decoding

Della Daiyi Luo¹, Bapun Giri², Kamran Diba², Caleb Kemere¹

¹Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA.

²Department of Anesthesiology, Neuroscience Graduate Program, University of Michigan Medical School, Ann Arbor, MI 48109, USA.

Keywords: Dimension reduction, neural decoding, sharp wave ripples

Abstract

Dimension reduction on neural activity paves a way for unsupervised neural decoding by dissociating the measurement of internal neural state repetition from the measurement of external variable tuning. With assumptions only on the smoothness of latent dynamics and of internal tuning curves, the Poisson Gaussian-process latent variable model (P-GPLVM) (Wu et al., 2017) is a powerful tool to discover the low-dimensional

latent structure for high-dimensional spike trains. However, when given novel neural data, the original model lacks a method to infer their latent trajectories in the learned latent space, limiting its ability for estimating the internal state repetition. Here, we extend the P-GPLVM to enable the latent variable inference of new data constrained by previously learned smoothness and mapping information. We also describe a principled approach for the constrained latent variable inference for temporally-compressed patterns of activity, such as those found in population burst events (PBEs) during hippocampal sharp-wave ripples, as well as metrics for assessing whether the inferred new latent variables are congruent with a previously learned manifold in the latent space. Applying these approaches to hippocampal ensemble recordings during active maze exploration, we replicate the result that P-GPLVM learns a latent space encoding the animal's position. We further demonstrate that this latent space can differentiate one maze context from another. By inferring the latent variables of new neural data during running, certain internal neural states are observed to repeat, which is in accordance with the similarity of experiences encoded by its nearby neural trajectories in the training data manifold. Finally, repetition of internal neural states can be estimated for neural activity during PBEs as well, allowing the identification for replay events of versatile behaviors and more general experiences. Thus, our extension of the P-GPLVM framework for unsupervised analysis of neural activity can be used to answer critical questions related to scientific discovery.

1 Introduction

Memory critically requires firing of neurons in the hippocampus both during ongoing experiences and afterwards, as the resultant memories are consolidated. While rodent studies have focused on spatial memories, hippocampal neurons can be generally understood to represent the conjunction of the sensory features associated with a particular context (Moser et al., 2015), and the temporal sequences that connect local contexts across time during an experience (Eichenbaum and Cohen, 2014; Eichenbaum, 2017). Importantly, sequential firing patterns of neural ensembles reactivate in a time-compressed manner during some of the population burst events (PBEs) that occur during sharp-wave ripple oscillations in sleep or quiet wakefulness (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999). By decoding those events, it has been shown the replay trajectories show a continuum of conformity to the original experience, including variability in momentum and both forward and reverse re-expression (Lee and Wilson, 2002; Foster and Wilson, 2006; Diba and Buzsáki, 2007; Csicsvari et al., 2007; Davidson et al., 2009; Krause and Drugowitsch, 2022). Traditionally, individual replay events have been identified based on a strong assumption of ordered consistency with patterns expressed during exploration. Consequently, un-ordered replay of contexts, or ordered replay of more complicated routes are often excluded from subsequent analysis. Thus, while much has been learned about memory consolidation and recall from the study of replay, existing approaches have colored our understanding. Therefore, a technique for decoding neural activity without strongly stereotyping the patterns represented or requiring a specifically spatial encoding model would be a powerful tool for understanding memory.

To extract the information from spike trains with minimal prior assumptions, one practicable approach is to find a low-dimensional embedding that can reveal the underlying dynamics. The Poisson Gaussian-process latent variable model (P-GPLVM) proposed by Wu et al. (2017) is a probabilistic, nonlinear, and dynamic dimension reduction approach. It infers temporally smooth low-dimensional latent neural trajectories and smooth, non-parametric internal tuning curves from spike trains without referring to external variables. This model consists of Poisson spiking observations and two Gaussian processes, one governing the temporal evolution of latent variables and another governing the nonlinear mapping from high-dimensional neural data to the low-dimensional latent variables.

In the learned low-dimensional latent space, (1) by mapping any possible external variable unto the embedding, how external variables are represented in this latent space can be revealed; (2) by measuring the repetition of internal neural state (relative locations to the low-dimensional training data embedding), the repetition of neural activity encoding external experiences can be detected (Yu et al., 2009; Rubin et al., 2019; Nieh et al., 2021). Unsupervised neural decoding can be achieved by dissociating the measurement of internal neural state repetition from the measurement of external variable tuning. However, an approach for inferring latent variables of new data points in the learned latent space is lacking from the original P-GPLVM model, limiting its utility for decoding.

In this paper, we extend the P-GPLVM framework to enable the latent variable inference of new inputs constrained by the learned smoothness parameters and tuning curves (Figure 1), and develop subsequent analyses to evaluate the new input data in

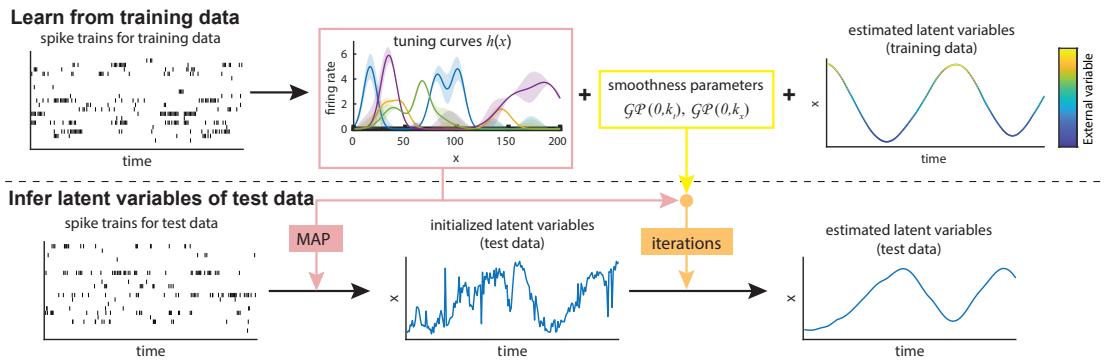


Figure 1: Schematic diagram of the extended P-GPLVM model. Internal tuning curves and smoothness information learned from training data are then used to constrain the inference of test data latent variables in the same latent space.

the latent space. We also present a preprocessing pipeline for PBE decoding when using models trained from behavioral data. The original P-GPLVM can be used to reveal encoded information in neural activity and discover neural trajectory evolution but only within the given training data. This extended model and PBE preprocessing pipeline, especially, makes use of information learned in the model and enables the effective and unsupervised decoding of new neural activity both during behavior and during PBEs.

2 Poisson Gaussian-process latent variable model

2.1 Model structure

Binning spike counts into M time bins from N neurons creates the matrix of spike trains $\mathbf{Y} = \{\mathbf{y}_m\}_{m=1}^M, \mathbf{y}_m \in \mathbb{R}^{N \times 1}$. The m^{th} time bin was recorded at time $t_m, m \in (1, \dots, M)$. In this model, two latent variable matrices will be learned: the log firing rate for Poisson spiking, $\mathbf{F} = \{\mathbf{f}_m\}_{m=1}^M, \mathbf{f}_m \in \mathbb{R}^{N \times 1}$, and the P -dimensional latent variables, $\mathbf{X} = \{\mathbf{x}_m\}_{m=1}^M, \mathbf{x}_m \in \mathbb{R}^{P \times 1}$.

Latent dynamics Each feature of the latent variable matrix \mathbf{X} evolves according to a Gaussian process depending on time t , $x_p(t) \sim \mathcal{GP}(0, k_t)$, $p \in (1, \dots, P)$, where $k_t(t, t') = r \exp(-|t - t'|/l)$, governing the temporal smoothness of x_p . Writing as a multivariate normal distribution,

$$\mathbf{x}_{p,1:M} | t_{1:M} \sim \mathcal{N}(\mathbf{0}, K_t). \quad (1)$$

The covariance K_t is an $M \times M$ matrix with entries k_t at all pairs of time bins.

Nonlinear mapping Let $h : \mathbb{R}^P \rightarrow \mathbb{R}$ be a nonlinear mapping function, describing the firing rate of a neuron in the m^{th} time bin as $\lambda_m = h(\mathbf{x}_m)$. The log tuning curve of the n^{th} cell in response to the latent variable \mathbf{x} is modeled as another Gaussian process as $f_n(\mathbf{x}) = \log h_n(\mathbf{x}) \sim \mathcal{GP}(0, k_x)$, $n \in (1, \dots, N)$. This process has $k_x(\mathbf{x}, \mathbf{x}') = \rho \exp(-\|\mathbf{x} - \mathbf{x}'\|_2^2/2\delta^2)$ as a spatial covariance function. Therefore, the log firing rates of the n^{th} neuron in all time bins have the multivariate normal distribution

$$\mathbf{f}_{n,1:M} | \mathbf{x}_{1:M} \sim \mathcal{N}(\mathbf{0}, K_x) \quad (2)$$

with the $M \times M$ covariance matrix K_x , whose entries are k_x evaluated for all pairs of \mathbf{x} . By combining \mathbf{f} over all the neurons, $\mathbf{F} \in \mathbb{R}^{N \times M}$ as firing rates in units of spike counts per time bin are obtained. *Note that this is per bin, rather than per second. This has consequences for PBE decoding, as described below.*

Poisson spiking Finally, for the n^{th} neuron in the m^{th} time bin, observed spike counts $y_{n,m}$ are assumed to be drawn from a Poisson process given the latent firing rate $\lambda_{n,m} = \exp(f_n(\mathbf{x}_m))$,

$$y_{n,m} | f_n, \mathbf{x}_m \sim \text{Poisson}(\exp(f_n(\mathbf{x}_m))). \quad (3)$$

2.2 Model training

During training, the model iteratively infers latent variables of training data without mapping constraints and optimizes smoothness parameters. In each iteration, \mathbf{X} is first fixed and \mathbf{F} is optimized. For each neuron, the posterior over \mathbf{f}_n is given by

$$p(\mathbf{f}_n | \mathbf{y}_n, \mathbf{X}) \propto p(\mathbf{y}_n | \mathbf{f}_n) p(\mathbf{f}_n | \mathbf{X}). \quad (4)$$

The optimal $\hat{\mathbf{f}}_n$ is found by maximizing the log conditional distribution

$$\log p(\mathbf{f}_n | \mathbf{y}_n, \mathbf{X}) = \log p(\mathbf{y}_n | \mathbf{f}_n) - \frac{1}{2} \mathbf{f}_n^\top K_x^{-1} \mathbf{f}_n - \frac{1}{2} \log |K_x| + \text{constant}. \quad (5)$$

Then, fixing \mathbf{F} , the optimal \mathbf{X} is discovered by maximizing the conditional likelihood

$$\mathbf{X}_{\text{MAP}} = \operatorname{argmax}_{\mathbf{X}} \sum_{n=1}^N p(\mathbf{y}_n | \mathbf{X}) p(\mathbf{X}), \quad (6)$$

where $p(\mathbf{y}_n | \mathbf{X})$ is given by

$$p(\mathbf{y}_n | \mathbf{X}) \propto \int p(\mathbf{y}_n | \mathbf{f}_n) p(\mathbf{f}_n | \mathbf{X}) d\mathbf{f}_n. \quad (7)$$

Using Laplace's method, the approximated log likelihood conditioned on \mathbf{X} is

$$\log q(\mathbf{y}_n | \mathbf{X}) = \log p(\mathbf{y}_n | \hat{\mathbf{f}}_n) - \frac{1}{2} \hat{\mathbf{f}}_n^\top K_x^{-1} \hat{\mathbf{f}}_n - \frac{1}{2} \log |I_M + K_x W_n|. \quad (8)$$

where $W_n = -\nabla \nabla \log p(\mathbf{y}_n | \mathbf{f}_n)$. Hyperparameters $\boldsymbol{\theta} = \{\rho, \delta, r, l\}$ are found by maximizing the same likelihood function.

3 Extensions of P-GPLVM

3.1 Constrained latent variable inference of new data

Latent variable initialization During model training, P-GPLVM learns smoothness parameters, $\boldsymbol{\theta}$, and internal mapping function, $\log h_n : \mathbb{R}^P \rightarrow \mathbb{R}$, parameterized by

latent variables $\hat{\mathbf{f}}_{\text{train}} = \{\hat{f}_m\}_{m=1}^M$ for each cell and $\hat{\mathbf{X}}_{\text{train}} = \{\hat{\mathbf{x}}_m\}_{m=1}^M$ from training data. In the original P-GPLVM paper, the tuning curve vectors \mathbf{f}_{grid} are evaluated at the grid of latent variables $\mathbf{X}_{\text{grid}} = \{\mathbf{x}_g\}_{g=1}^G$ using a joint Gaussian distribution with $\hat{\mathbf{X}}_{\text{train}}$ and $\hat{\mathbf{f}}_{\text{train}}$.

$$\begin{bmatrix} \hat{\mathbf{f}}_{\text{train}} \\ \mathbf{f}_{\text{grid}} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \begin{bmatrix} K_{\text{train}} & \mathbf{k} \\ \mathbf{k}^\top & K_{\text{grid}} \end{bmatrix}), \quad (9)$$

where $K_{\text{train}} = k_x(\hat{\mathbf{X}}_{\text{train}}, \hat{\mathbf{X}}_{\text{train}})$, $K_{\text{grid}} = k_x(\mathbf{X}_{\text{grid}}, \mathbf{X}_{\text{grid}})$, and $\mathbf{k} = k_x(\hat{\mathbf{X}}_{\text{train}}, \mathbf{X}_{\text{grid}})$, are the covariance matrices. The entry $k_x(\mathbf{X}_1, \mathbf{X}_2)_{m,g} = k_x(\mathbf{x}_{1m}, \mathbf{x}_{2g})$. The posterior distribution of \mathbf{f}_{grid} can then be written as

$$\mathbf{f}_{\text{grid}} | \mathbf{X}_{\text{grid}}, \hat{\mathbf{f}}_{\text{train}}, \hat{\mathbf{X}}_{\text{train}} \sim \mathcal{N}(\mathbf{k}^\top K_{\text{train}}^{-1} \hat{\mathbf{f}}_{\text{train}}, \text{diag}(K_{\text{grid}}) - \mathbf{k}^\top K_{\text{train}}^{-1} \mathbf{k}). \quad (10)$$

Given new input data from the same group of neurons, $\mathbf{Y}_{\text{new}} = \{\mathbf{y}_t\}_{t=1}^T, \mathbf{y}_t \in \mathbb{R}^{N \times 1}$, the latent variables $\mathbf{X}_{\text{new}} = \{\mathbf{x}_t\}_{t=1}^T, \mathbf{x}_t \in \mathbb{R}^{P \times 1}$, are initialized with elements in \mathbf{X}_{grid} that maximize the posterior probability using a Bayesian approach. At the t^{th} time bin, the posterior over each element in \mathbf{X}_{grid} is

$$p(\mathbf{x}_g | \mathbf{y}_t, \hat{\mathbf{F}}_{\text{train}}, \hat{\mathbf{X}}_{\text{train}}) \propto p(\mathbf{y}_t | \mathbf{x}_g, \hat{\mathbf{F}}_{\text{train}}, \hat{\mathbf{X}}_{\text{train}}) p(\mathbf{x}_g), \quad (11)$$

where we've assumed a uniform prior over \mathbf{x}_g during initialization. The prior over \mathbf{y}_t is

$$p(\mathbf{y}_t | \mathbf{x}_g, \hat{\mathbf{F}}_{\text{train}}, \hat{\mathbf{X}}_{\text{train}}) = \prod_{n=1}^N p(y_{n,t} | f_{n,g}) p(f_{n,g} | \mathbf{x}_g, \hat{\mathbf{f}}_{\text{train},n}, \hat{\mathbf{X}}_{\text{train}}), \quad (12)$$

The latent variable \mathbf{x}_t corresponding to \mathbf{y}_t is initialized with the \mathbf{x}_g with maximum posterior.

Impose smoothness constraints Merely using MAP tuning curves to constrain the latent variable inference of new input data omits the learned temporal smoothness in the

latent space. To impose both temporal smoothness and tuning constraints, subsequent iterations are needed. We noted that employing \mathbf{f}_{grid} and \mathbf{X}_{grid} as tuning curve vectors is computationally expensive – considering a samples per dimension, a P -dimensional \mathbf{X}_{grid} contains a^P elements. Moreover, a majority of the grid elements contain minimal information as no $\hat{\mathbf{x}}_{\text{train}}$ locate nearby. To save computational costs during iterations and to preserve the tuning details of $\hat{\mathbf{f}}_{\text{train}}$ in response to $\hat{\mathbf{X}}_{\text{train}}$ as much as possible, $\hat{\mathbf{F}}_{\text{train}}$ and $\hat{\mathbf{X}}_{\text{train}}$ were directly used as the tuning curve vectors (TC), $\hat{\mathbf{F}}_{\text{TC}}$ and their encoded latent variables $\hat{\mathbf{X}}_{\text{TC}}$.

Iterations are similar to those during unconstrained inference as described in Section 2.2, but the smoothness parameters are fixed as learned and the mapping function $\log h_n$ is constrained by $\hat{\mathbf{f}}_{\text{TC}}$ and $\hat{\mathbf{X}}_{\text{TC}}$. This is achieved by substituting $p(\mathbf{f}_n|\mathbf{X}, \hat{\mathbf{f}}_{\text{TC}}, \hat{\mathbf{X}}_{\text{TC}})$ for $p(\mathbf{f}_n|\mathbf{X})$ in Eq. 4 and 7. As in Eq. 9, the joint distribution with $\hat{\mathbf{f}}_{\text{TC}}$ and $\hat{\mathbf{X}}_{\text{TC}}$ gives posterior of \mathbf{f}_n as

$$\mathbf{f}_n|\mathbf{X}, \hat{\mathbf{f}}_{\text{TC}n}, \hat{\mathbf{X}}_{\text{TC}} \sim \mathcal{N}(\boldsymbol{\mu}_n, \sigma^2 I_T) \quad (13)$$

where $\boldsymbol{\mu}_n = \mathbf{k}^\top K_{\text{TC}}^{-1} \hat{\mathbf{f}}_{\text{TC}n}$, $\mathbf{k} = k_x(\hat{\mathbf{X}}_{\text{TC}}, \mathbf{X})$, $K_{\text{TC}} = k_x(\hat{\mathbf{X}}_{\text{TC}}, \hat{\mathbf{X}}_{\text{TC}})$ and σ^2 is the observation noise. Correspondingly, Eq. 5 is modified as

$$\log p(\mathbf{f}_n|\mathbf{y}_n, \mathbf{X}, \hat{\mathbf{f}}_{\text{TC}n}, \hat{\mathbf{X}}_{\text{TC}}) = \log p(\mathbf{y}_n|\mathbf{f}_n) - \frac{1}{2\sigma^2}(\mathbf{f}_n - \boldsymbol{\mu}_n)^\top(\mathbf{f}_n - \boldsymbol{\mu}_n) - \frac{T}{2} \log \sigma^2. \quad (14)$$

Finally, obtaining the optimal $\hat{\mathbf{f}}_n$, Eq. 8 is modified as

$$\log q(\mathbf{y}_n|\mathbf{X}, \hat{\mathbf{f}}_{\text{TC}n}, \hat{\mathbf{X}}_{\text{TC}}) = \log p(\mathbf{y}_n|\hat{\mathbf{f}}_n) - \frac{1}{2\sigma^2}(\hat{\mathbf{f}}_n - \boldsymbol{\mu}_n)^\top(\hat{\mathbf{f}}_n - \boldsymbol{\mu}_n) - \frac{1}{2} \log |I_T + \sigma^2 W_n|. \quad (15)$$

3.2 Preprocessing for PBE decoding

Critically for our model, neural activity during PBEs is temporally-compressed relative to that expressed during exploration. Consequently, the firing rate model learned on exploration-related neural activity will not properly model PBEs. This is addressed in two steps, choosing a shorter time bin (implicitly compressing time) and explicitly scaling the model parameters. Temporal compression implies that the time intervals between pairs of place cells activating during exploration running are expected to be proportional to those during replay. Thus, the cross-correlation histogram (Harrison et al., 2013; Karlsson and Frank, 2009) of spike trains from place cell pairs is used to find the PBE time bin size that would best match the binned behavioral data. Place cells are identified in either running direction and then are pooled. For each place cell pair with spike trains s_1, \dots, s_M from one cell and t_1, \dots, t_N from another, time lags between any pairs of spikes (s_m, t_n) within 4 seconds are recorded. The histogram of all these spike time lags is the cross-correlation histogram (CCH) for this cell pair. All histograms are centered at 0s time lag and share the same number of bins. For ease, the histogram bin size of CCH during active exploration is set as the same as that used for P-GPLVM analysis of this data. Among all pairwise combinations of place cells, cell pairs with place field (where peak firing rate occurs) distances larger than the 80th percentile of the distribution of travel distance in 4 seconds are excluded from subsequent analysis. For each included cell pair, the Pearson correlation between CCH during exploration and CCH during PBEs is evaluated. The optimal bin size for PBEs is found by maximizing the sum of all significant positive Pearson correlation coefficients using a range of possible PBE bin sizes.

As the unit of our tuning curves is spike counts per time bin, for PBE-decoding, the firing rates of P-GPLVM tuning curves learned using neural activity during active exploration must be further scaled. Scaling by the ratio of the time bin sizes is found to yield good PBE-decoding performance.

3.3 Analysis in latent space

3.1 Number of well-separated manifolds

The hippocampus "remaps" between different environments, meaning that whether a given neuron is active, and if so, how its spatial tuning will relate to that of other neurons is essentially random between different environments (Moser et al., 2015; Alme et al., 2014) (though recent studies have challenged this concept (Cai et al., 2016)). This implies that different environments should correspond to well-separated manifolds in latent space. This idea intuitively leads to the investigation that if the low-dimensional latent variable $\hat{\mathbf{X}}_{\text{train}}$ is organized in a single manifold or in multiple well-separated manifolds. This procedure consists of constructing a K -nearest neighbor graph in $\hat{\mathbf{X}}_{\text{train}}$ and then counting the number of connected components in the graph. Components containing less than 3% of total number of points are considered "residual components". Starting with a small number, K is gradually increased until there are no residual components left. The number of well-separated manifolds will stabilize as K continues to increase for a multiple times. This stabilized number is defined as the number of connected components in $\hat{\mathbf{X}}_{\text{train}}$. Each separated manifold is denoted as $\hat{\mathbf{X}}(i)$ hereafter.

3.2 Congruence with the learned manifold(s)

In the latent space, similar trajectories indicate repetition of neural state encoding similar external experiences. To assess new test data, whether or not their latent variables have similar spatial distribution and dynamics to those of the training data is investigated. Training and test data can contain multiple continuous recording segments. Each continuous segment is a single neural trajectory in the latent space. When the neural patterns and dynamics of new inputs \mathbf{Y}_{new} match the learned model perfectly, their estimated latent trajectories \mathbf{X}_{new} will (i) span the manifold of $\hat{\mathbf{X}}_{\text{TC}}$ and (ii) progress smoothly along it. Three measures are proposed to evaluate congruence of new input data with a learned model.

Log likelihood We use log likelihood to measure how well the learned P-GPLVM model fits the test data. The joint probability of \mathbf{Y} , \mathbf{F} , \mathbf{X} for test data is computed as

$$p(\mathbf{Y}, \mathbf{F}, \mathbf{X} | \hat{\mathbf{F}}_{\text{TC}}, \hat{\mathbf{X}}_{\text{TC}}, \boldsymbol{\theta}) = p(\mathbf{Y} | \mathbf{F}) p(\mathbf{F} | \mathbf{X}, \hat{\mathbf{F}}_{\text{TC}}, \hat{\mathbf{X}}_{\text{TC}}, \rho, \delta) p(\mathbf{X} | r, l). \quad (16)$$

The log likelihood (LLH) of a continuous segment in test data is written as

$$\begin{aligned} LLH &= \log p(\mathbf{Y}, \mathbf{F}, \mathbf{X} | \hat{\mathbf{F}}_{\text{TC}}, \hat{\mathbf{X}}_{\text{TC}}, \boldsymbol{\theta}) \\ &= \sum_{n=1}^N \sum_{t=1}^T \log p(y_{n,t} | f_{n,t}) + \sum_{n=1}^N \log p(\mathbf{f}_n | \mathbf{X}, \hat{\mathbf{f}}_{\text{TC}_n}, \hat{\mathbf{X}}_{\text{TC}}, \rho, \delta) + \sum_{p=1}^P \log p(\mathbf{x}_p | r, l). \end{aligned} \quad (17)$$

Recall that

$$y_{n,t} | f_{n,t} \sim \text{Poisson}(\exp(f_{n,t})), \quad (18)$$

$$\mathbf{f}_n | \mathbf{X}, \hat{\mathbf{f}}_{\text{TC}_n}, \hat{\mathbf{X}}_{\text{TC}}, \rho, \delta \sim \mathcal{N}(k_x(\hat{\mathbf{X}}_{\text{TC}}, \mathbf{X})^\top k_x(\hat{\mathbf{X}}_{\text{TC}}, \hat{\mathbf{X}}_{\text{TC}})^{-1} \hat{\mathbf{f}}_{\text{TC}_n}, \sigma^2 I_T), \quad (19)$$

$$\mathbf{x}_p | r, l \sim \mathcal{N}(\mathbf{0}, K_t). \quad (20)$$

Spatial consistency This metric measures how well a latent trajectory, $\mathbf{X}_{\text{new}} = \{\mathbf{x}_t\}_{t=1}^T$,

spans the latent variable manifold(s) of the tuning curve vector, $\hat{\mathbf{X}}_{\text{TC}} = \{\hat{\mathbf{x}}_m\}_{m=1}^M$.

Firstly, for each point \mathbf{x}_t in the trajectory, its K -nearest neighbors are searched in the manifold $\hat{\mathbf{X}}_{\text{TC}}$, obtaining the set of its nearest neighbors $S_t = \{\hat{\mathbf{x}}_{m_k}\}_{k=1}^K$ and its distances to those neighbors $\{\|\mathbf{x}_t - \hat{\mathbf{x}}_{m_k}\|_2\}_{k=1}^K$. When the trajectory spans the manifold nicely, the union of neighbor sets, $\bigcup_{t=1}^T S_t$, should have much more elements than by chance. The number of all identified neighbors is then weighted by the neighbor distances to integrate the neighborhood quality into the measure.

For each $\hat{\mathbf{x}}_m \in \bigcup_{t=1}^T S_t$, its distance to its closest point in \mathbf{X}_{new} , $d_m = \min_{t,k} \|\mathbf{x}_t - \hat{\mathbf{x}}_{m_k}\|_2$ is estimated. Then, the spatial consistency of this trajectory in the latent space is computed as

$$\text{consistency} = \sum_{m=1}^M w_m \quad (21)$$

where

$$w_m = \begin{cases} \frac{c}{c+d_m}, & \text{if } \hat{\mathbf{x}}_m \in \bigcup_{t=1}^T S_t \\ 0, & \text{otherwise} \end{cases} \quad (22)$$

which means, when one $\hat{\mathbf{x}}_m$ is identified as a neighbor, it contributes to the spatial consistency. This contribution is weighted by the distance to its nearest neighbor on the trajectory, which is 1 when the distance is 0, and then decay as the distance increases, reaching at 0.5 when the distance is c . The parameter c is set as 20% of the standard deviation of $\hat{\mathbf{X}}_{\text{TC}}$. When $\hat{\mathbf{X}}_{\text{TC}}$ contains more than one separated manifold, the standard deviation is computed as the average within-manifold standard deviation:

$$c = 20\% \sqrt{\frac{1}{M} \sum_i \sum_{\hat{\mathbf{x}}_m \in \hat{\mathbf{X}}(i)} \|\hat{\mathbf{x}}_m - \boldsymbol{\mu}_i\|_2^2} \quad (23)$$

where $\boldsymbol{\mu}_i$ is the center of manifold $\hat{\mathbf{X}}(i)$.

Higher value of spatial consistency indicates that this trajectory travels a long way along the tuning curve manifold(s) rather than only jiggling at a small local area or locating far away. The spatial consistency to $\hat{\mathbf{X}}_{\text{TC}}$ is first used to evaluate the neural trajectory behavior in the latent space. Further, when $\hat{\mathbf{X}}_{\text{TC}}$ contains more than one well-separated manifold $\hat{\mathbf{X}}(i)$, we also estimate the manifold contribution ratio to the consistency value, which is computed as in Eq. 21 but only for $\hat{\mathbf{x}}_m \in \hat{\mathbf{X}}(i)$. The ratio between these contribution values from manifolds indicate to which manifold the trajectory belongs.

Average step distance This metric reveals the consistency between temporal adjacency and pattern-matched location adjacency. The average Euclidean distance of each step in each latent trajectory of \mathbf{X}_{new} is estimated. In general, the tuning curve constraint tends to drive \mathbf{X}_{new} towards their pattern-matched locations on the $\hat{\mathbf{X}}_{\text{TC}}$ manifold. When the new data have different temporal sequences from the training data, pattern-matched locations of temporally adjacent points are mostly distant, leading to longer step distances.

4 Experiments and results

The dataset used in this paper is a neural recording collected in the dorsal hippocampal CA1 and CA3 areas from a rat. The recording started with a 3 hours rest session (pre),

and then the animal was exposed to a novel maze (maze1) for 50 minutes, running back and forth to get rewards at the two ends. Following another two hours rest (post1), the animal explored another novel maze (maze2) for 50 minutes. Lastly, the animal went through a 5 hours rest session (post2). Spike trains during the two maze exploration sessions were binned into 500 ms time bins. Animal positions were linearized and then running periods (speed>3cm/s and peak speed>5cm/s) were extracted along with their corresponding time bin indices.

4.1 Exploration in one maze

We first examined this P-GPLVM with the neural activity when the animal was exposed to only one maze (maze1). The linearized animal positions during running in maze1 exploration session are shown in Figure 2A and corresponding original positions in Figure 2B. One chunk of spike trains was selected as training data to train the model (orange lines in Figure 2A, colored scatters in Figure 2B) and then the spike trains of a continuous running segment in maze1 was used as test data to evaluate the decoding performance (black line in Figure 2A,B).

After fitting the model, smoothness parameters θ and latent variables of training data, $\hat{\mathbf{X}}_{\text{train}}$ and $\hat{\mathbf{F}}_{\text{train}}$, were obtained. Only one well-separated manifold is found in $\hat{\mathbf{X}}_{\text{train}}$, suggesting that there is no multiple distinctly different neural processes found in this training data. Given the known spatial tuning in hippocampus, to examine the information encoded in the latent space, $\hat{\mathbf{X}}_{\text{train}}$ are color-coded with the corresponding animal linearized positions. Position information appears to be encoded along the manifold smoothly (Figure 2C). The latent neural trajectories seem to be a bit unkempt,

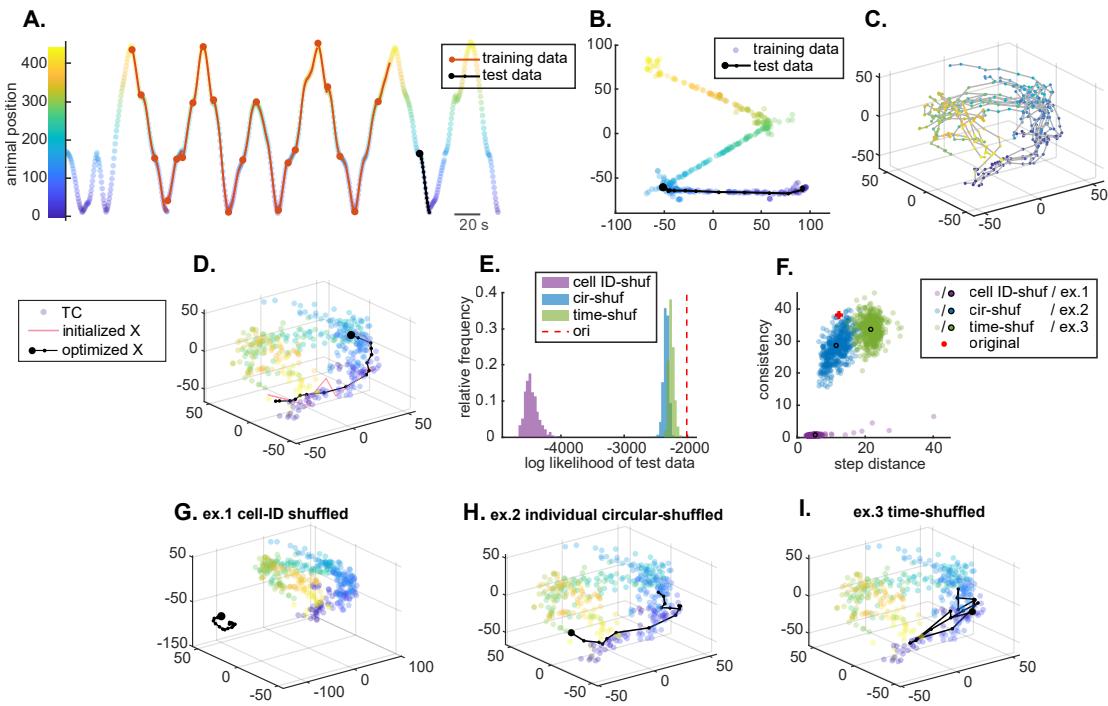


Figure 2: **(A)** Linearized animal position during running periods in maze1 exploration session. The orange / black lines indicate the segments used as training / test data. The starting of each continuous segment is marked by a big filled circle. Each step in the test data is marked by small dots. Scatters are color-coded by linearized animal positions, same in subsequent panels (B-D, G-I). **(B)** Colored scatters of original animal positions in training data. The black line indicates the same test data segment in (A). **(C)** Latent variables of training data in the learned P-GPLVM. Continuous segments are connected by lines as latent neural trajectories. **(D)** Inferred neural trajectory of test data in the learned latent space. The semi-transparent markers are identical with the dot scatters in (C), which, in this panel, indicates the latent variable manifold of tuning curve vectors, \hat{X}_{TC} . The black line depicts the inferred neural trajectory of test data, progressing along the manifold \hat{X}_{TC} in accordance with the real animal positions (from blue to purple). **(E)** Log likelihood of the original test data (dashed red line) and its shuffled versions (histograms). **(F)** Estimated step distances vs. spatial consistency to \hat{X}_{TC} of the inferred neural trajectory in (D) compared with its cell ID-shuffled, individual circular-shuffled, and time-shuffled versions. The three black-edged circles indicate one example of cell ID-shuffled test data (ex.1), of individual circular-shuffled test data (ex.2), and of time-shuffled test data (ex.3), respectively. **(G, H, I)** Same presentation as in (D), but the inferred latent trajectories of ex.1-3 in (F).

which is still reasonable because this was the animal's first exposure to maze1.

Next, we estimated the internal state repetition of the test data in the learned latent space. In Figure 2D, the manifold of tuning curve vector latent variables $\hat{\mathbf{X}}_{TC}$, which is $\hat{\mathbf{X}}_{train}$ in this setting, is shown as the colored semi-transparent markers. Using the Bayesian approach, the initialized latent variables of test data appears as a bumpy trajectory (light red line). Applying the constrained inference, the latent variables converged to a smooth latent trajectory after 15 iterations (black line). By comparing the color in Figure 2B and Figure 2D, it is obvious that, the inferred neural trajectory of test data proceeds closely along those of training data with similar experiences, where the animal traveled from positions in blue to positions in purple, in accordance with the repetition of similar internal neural state.

To validate this internal neural state repetition and contrast the inferred latent variables of test data, three kinds of surrogate data were generated:

- Cell identity(ID)-shuffled data, by randomly permuting the cell identity of the test data, which scrambles the original neural coactivation patterns and permutes neuronal firing rates, but preserves the temporal sequences.
- Local individual circular-shuffled (cir-shuf) data, by randomly and circularly permuting the test data of each cell individually within each continuous segment, which scrambles coactivation patterns but preserves the temporal smoothness and local neuronal firing rates.
- Local time-shuffled data, by randomly permuting the temporal order of the test data within each continuous segment, which preserves the coactivation patterns

but scrambles the time sequences.

For each shuffle type, we generated 500 surrogate test data. Metric values and example latent trajectories of the original and surrogate data are shown in Figure 2.

Figure 2E depicts the distribution of LLH values of the inferred latent variables of original test data and of the surrogate test data. Overall, the original test data has a higher LLH value than all the surrogate data. Figure 2F shows their step distance and spatial consistency values in scatters. Inferred latent trajectories of example surrogate data are shown in Figure 2G-I, whose metric values are indicated in the black-edged scatter in Figure 2F. In the latent space, the trajectories of cell ID-shuffled and of time-shuffled test data look distinct from that of the original test data trajectory, while the circular-shuffled test data trajectory has visually similar behavior.

Note that, the iterative inference process searches for an optimal balance between temporal smoothness and tuning curve constraints. Because of the changed relative firing rates among cells in cell ID-shuffled data, the tuning curve constraints push the trajectories far away from $\hat{\mathbf{X}}_{TC}$, where temporal smoothness constraint dominates, leading to small step distances, diminutive spatial consistency values, and the lowest LLHs among all types of shuffled data. With preserved local neuronal firing rates, circular-shuffled data trajectories stay in the similar area as the original data, but due to the scrambled coactivation patterns and preserved temporal smoothness, the tuning curve constraints are compromised by the temporal smoothness constraint, resulting in less likely but smooth trajectories with step distances comparable to the original test data, slightly lower spatial consistency values, and significantly lower LLH values. As for time-shuffled data, with intact coactivation patterns at each time point, the tuning curve

constraints are dominating during the inference, driving the latent variables to the corresponding locations on the manifold of $\hat{\mathbf{X}}_{TC}$ as the original test data does, while the temporal smoothness is sacrificed, resulting in spiky trajectories with large step distances, comparable spatial consistency, and slightly lower LLH values.

We contrast the original test data with its shuffled versions using a modified z-scored LLH, calculated by dividing the difference from the median by the median absolute deviation of shuffled data LLH. The z-scored LLH of the original test data among its cir-shuf versions is 13.56, and among its time-shuffled versions, 8.98, indicating that the test data is a valid sequential repetition of internal neural states captured in the training data. This measure separately verifies the matching of coactivation patterns between test data and tuning curves (when compared with cir-shuf test data), and the matching of temporal sequences (when compared with time-shuffled test data).

4.2 Exploration in two mazes

Next, we asked if our approach can still work well for neural activity during exploration in more than one environment context. One chunk of neural data during running in maze1 (orange lines in Figure 3A, colored scatters on the left in Figure 3B) and another in maze2 (blue lines in Figure 3A, colored scatters on the right in Figure 3B) are selected and concatenated as the training data to train a new P-GPLVM. The test data is from maze1, the same as in the previous section (black line in Figure 3A, B). All the semi-transparent markers in the scatter plot are colored by the corresponding linearized animal positions.

In $\hat{\mathbf{X}}_{train}$ of the newly learned model, two well-separated manifolds, m1 ($\hat{\mathbf{X}}(1)$)

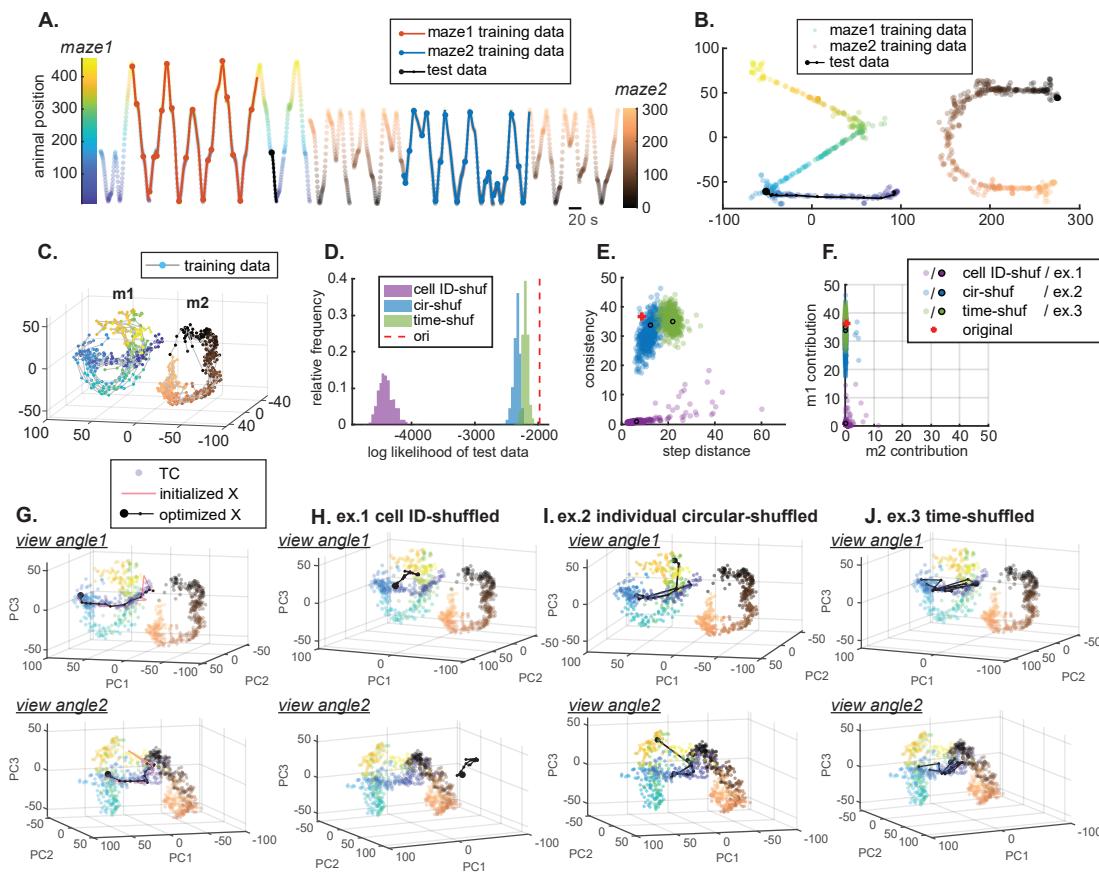


Figure 3: **(A)** Linearized animal position during running periods in maze1 and maze2 exploration sessions. The orange / blue lines indicate the segments in maze1 / maze2 used as training data. The black line indicates the test data. Scatters are color-coded by contexts and linearized animal positions, same in (B-C, G-J). **(B)** Scatters of original animal positions in training data from maze1 (left) and maze2 (right). The black line indicates the same test data segment as in (A). **(C)** Following fitting the training data, latent variables are separated into two distinct manifolds, m1 and m2. Each data point is colored by its corresponding linearized animal position as in (A). **(D)** Log likelihoods (LLH) of test data (red dashed line), cell ID-shuffled (purple histogram), individual circular-shuffled (blue histogram), and time-shuffled surrogate data (green histogram). **(E)** Inferred latent neural trajectory step distance vs. spatial consistency to \hat{X}_{TC} of original test data (red +), of cell ID-shuffled (purple scatters), of individual circular-shuffled (blue scatters), and of time-shuffled (green scatters) surrogate data. Example surrogate data indicated in one black-edged circle from either shuffle type is then shown for visualization. **(F)** Manifold contribution to neural trajectory spatial consistency from m1 vs. from m2, of original and surrogate test data. **(G)** Latent neural trajectories as initialized (light red line) and optimized after iterations (black line) of test data in the latent space. Latent variables of tuning curve vectors, \hat{X}_{TC} , are shown in semi-transparent markers color-coded by the animal positions, same as in (C). Both initialized and optimized trajectories of test data are associated with the corresponding context manifold. Unlike the bumpy initialization, optimized neural trajectory is smooth and progresses according to the animal positions. **(H-J)** Same representations as in (G), inferred trajectory of (H) ex.1 for cell ID-shuffled test data, (I) ex.2 for individual circular-shuffled test data, and (J) ex.3 for time-shuffled test data, as indicated in black-edged circles in (E,F).

and $m2(\hat{X}(2))$, are found, each of whom turns out to contain all data points from one of the maze contexts. Using the same color code as in Figure 3A and B on \hat{X}_{train} , position information appears to be encoded along the corresponding context manifold smoothly (Figure 3C). Thus, this approach is capable of detecting the presence of two environment contexts as well as capturing the encoded position information merely from the neural activity in the training data.

Again, cell ID-shuffled, individual circular-shuffled, and time-shuffled surrogate test data were generated. Along with the original test data, their latent variables in the new latent space are inferred, which behave quite similarly to those in Section 4.1. Figure 3D depicts the distribution of LLH values of the original test data and the surrogate test data. Cell ID-shuffled test data have much lower LLHs than all the other type of data. Circular-shuffled data have slightly smaller LLHs than time-shuffled data. The original test data has a higher LLH value than all surrogate data, where z-scored LLH among its circular-shuffled versions is 11.75, and among its time-shuffled versions, 8.40, confirming that the test data is a valid internal neural state repetition.

Next, we investigated the content of this state repetition by looking into the latent variable locations relative to \hat{X}_{train} . The inferred latent trajectory of the test data and its initialization in this new latent space are shown in Figure 3G. The trajectory is similar to those neural trajectories of training data with similar experiences, traversing the same arm of maze1, from positions blue to purple just as in Figure 3B.

As a high-level summary of the spatial distribution and dynamics of these neural trajectories, Figure 3E shows the step distance and spatial consistency values to \hat{X}_{train} , while Figure 3F shows the manifold contribution to spatial consistency values from $m1$

versus from m2. Examples are shown in Figure 3H-J (metric values indicated in black-edged scatters in Figure 3E,F). The inferred cell ID-shuffled test data trajectories tend to have small steps and locate far away from both manifolds, having diminutive spatial consistency values. Compared with original test data trajectory, circular-shuffled data have smooth trajectories around the same area, with slightly smaller spatial consistency values and comparable step distances. Latent variables of time-shuffled test data has similar location as the original test data does, but are mostly taking big jumps across time. Just like the original test data, both circular-shuffled and time-shuffled test data have similarly high m1 contribution to spatial consistency and diminutive m2 contribution, suggesting the test data is repeating neural state encoding experience in maze1 rather than maze2. Recall that circular-shuffled data have scrambled coactivation patterns, which indicates that local neuronal firing rates can already distinguish animal contexts.

4.3 PBE decoding

In the previous two sections, it has been demonstrated that the model learned from training data during running can capture the neural state repetition in test data during running well. Next, we evaluated whether this running-state model could be used to estimate the internal state repetition in neural data during PBEs, which occurred when the animal had paused running in the exploration session.

First, the best time bin size for PBE data was searched. Both in the maze1 and maze2 exploration sessions, the 80th percentile of 4 seconds' travel distance is 86-87cm. Therefore, only place cell pairs with their place field distances less than 87cm are in-

cluded in the subsequent analysis. It turns out that, for both maze sessions, the sum of significant ($p < 0.01$) positive Pearson correlation coefficient r between CCH during running and CCH during PBEs across cell pairs reached its peak when the PBE time bin size is 17 ms. Therefore, PBEs were binned into 17 ms time bins.

Next, in the latent space learned using running data from the two maze exploration sessions described in Section 4.2, we examined whether the inferred latent trajectories of binned PBE data could be identified as neural state repetitions. Firing rates of the tuning curve vectors were scaled by the time bin size ratio 17/500. Among 562 PBEs detected within the maze1 exploration session, we randomly selected 100 PBEs as test data. To evaluate the resulting latent trajectories, for each PBE, 200 surrogate data for each of the three shuffling types were generated and metric values were assessed.

The first question to ask about the PBE latent trajectories is the validity of neural state repetition. We consider a PBE to be a significant replay event, when its LLH value is larger than all of those of its circular-shuffled versions and its modified z-scored LLH is greater than 3. Moreover, a significant replay event should have a latent trajectory located closely to the tuning curve latent variable $\hat{\mathbf{X}}_{\text{TC}}$ manifolds to guarantee the strength of the tuning curve constraint. Therefore, the average distance of the latent trajectory to its nearest neighbors identified during measuring the spatial consistency value is estimated. PBEs with average nearest neighbor distance larger than the standard deviation of $\hat{\mathbf{X}}_{\text{TC}}$ are excluded. Given that neuron firings of a replay are not exactly the time-compressed version of those during corresponding behavior, this criterion is a strict indicator of the ensemble activity pattern reactivation, and therefore, a strict indicator of the validity of being a replay event. Following the identification as a replay, z-scored

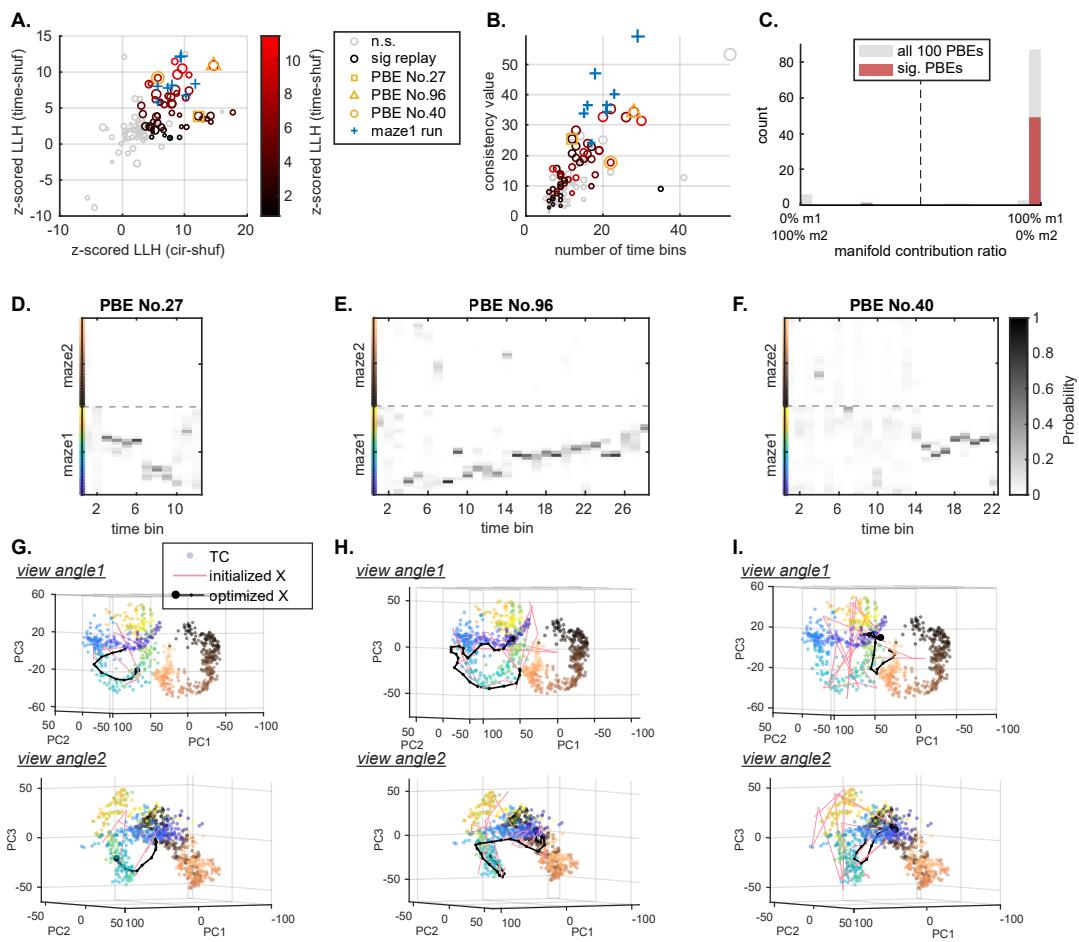


Figure 4: Inferring latent trajectories of PBE data in the P-GPLVM latent space learned from running data. **(A)** Modified z-scored LLH values among circular-shuffled versions versus among time-shuffled versions, of the 100 selected PBEs and of 8 continuous running segments in maze1 for comparison. The size of each circle indicates the spatial consistency value of the PBE. PBEs not meeting the replay criterion are in gray and PBEs identified as significant replay events are colored according to their z-scored LLH among time-shuffled versions. **(B)** Number of time bins of each test data versus their spatial consistency to the latent variable manifolds of tuning curve vector, \hat{X}_{TC} . Dots are presented in the same way as in (A). **(C)** Histogram of ratio of manifold contribution to spatial consistency. Most of the PBEs are associated with maze1. **(D-F)** For comparison, we use traditional Bayesian decoding to render the posterior position distribution of the example PBEs (No.27, No.96, No.40) indicated in markers in (A-B). Corresponding context and position colors are indicated on the left side, same color code as Figure 3A. **(G-I)** Inferred latent trajectories of the example PBEs in (D-F) in the P-GPLVM latent space. While the initialized trajectories are jumpy, after iterations, the trajectories of the PBEs all converge to the maze1 manifold (same representation as in Figure 3G).

LLH among its time-shuffled versions suggests the likeliness of being a *sequential* replay. Figure 4A shows the LLH values z-scored by circular-shuffled versions versus z-scored by time-shuffled versions, of the 100 selected PBEs and of 8 continuous running segments in maze1 for comparison. Each PBE is depicted as a dot, whose size stands for its spatial consistency value. 50 out of the selected 100 PBEs meet the replay criterion and are considered significant replay events, colored according to their z-scored LLH among time-shuffled versions. Insignificant PBEs are colored in gray.

Next, the spatial distribution and dynamics features of each PBE were investigated. Since it had been shown that position information is encoded along the manifolds in latent space, for comparison, we estimated spatial tuning curves in the two running directions separately and used a Bayesian model to decode the PBE positions following the traditional method. Among those probable replay events, three example PBEs of different trajectory behaviors are highlighted by different markers in Figure 4A-B, whose Bayesian decoding results and inferred latent trajectories are visualized subsequently in Figure 4D-I. Figure 4B shows the number of time bins of each PBE and its spatial consistency to \hat{X}_{TC} . Visually, there is a trend that more time bins lead to larger spatial consistency values, which is intuitive because longer trajectory can travel through more area. PBEs in this trend indicate one-directional replays of experiences. For example, PBE No.96 in Figure 4E and H is a continuous replay of experience traveling from location in purple all the way to location in green, consistent with the Bayesian decoded results. PBEs with fewer time bins relative to those with comparable spatial consistency values (top left side of the chart) indicate that their trajectories might be of high speed or even have leaps. PBE No.27 in Figure 4D and G is shown as an example of this kind,

whose Bayesian decoded results suggest a travel from location green to location purple with a leap in the middle. On the contrary, PBEs with more time bins relative to those with comparable spatial consistency values (bottom right side of the chart) indicate that their trajectories might be stationary or slightly oscillating. One example is PBE No.40 in Figure 4F and I, whose Bayesian decoded positions are first unfathomable (which might not be able to be decoded by a spatial decoder) and then travel around location green. Its latent trajectory at the first major part jiggled around location purple and then traveled to location green. Figure 4C shows the manifold contribution to spatial consistency value from m1 manifold versus from m2 manifold of each PBE. As expected, since the animal hadn't experienced maze2 yet, almost all of the probable replay events belong to m1 manifold rather than to m2 manifold, indicating that they are replay events of maze1 experience.

Combining the observation of metric values, we can estimate the validity of neural state repetition, evaluate the pattern of this repetition, and infer the external variables of the replayed experience. *Note that we have not used position data in any of our analyses. This approach is the first presentation of a method for continuously decoding and identifying replay events without relying on external variables.*

Finally, with no need of assuming external variables that drive the neural activity or predetermining replay behavior patterns, the summary of replay behavior through out all sessions in this experiment can be revealed by applying the extended P-GPLVM model to all PBEs detected. Latent neural trajectories of all PBEs are inferred. For each PBE, 100 circular-shuffled versions were generated to for identifying significant replay events. PBE occurrence time vs. the ratio of manifold contribution to PBE spatial

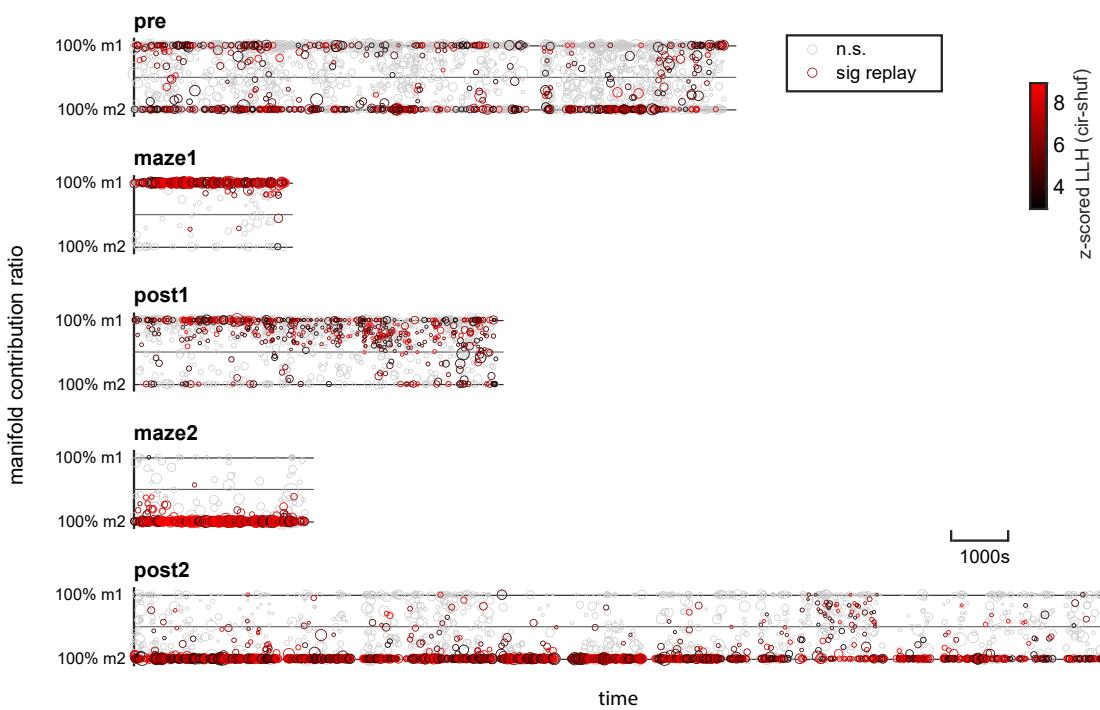


Figure 5: PBE occurrence time vs. ratio of manifold contribution for PBE latent trajectories in all sessions. The top solid line indicates 100% of the trajectory nearest neighbors are from maze1 manifold (m1). The bottom solid line indicated 100% of the nearest neighbors are from maze2 manifold (m2). Each circle indicate one detected PBE. Significant replay events are colored in black to light red according to their z-scored LLH values among their circular-shuffled versions.

consistency value is shown in Figure 5 from pre session to post2 session in time order, where each PBE is depicted as a circle, whose size indicates its spatial consistency value to \hat{X}_{TC} . Circles on the top / bottom line means 100% of identified nearest neighbors for the latent neural trajectories are from m1 / m2 manifold, indicating affiliation with maze1 / maze2 context experience. Circles located in between suggest mixtures. PBEs identified as significant replay events are colored according to its z-scored LLH value among their circular-shuffled versions. In the pre rest session, identified significant replay events are comparably sparse and have small spatial consistency values. In maze1 and maze2 exploration sessions, as expected, almost all the significant replay events are associated with their current contexts and have comparably larger consistency and z-scored LLH values, which indicates long and sequential experience replay. In post1 rest session, the majority of significant replay events are associated with maze1 context and most of them have small consistency values. Those replay events became less associated with maze1 as the rest time increased. In post2 rest session, replay events seem to be more frequent and most of them are associated with maze2 context.

Conclusion

As the results have shown, without referring to any external variables, P-GPLVM is a powerful tool to capture nonlinear neural population dynamics in the hippocampus by discovering the low-dimensional structure underlying the neural activity, which reveals important encoded external variables such as the number of contexts and the topology of animal spatial behavior (i.e. position). In this paper, this model is extended to leverage

those advantages in the neural decoding scenario by enabling the constrained latent variable inference of new neural data and proposing a family of analyses for result evaluations. This constrained inference approach requires much less computational cost than relearning the model with training data combining old and new data. This extended model is flexible that, for new neural data either during running or during PBEs, neural trajectories can be inferred in the latent space learned from training data during running and internal neural state repetition can be evaluated. External variables can then be decoded based on the external experiences corresponding to the repeated neural states. Metrics are defined that for the first time enable the identification of continuously-decoded replay using a model trained only with neural activity (and refine this definition to allow for both sequential and non-sequential events).

Acknowledgments

This research was supported by R01NS115233.

References

Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., and Moser, M.-B. (2014). Place cells in the hippocampus: eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences*, 111(52):18428–18435.

Cai, D. J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., Wei, B., Veshkini, M., La-Vu, M., Lou, J., Flores, S. E., Kim, I., Sano, Y., Zhou, M., Baumgaertel, K., Lavi, A., Kamata, M., Tuszyński, M., Mayford, M., Golshani, P., and Silva, A. J.

(2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature*, 534(7605):115–118.

Csicsvari, J., O’Neill, J., Allen, K., and Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of ca1 pyramidal cells during sharp waves in open-field exploration. *European Journal of Neuroscience*, 26(3):704–716.

Davidson, T. J., Kloosterman, F., and Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, 63(4):497–507.

Diba, K. and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241–1242.

Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. *Journal of Neurophysiology*, 117(4):1785–1796.

Eichenbaum, H. and Cohen, N. J. (2014). Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron*, 83(4):764–770.

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

Harrison, M. T., Amarasingham, A., and Kass, R. E. (2013). of synchronous spiking. *Spike Timing: Mechanisms and Function*, page 77.

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

Krause, E. L. and Drugowitsch, J. (2022). A large majority of awake hippocampal

sharp-wave ripples feature spatial trajectories with momentum. *Neuron*, 110(4):722–733.e8.

Kudrimoti, H. S., Barnes, C. A., and McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and eeg dynamics. *Journal of Neuroscience*, 19(10):4090–4101.

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

Moser, M.-B., Rowland, D. C., and Moser, E. I. (2015). Place Cells, Grid Cells, and Memory. *Cold Spring Harbor Perspectives in Biology*, 7(2):a021808.

Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience*, 19(21):9497–9507.

Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., Pinto, L., Gauthier, J. L., Brody, C. D., and Tank, D. W. (2021). Geometry of abstract learned knowledge in the hippocampus. *Nature*, 595(7865):80–84.

Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., and Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. *Nature communications*, 10(1):1–14.

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

Wilson, M. A. and McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679.

Wu, A., Roy, N. A., Keeley, S., and Pillow, J. W. (2017). Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Advances in neural information processing systems*, 30.

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, 102(1):614–635. PMID: 19357332.