A nonparametric Bayesian approach for uncovering rat hippocampal population codes during spatial navigation

Scott W. Linderman^a, Matthew J. Johnson^{a,b}, Matthew A. Wilson^c, Zhe Chen^{*d}

^a*Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA*

^b*Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA*

^c*Picower Institute for Learning and Memory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

^d*Department of Psychiatry, Department of Neuroscience and Physiology, New York University School of Medicine, New York, NY 10016, USA*

Abstract

Background: Rodent hippocampal population codes represent important spatial information about the environment during navigation. Computational methods have been developed to uncover the neural representation of spatial topology embedded in rodent hippocampal ensemble spike activity.

New method: We extend our previous work and propose a novel nonparametric Bayesian approach to infer rat hippocampal population codes during spatial navigation. To tackle the model selection problem, we leverage a nonparametric Bayesian model. Specifically, we apply a hierarchical Dirichlet process-hidden Markov model (HDP-HMM) using two Bayesian inference methods, one based on Markov chain Monte Carlo (MCMC) and the other based on variational Bayes (VB).

Results: The effectiveness of our Bayesian approaches is demonstrated on recordings from a freely-behaving rat navigating in an open field environment.

Comparison with existing methods: The HDP-HMM outperforms the finite-state HMM in both simulated and experimental data. For HPD-HMM, the MCMC-based inference with Hamiltonian Monte Carlo (HMC) hyperparameter sampling is flexible and efficient, and outperforms VB and MCMC approaches with hyperparameters set by empirical Bayes.

Conclusion: The nonparametric Bayesian HDP-HMM method can efficiently perform model selection and identify model parameters, which can be used for modeling latent-state neuronal population dynamics.

Keywords: Hidden Markov model, Stick-breaking process, Hierarchical Dirichlet process, Nonparametric Bayesian inference, Variational Bayes,

Email addresses: `slinderman@seas.harvard.edu` (Scott W. Linderman), `mattjj@csail.mit.edu` (Matthew J. Johnson), `mwilson@mit.edu` (Matthew A. Wilson), `Zhe.Chen3@nyumc.org` (*Corresponding author) (Zhe Chen*)

Preprint submitted to Journal of Neuroscience Methods

January 25, 2016

1. Introduction

A fundamental goal in neuroscience is to understand how populations of neurons represent and transmit information about the external world. The hippocampus is known to encode information relevant to spatial navigation and episodic memory. Spatial representation of the environment is pivotal for navigation in rodents (O’Keefe and Nadel, 1978). One type of spatial representation is a topological map, which contains only relative ordering or connectivity information between spatial locations and is invariant to orientation or deformation. A relevant question of interest is: how can neurons downstream of the hippocampus infer representations of space from hippocampal spike activity without *a priori* place field information (namely, without the measurement of spatial correlates)? Several reports have been dedicated to the mathematical analysis of this problem (Curto and Itskov, 2008; Dabaghian et al., 2012, 2014); however, a data-driven approach for analyzing ensemble hippocampal spike data remains missing. This paper employs probabilistic modeling and inference methods to uncover the spatial representation (or topological map) based on the ensemble spike activity.

Bayesian statistical modeling is a consistent and principled framework for dealing with uncertainties about the observed data (Scott, 2002). The goal of Bayesian inference is to incorporate prior knowledge and constraints of the problem and to infer the posterior distribution of unobserved variables of interest (Gelman et al., 2013). In recent years, cutting-edge Bayesian methods have become increasingly popular for data analyses in neuroscience, medicine and biology (Mishchenko et al., 2011; Chen et al., 2011; Chen, 2013; Davidson et al., 2009; Kloosterman et al., 2014; Yau et al., 2011). Specifically, thanks to ever-growing computing power, Markov chain Monte Carlo (MCMC) methods have been widely used in Bayesian inference.

In our previous work (Chen et al., 2012a, 2014), we have developed a *parametric Bayesian* approach to uncover the neural representation of spatial topology embedded in rodent hippocampal population codes during spatial navigation. Here we extend the preceding work and consider a *nonparametric Bayesian* approach. The nonparametric Bayesian method brings additional flexibility to the probabilistic model, which allows us to model the complex structure of neural data (Teh and Jordan, 2010; Wood and Black, 2008; Shalchyan and Farina, 2014). Specifically, we leverage the so-called a hierarchical Dirichlet process-HMM (HDP-HMM) (Teh et al., 2006), which extends the finite-state hidden Markov model (HMM) with a nonparametric, HDP prior and derive corresponding Bayesian inference algorithms. We consider both deterministic and stochastic approaches for fully Bayesian inference. Based on deterministic approximation, we extend the work of (Chen et al., 2012a; Johnson and Willsky,

2014) and use a variational Bayes (VB) method for approximate Bayesian inference. For MCMC, we adapt the Gibbs sampling approach of (Teh et al., 2006), and integrate it with a Hamiltonian Monte Carlo (HMC) method (Neal, 2010) for hyperparameter inference. To the best of our knowledge, the application of the HDP-HMM to hippocampal ensemble neuronal spike trains and the HMC hyperparameter inference algorithm is novel.

We test the statistical model and inference methods with both simulation data and experimental data. The latter consists of a recording of rat dorsal hippocampal ensemble spike activity during open field navigation. Using a decoding analysis and predictive likelihood, we verify and compare the performance of the proposed Bayesian inference algorithms. We also discuss the results of model selection related to the sample size and the choice of concentration parameter or hyperparameters. Our methods provide an extended tool to analyze rodent hippocampal population codes, which may further empower us to explore important neuroscience questions about neural representation, learning and memory.

2. Methods: modeling and inference

2.1. Basic probabilistic model

In our previous work (Chen et al., 2012a), we used a finite m -state HMM to characterize the population spiking activity from a population of C hippocampal place cells. It was assumed that first, the animal’s spatial location during locomotion, modeled as a latent state process, followed a first-order discrete-state Markov chain $\mathcal{S} = S_{1:T} \equiv \{S_t\} \in \{1, \dots, m\}$, and second, the spike counts of individual place cells at time t , conditional on the hidden state S_t , followed a Poisson probability with their respective tuning curve functions $\Lambda = \{\lambda_c\} = \{\{\lambda_{c,i}\}\}$. Essentially, we employed a Markov-driven population Poisson firing model with the following probabilistic models

$$\begin{aligned}
p(\mathbf{y}_{1:T}, S_{1:T} | \boldsymbol{\pi}, \mathbf{P}, \Lambda) &= p(S_1 | \boldsymbol{\pi}) \prod_{t=2}^T p(S_t | S_{t-1}, \mathbf{P}) \prod_{t=1}^T p(\mathbf{y}_t | S_t, \Lambda), \quad (1) \\
p(S_1 | \boldsymbol{\pi}) &= \text{Multinomial}(S_1 | \boldsymbol{\pi}), \\
p(S_t | S_{t-1}, \mathbf{P}) &= \text{Multinomial}(S_t | \mathbf{P}_{S_{t-1}, :}), \\
p(\mathbf{y}_t | S_t, \Lambda) &= \prod_{c=1}^C \text{Poisson}(y_{c,t} | \lambda_{c,S_t}).
\end{aligned}$$

where $\mathbf{P} = \{P_{ij}\}$ denotes an m -by- m state transition probability matrix, with P_{ij} representing the transition probability from state i to j (since $\sum_{k=1}^m P_{ik} = 1$, each row of matrix $\mathbf{P}$ specifies a multinomial likelihood); $y_{c,t}$ denotes the number of spike counts from cell c within the t -th temporal bin (here we assume that the rate is defined in the unit bin size of 250 ms) and $\mathbf{y}_{1:T} = \{y_{c,t}\}_{C \times T}$ denotes time series of C -dimensional population response vector; and $\text{Poisson}(y_{c,t} | \lambda_{c,i})$ defines a Poisson distribution with the rate parameter $\lambda_{c,i}$ when $S_t = i$. Finally,

$\log p(\mathbf{y}_{1:T}|\mathcal{S}, \boldsymbol{\theta})$ defines the observed data log likelihood given the hidden state sequence $\mathcal{S}$ and all parameters $\boldsymbol{\theta} = \{\boldsymbol{\pi}, \mathbf{P}, \boldsymbol{\Lambda}\}$ (where $\boldsymbol{\pi} = \{\pi_i\}$ denotes a probability vector for the initial state S_1).

The hidden variables $\mathcal{S} = \{S_{1:T}\}$ are treated as the missing data, $\mathbf{y}_{1:T}$ as the observed (incomplete) data, and their combination $\{S_{1:T}, \mathbf{y}_{1:T}\}$ as the complete data.

A Bayesian version of this model introduces prior distributions over the parameters. We use the following prior distributions,

$$\begin{aligned}\alpha_0 &\sim \text{Gamma}(a_{\alpha_0}, 1) \\ \boldsymbol{\pi} &\sim \text{Dir}(\alpha_0 \mathbf{1}), \\ \mathbf{P}_{i,:} &\sim \text{Dir}(\alpha_0 \mathbf{1}), \\ \lambda_{c,i} &\sim \text{Gamma}(a_c^0, b_c^0).\end{aligned}\tag{2}$$

where Dir denotes the Dirichlet prior distribution, and $\text{Gamma}(a_c^0, b_c^0)$ denotes the gamma prior distribution with shape parameter a_c^0 and scale parameter b_c^0 .

2.2. HDP-HMM

Model selection is an important issue for statistical modeling and data analysis. We have previously proposed a *Bayesian deviance information criterion* to select the model size m of HMM (Chen et al., 2012a, 2014). Here we extend the finite-state HMM to an HDP-HMM, a nonparametric Bayesian extension of the HMM that allows for a potentially infinite number of hidden states (Teh et al., 2006; Beal et al., 2002). Namely, the HDP-HMM treats the priors via a stochastic process. Instead of imposing a Dirichlet prior distribution on the rows of the finite state transition matrix $\mathbf{P}$, we use a HDP that allows for a countably infinite number of states.

Specifically, we sample a distribution over latent states, G_0 , from a Dirichlet process (DP) (Ferguson, 1973) prior, $G_0 \sim \text{DP}(\gamma, H)$, where γ is the concentration parameter and H is the base measure. Moreover, we place a prior distribution over the concentration parameter, $\gamma \sim \text{Gamma}(a_\gamma, 1)$. Given the concentration, one may sample from the DP via the “stick-breaking construction” (Sethuraman, 1994). First, sample the stick-breaking weights, $\boldsymbol{\beta}$,

$$\tilde{\beta}_i \sim \text{Beta}(1, \gamma), \quad \beta_i = \tilde{\beta}_i \prod_{j=1}^{i-1} (1 - \beta_j)\tag{3}$$

where $\beta_1 = \tilde{\beta}_1$, $\sum_{i=1}^{\infty} \beta_i = 1$, and $\text{Beta}(a, b)$ defines a beta distribution with two positive shape parameters a and b .

The stick-breaking construction of (3) is sometimes denoted as $\boldsymbol{\beta} \sim \text{GEM}(\gamma)$, after Griffiths, Engen, and McCloskey (Ewens, 1990). The name “stick-breaking” comes from the interpretation of β_i as the length of the piece of a unit-length stick assigned to the i -th value. After the first $i - 1$ values having their portions assigned, the length of the remainder of the stick is broken according to a sample $\tilde{\pi}_i$ from a beta distribution, and $\tilde{\beta}_i$ indicates the portion of the remainder to

be assigned to the i -th value. Therefore, the stick-breaking process $\text{GEM}(\gamma)$ also defines a DP—the smaller γ , the less (in a statistical sense) of the stick will be left for subsequent values.

After sampling β , we next sample the latent state variables, in this case λ_c , from the base measure H . Our draw from the $\text{DP}(\gamma, H)$ prior is then given by

$$G_0 = \sum_{j=1}^{\infty} \beta_j \delta_{\lambda_c^{(j)}}. \quad (4)$$

Thus, the stick breaking construction makes clear that draws from a Dirichlet process distribution are discrete with probability one.

Given a countably infinite set of shared states, we may then sample the rows of the transition matrix, $\mathbf{P}_{i,:} \sim \text{DP}(\alpha_0, \beta)$. We place the same prior over π . The base measure in this case is β , a countably infinite vector of stick-breaking weights, that serves as the mean of the DP prior over the rows of $\mathbf{P}$. The concentration parameter, α_0 , governs how concentrated the rows are about the mean. Since the base measure β is discrete, each row of $\mathbf{P}$ will be able to “see” the same set of states. By contrast, if we remove the HDP prior and treat each row of $\mathbf{P}$ as an independent draw from a DP with base measure H , each row would see a disjoint set of states with probability one. In other words, the hierarchical prior is required to provide a discrete (but countably infinite) set of latent states for the HMM.

2.3. Overdispersed Poisson model

An interesting consequence of this Bayesian model is that it naturally leads to a distribution of spike counts that is overdispersed relative to simple Poisson model, a feature that has been observed in neural recordings (Goris et al., 2014). Recent work has explored the negative binomial (NB) distribution as an alternative to Poisson model, since its two parameters allow for Fano factors greater than one. The NB distribution can also be seen as a continuous mixture of Poisson distributions (i.e., a compound probability distribution) where the mixing distribution of the Poisson rate is a gamma distribution (Gelman et al., 2013). In other words, the NB distribution is viewed as a gamma-Poisson (mixture distribution): a $\text{Poisson}(\lambda)$ distribution whose rate λ is itself a gamma random variable. In our case, the gamma prior over firing rates leads to a negative binomial marginal distribution over $y_{c,t}$.

Though the marginal spike count at a particular time t may be marginally distributed according to a negative binomial distribution, it is not necessarily true that a sequence of time bins, $\mathbf{y}_{1:T}$, will be i.i.d. negative binomials. This arises from the correlations induced by the state transition matrix. Instead, $\mathbf{y}_{1:T}$ will follow a finite mixture of Poisson distributions, with one component for each latent state. The mixture will be weighted by the marginal probability of the corresponding latent state. However, as the number of visited states grows, and the marginal probability of latent states becomes more uniform, the resulting marginal distribution over the spike count sequences inherits the over dispersed

nature of the NB distribution. This is particularly true of an HDP-HMM with a high concentration.

2.4. Markov chain Monte Carlo (MCMC) inference

Several MCMC-based inference methods have been developed for the HDP-HMM (Beal et al., 2002; Teh et al., 2006; Van Gael et al., 2008). Some of these previous works use a collapsed Gibbs sampler in which the transition matrix P and the observation parameters $\mathbf{\Lambda}$ are integrated out (Teh et al., 2006; Van Gael et al., 2008). In this work, however, we use a “weak limit” approximation in which the DP prior is approximated with a symmetric Dirichlet prior. Specifically, we let

$$\begin{aligned}\gamma &\sim \text{Gamma}(a_\gamma, 1) \\ \alpha_0 &\sim \text{Gamma}(a_{\alpha_0}, 1) \\ \beta | \gamma &\sim \text{Dir}(\gamma/M, \dots, \gamma/M), \\ \pi | \alpha_0, \beta &\sim \text{Dir}(\alpha_0\beta_1, \dots, \alpha_0\beta_M), \\ \mathbf{P}_{i,:} | \alpha_0, \beta &\sim \text{Dir}(\alpha_0\beta_1, \dots, \alpha_0\beta_M).\end{aligned}\tag{5}$$

where M denotes a truncation level for approximating the infinity (which is different from m in the finite-state setting). It can be shown that this prior will weakly converge to the DP prior as the dimensionality of the Dirichlet distribution approaches infinity (Johnson and Willsky, 2014; Ishwaran and Zarepour, 2002). With this approximation we can capitalize on forward-backward sampling algorithms to jointly update the latent states $\mathcal{S}$.

Previous work has typically been presented with Gaussian or multinomial likelihood models, with the acknowledgement that the same methods work with any exponential family likelihood when the base measure H is a conjugate prior. Here we present the Gibbs sampling algorithm of (Teh et al., 2006) for the HDP-HMM applied to the special case of independent Poisson observations, and we derive Hamiltonian Monte Carlo (HMC) transitions to sample the cell-specific hyperparameters of the firing rate priors.

We begin by defining Gibbs updates for the neuronal firing rates $\mathbf{\Lambda}$. Since we are using gamma priors with independent Poisson observations, the model is fully conjugate and simple Gibbs updates suffice. Therefore, we have

$$\lambda_{c,i} | \mathbf{y}, \mathcal{S} \sim \text{Gamma} \left(\alpha_c^0 + \sum_{t=1}^T y_{c,t} \mathbb{I}[S_t = i], \beta_c^0 + \sum_{t=1}^T \mathbb{I}[S_t = i] \right). \tag{6}$$

Under the weak limit approximation, the priors on $\mathbf{P}_{i,:}$ and π reduce to Dirichlet distributions, which are also conjugate with the finite HMM. Hence we can derive conjugate Gibbs updates for these parameters as well. They take

the form:

$$\begin{aligned}\boldsymbol{\pi} | \alpha_0, \boldsymbol{\beta} &\sim \text{Dir}(\alpha_0 \boldsymbol{\beta} + \mathbf{1}_{S_1}), \\ \mathbf{P}_{i,:} | \alpha_0, \boldsymbol{\beta} &\sim \text{Dir}(\alpha_0 \boldsymbol{\beta} + \mathbf{n}_i), \\ n_{i,j} &= \sum_{t=1}^{T-1} \mathbb{I}[S_t = i, S_{t+1} = j],\end{aligned}\tag{7}$$

where $\mathbf{1}_j$ is a unit vector with a one in the j -th entry.

Conditioned upon the firing rates, the initial state distribution, and the transition matrix, we can jointly update the latent states of the HDP-HMM using a *forward filtering, backward sampling* algorithm. Details of this well-known algorithm can be found in (Johnson, 2014); the intuition is that in the backward pass of the algorithm, we have a conditional distribution over S_t given $S_{t+1:T}$. We can iteratively sample from these distributions as we go backward in time to generate a full sample from $p(\mathcal{S} | \mathbf{P}, \boldsymbol{\pi}, \boldsymbol{\Lambda})$. Jointly sampling these latent states allows us to avoid issues with mixing when individually sampling states that are highly correlated with one another.

Finally, the Dirichlet parameters $\boldsymbol{\beta}$ and the concentration parameters α_0 and γ can be updated as in (Teh et al., 2006). A single iteration of the final algorithm consists of an update for each parameter of the model. The aforementioned updates are based upon previous work; one novel direction that we explore in this work is the sampling of the hyperparameters of the gamma firing rate priors.

2.4.1. Setting firing rate hyperparameters

We consider three approaches to setting the hyperparameters of the gamma priors for Poisson firing rates, namely, $\{\alpha_c^0, \beta_c^0\}$ for cell c . Note that these parameters are distinct from the parameters of the DP prior.

- In the first approach, we estimate these parameters using an empirical Bayesian (EB) procedure, that is, by maximizing the marginal likelihood of the spike counts. For each cell, this may be easily done using standard maximum likelihood estimation for the negative binomial model. In practice, we found that without regularization this approach leads to extreme values of the hyperparameters.
- Our second approach samples these hyperparameters using Hamiltonian Monte Carlo (HMC) (Neal, 2010). We note that for fixed values of the “shape” parameter α_c^0 , the conditional distribution of the “scale” parameter, β_c^0 is conjugate with a gamma prior distribution. However, setting the shape parameter *a priori* is challenging because it can have a strong influence on the firing rate distribution. HMC allows us to jointly sample both the shape and the scale parameters simultaneously. To implement HMC, we must have access to both the log probability of the parameters as well as its gradient. Since both parameters are restricted to be positive,

we instead reparameterize the problem in terms of their logs. For cell c , the conditional log probability equal to,

$$\begin{aligned}\mathcal{L} &= \log p(\log \alpha_c^0, \log \beta_c^0 \mid \mathbf{\Lambda}_{c,\cdot}) \\ &= \sum_{i=1}^m \log p(\lambda_{c,i} \mid \alpha_c^0, \beta_c^0) + \text{const.} \\ &= \sum_{i=1}^m \alpha_c^0 \log \beta_c^0 - \log \Gamma(\alpha_c^0) + (\alpha_c^0 - 1) \log \lambda_{c,i} - \beta_c^0 \lambda_{c,i}.\end{aligned}\tag{8}$$

Taking gradients with respect to both parameters yields,

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial \log \alpha_c^0} &= \sum_{i=1}^m [\log \beta_c^0 - \Psi(\alpha_c^0) + \log \lambda_{c,i}] \times \alpha_c^0, \\ \frac{\partial \mathcal{L}}{\partial \log \beta_c^0} &= \sum_{i=1}^m \left[\frac{\alpha_c^0}{\beta_c^0} - \lambda_{c,i} \right] \times \beta_c^0.\end{aligned}\tag{9}$$

The HMC algorithm uses these gradients to inform a stochastic walk over the posterior distribution. With knowledge of the gradients, HMC can sometimes make large updates to parameters, especially in cases where the parameters are highly correlated under the posterior.

- In the final approach, we fix the shape hyperparameter, α_c^0 , and infer the scale hyperparameter, β_c^0 . We place a gamma prior on the scale, $\beta_c^0 \sim \text{Gamma}(\mu, \nu)$. Given α_c^0 , the conditional distribution of the scale hyperparameter is

$$\beta_c^0 \mid \alpha_c^0, \{\lambda_{c,i}\}, \mathcal{S} \sim \text{Gamma}\left(\mu + \sum_{i=1}^m \mathbb{I}[n_i > 0] \cdot \alpha_c^0, \nu + \sum_{i=1}^m \mathbb{I}[n_i > 0] \cdot \lambda_{c,i}\right)\tag{10}$$

$$n_i = \sum_{t=1}^T \mathbb{I}[S_t = i].$$

In the following experiments, we set the shape parameter to $\alpha_c^0 = 1$, and we set the scale prior parameters to $\mu = 1$ and $\nu = 1$. This is equivalent to an exponential prior on rates, $\lambda_{c,i} \sim \text{Exp}(\beta_c^0)$, and an exponential prior on the scale $\beta_c^0 \sim \text{Exp}(1)$. One could perform cross validation over the shape parameter, but the exponential prior is a rather weak assumption that enables fully-Bayesian inference.

2.4.2. Predictive log likelihood

Upon completing the parameter and hyperparameter inference, we evaluate the performance of our algorithm in terms of its predictive log likelihood on held

out test data. We approximate the predictive log likelihood with samples from the posterior distribution generated by our MCMC algorithm. That is,

$$\begin{aligned} \log p(\mathbf{y}_{test} | \mathbf{y}_{1:T}) &= \log \sum_{\mathcal{S}_{test}} \int_{\Theta} p(\mathbf{y}_{test}, \mathcal{S}_{test} | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \mathbf{y}_{train}) d\boldsymbol{\theta}, \\ &\approx \log \frac{1}{N} \sum_{n=1}^N \sum_{\mathcal{S}_{test}} p(\mathbf{y}_{test}, \mathcal{S}_{test} | \boldsymbol{\theta}_n), \end{aligned} \quad (11)$$

where $\boldsymbol{\theta} = (\boldsymbol{\Lambda}, \mathbf{P}, \boldsymbol{\pi})$ and $\{\boldsymbol{\theta}_n\}_{n=1}^N \sim p(\boldsymbol{\theta} | \mathbf{y}_{train})$. The summation over latent state sequences for the test data is performed with the message-passing algorithm for HMMs.

2.5. Variational Bayes (VB) inference

We build upon our previous work (Chen et al., 2012a, 2014; Johnson and Willsky, 2014) to develop a variational inference algorithm for fitting the HDP-HMM to hippocampal spike trains. Our objective is to approximate the posterior distribution of the HDP-HMM with a distribution from a more tractable family. As usual, we choose a factorized approximation that allows for tractable optimization of the parameters of the variational model. Specifically, we let,

$$p(\mathcal{S}, \boldsymbol{\Lambda}, \mathbf{P}, \boldsymbol{\pi}, \boldsymbol{\beta} | \mathbf{y}_{1:T}) \approx q(\mathcal{S})q(\boldsymbol{\Lambda})q(\mathbf{P})q(\boldsymbol{\pi})q(\boldsymbol{\beta}). \quad (12)$$

Since the independent Poisson observations are conjugate with the gamma firing rate prior distributions, choosing a set of independent gamma distributions for $q(\boldsymbol{\Lambda})$ allows for simple variational updates.

$$\begin{aligned} q(\boldsymbol{\Lambda}) &= \prod_{i=1}^M \prod_{c=1}^C \text{Gamma}(\tilde{\alpha}_{c,i}, \tilde{\beta}_{c,i}), \\ \tilde{\alpha}_{c,i} &\leftarrow \alpha_c^0 + \sum_{t=1}^T y_{c,t} \mathbb{E}_q[\mathbb{I}[S_t = i]], \\ \tilde{\beta}_{c,i} &\leftarrow \beta_c^0 + \sum_{t=1}^T \mathbb{E}_q[\mathbb{I}[S_t = i]]. \end{aligned} \quad (13)$$

Following (Johnson and Willsky, 2014), we use a “direct assignment” truncation for the HDP (Bryant and Sudderth, 2012; Liang et al., 2007). In this scheme, a truncation level M is chosen *a priori* and $q(\mathcal{S})$ is limited to support only states $S_t \in \{1, \dots, M\}$. The advantage of this approximation is that conjugacy is retained with $\boldsymbol{\Lambda}$, $\mathbf{P}$, and $\boldsymbol{\pi}$, and the variational approximation $q(\mathcal{S})$ reduces to

$$\begin{aligned} q(\mathcal{S}) &= \text{HMM}(\tilde{\mathbf{P}}, \tilde{\boldsymbol{\pi}}, \tilde{\boldsymbol{\Lambda}}), \\ \tilde{\mathbf{P}} &= \exp \{ \mathbb{E}_q[\ln \mathbf{P}] \}, \\ \tilde{\boldsymbol{\pi}} &= \exp \{ \mathbb{E}_q[\ln \boldsymbol{\pi}] \}, \\ \tilde{\boldsymbol{\Lambda}} &= \exp \{ \mathbb{E}_q[\ln p(\mathbf{y} | \boldsymbol{\Lambda})] \}. \end{aligned} \quad (14)$$

Expectations $\mathbb{E}_q[S_t = i]$ can then be computed using standard the message-passing algorithm for the HMM.

With the direct assignment truncation, the variational factors for $\mathbf{P}_{i,:}$ and $\boldsymbol{\pi}$ take on Dirichlet priors. Unlike in the finite-state HMM, however, these Dirichlet priors are now over $M + 1$ dimensions since the final dimension accounts for all states $i > M$. Under the HDP prior we had $P_{i,:} \sim \text{DP}(\alpha_0 \cdot \boldsymbol{\beta})$, and under the truncation the DP parameter becomes $\alpha_0 \cdot \boldsymbol{\beta}_{1:M+1}$. Again, leveraging the conjugacy of the model, we arrive at the following variational updates:

$$q(\mathbf{P}) = \prod_{i=1}^M \text{Dir}(\tilde{\boldsymbol{\alpha}}_P^{(i)}), \quad (15)$$

$$(\tilde{\boldsymbol{\alpha}}_P^{(i)})_j \leftarrow \alpha_0 \beta_j + \mathbb{E}_q[\mathbb{I}[S_t = i] \cdot \mathbb{I}[S_{t+1} = j]].$$

We use an analogous update for $\boldsymbol{\pi}$.

The principal drawback of the direct assignment truncation is that the prior for $\boldsymbol{\beta}$ is no longer conjugate. This could be avoided with the fully conjugate approach of (Hoffman et al., 2013), however, this results in extra bookkeeping and the duplication of states. Instead, following (Johnson and Willsky, 2014; Bryant and Sudderth, 2012; Liang et al., 2007), we use a point estimate for this parameter by setting $q(\boldsymbol{\beta}) = \delta_{\boldsymbol{\beta}^*}$ and use gradient ascent with backtracking line search to update this parameter during inference.

There are a number of hyperparameters to set for the VB approach. The hyperparameters α_c^0 and β_c^0 of gamma prior on firing rates can be set with empirical Bayes, as above. We resort to cross validation to set the Dirichlet parameter α_0 and the GEM parameter γ .

Finally, in order to compute the predictive log likelihood on held out test data, we draw multiple samples $\{(\boldsymbol{\Lambda}, \mathcal{S}, \mathbf{P}, \boldsymbol{\pi}, \boldsymbol{\beta})_n\}_{n=1}^N$ from the variational posterior and approximate the predictive log likelihood as

$$\begin{aligned} \ln p(\mathbf{y}_{test} | \mathbf{y}_{1:T}) &\approx \ln \mathbb{E}_q [p(\mathbf{y}_{test} | \mathcal{S}, \boldsymbol{\Lambda}, \mathbf{P}, \boldsymbol{\pi}, \boldsymbol{\beta})] \\ &\approx \ln \frac{1}{N} \sum_{n=1}^N p(\mathbf{y}_{test} | (\mathcal{S}, \boldsymbol{\Lambda}, \mathbf{P}, \boldsymbol{\pi}, \boldsymbol{\beta})_n). \end{aligned} \quad (16)$$

3. Results

The inference algorithms were implemented based upon the PyHSMM framework of (Johnson, 2014). The codebase was written in Python with C offloads for the message passing algorithms. We have extended the codebase to perform hyperparameter inference using the methods described above, and expanded it to tailor to neural spike train analysis. Our source code is publicly available (https://github.com/slinderman/pyhsmm_spiketrains).

3.1. Simulation data

Setup. First, we simulate synthetic spike count data using an HDP-HMM with $C = 50$ neurons, $T = 2000$ time bins, and Dirichlet concentration parameters $\alpha_0 =$

Figure 1: An example of a synthetic dataset drawn from an HDP-HMM. (A) Simulated population spike trains or spike counts. (B) Inferred latent state sequence. (C) Inferred state transition matrix P . (D) Inferred neuronal firing rate vectors λ_i specific to each state.

12.0 and $\gamma = 12.0$. These configuration yield state sequences that tend to visit 30-45 states. All of neuronal firing rate parameters are drawn from a gamma distribution: $\text{Gamma}(\alpha_c^0 = 1, \beta_c^0 = 1)$ (with mean 1.0 and standard deviation 1.0).

An example of one such synthetic dataset is shown in Fig. 1. The states have been ordered according to their occupancy (i.e., how many times they are visited during the simulation), such that the columns of the transition matrix exhibit a decrease in probability as the incoming state number, S_{t+1} , increases. This is a characteristic of the HDP-HMM, indicating the tendency of the model to reuse states with high occupancy.

We compare six combinations of model, inference algorithm, and hyperparameter selection approaches: (i) HMM with the correct number of states, fit by Gibbs sampling with fixed $\alpha_c^0 = 1$; (ii) HMM with the correct number of states, fit by VB with hyperparameters set by empirical Bayes; (iii) HDP-HMM fit by Gibbs sampling with fixed $\alpha_c^0 = 1$; (iv) HDP-HMM fit by Gibbs sampling and HMC for hyperparameter updates; (v) HDP-HMM fit by MCMC with hyperparameters set by empirical Bayes; and (vi) HDP-HMM fit by VB with hyperparameters set by empirical Bayes. For the MCMC methods, we set gamma priors over the concentration parameters (α_0 and γ); for the VB methods, we set α_0 and γ to their true values. Alternatively, they can be selected by cross validation. We set both the weak limit approximation for MCMC and the direct assignment truncation level for VB to $M = 100$.

We collect 5000 samples from the MCMC algorithms and use the last 2000 samples for computing predictive log likelihoods. For visualization, we use the final sample to extract the transition matrix and the firing rates. The number of samples and the amount of burn-in iterations were chosen by examining the log probability and parameter traces for convergence. It is found that the MCMC algorithm converges within hundreds of iterations. For further convergence diagnosis of a single Gibbs chain, one may use the autocorrelation tools suggested in (Raftery and Lewis, 1992; Cowles and Carlin, 1996).

We run the VB algorithm for 200 steps to guarantee convergence of the variational lower bound. Again, this is assessed by examining the variational lower bound and is found to converge to a local maxima within tens of iterations.

Assessment. We use two criteria for result assessment with simulation data. The first criterion is based on the Hamming error between the true and inferred state sequences. To compute this, we first relabel the inferred states in order to maximize overlap with the true states. Let $\mathcal{S}$ be the true state sequence and $\mathcal{S}'$ be the inferred state sequence. We define the overlap matrix $O \in \mathbb{N}^{M \times M}$ whose entries $O_{i,j}$ is the number of times the true state is i and the inferred state is j :

$$O_{i,j} = \sum_{t=1}^T \mathbb{I}[S_t = i] \mathbb{I}[S'_t = j]. \quad (17)$$

We use the Hungarian method (Kuhn, 1955) to find a relabeling of the inferred states that maximizes overlap, and then we measure the Hamming error between

Table 1: Comparison of Hamming error (see Eq. 18) computed from the same nine simulated data sets as above. The VB inference methods tend to overestimate the number of states and therefore have much larger Hamming error.

Dataset	1	2	3	4	5
HMM (Gibbs)	9	401	13	24	615
HMM (VB)	166	290	295	123	124
HDP-HMM (Gibbs)	2	3	5	1	6
HDP-HMM (HMC)	3	4	3	2	4
HDP-HMM (EB)	1	3	2	3	12
HDP-HMM (VB)	432	586	340	264	675

the true state sequence $\mathcal{S}$, and the relabeled sequence of inferred states, $\tilde{\mathcal{S}}'$:

$$\text{err}(\mathcal{S}, \tilde{\mathcal{S}}') = \sum_{t=1}^T \mathbb{I}[S_t \neq \tilde{S}'_t]. \quad (18)$$

Table 1 summarizes the Hamming error for all six models on five synthetic datasets. We see that the HDP-HMM fit via Gibbs sampling with firing rate hyperparameters set via empirical Bayes outperforms the other models and inference algorithms on three of five datasets, but the HDP-HMM with hyperparameter HMC sampling are very comparable. By contrast, when the models are fit with VB inference, the inferred state sequences tend to use more than the true number of states, which results in very poor Hamming error. Similarly, the HMM fit via Gibbs sampling does not factor in the penalty on additional states and instead tends to use all states equally, resulting in large Hamming error.

The second criterion is the model’s predictive log likelihood (unit: bits/spike) on a held out sequence of $T_{\text{test}} = 1000$ time steps. We compare the predictive log likelihood to that of a set of independent Poisson processes. Their rates and the corresponding predictive log likelihood are given by,

$$\hat{\lambda}_c = \frac{1}{T_{\text{train}}} \sum_{t=1}^{T_{\text{train}}} y_{c,t}, \quad (19)$$

$$\log p(\mathbf{y}_{\text{test}} | \mathbf{y}_{\text{train}}) = \sum_{c=1}^C \left[-T_{\text{test}} \hat{\lambda}_c + \sum_{t=1}^{T_{\text{test}}} y_{c,t} \log \hat{\lambda}_c \right]. \quad (20)$$

The improvement obtained by a model is measured in bits, and is normalized by the number of spikes in the test dataset in order to obtain comparable units for each of the test datasets.

Table 2 summarizes the predictive log likelihood comparison. For all five datasets, the HDP-HMM fit via Gibbs sampling with fixed α_c^0 performs best, though in general the increase over fitting the HDP-HMM when using HMC or EB for hyperparameter selection is small. By contrast, the improvement

Table 2: Comparison of predictive log likelihood computed from 5 simulated data sets, measured in bits/spike improvement over a baseline of independent, homogeneous Poisson processes (the best result in each data set is marked in bold font).

Dataset	1	2	3	4	5
HMM (Gibbs)	0.315	0.300	0.312	0.310	0.250
HMM (VB)	0.298	0.290	0.313	0.306	0.252
HDP-HMM (Gibbs)	0.323	0.307	0.321	0.318	0.259
HDP-HMM (HMC)	0.323	0.306	0.320	0.318	0.259
HDP-HMM (EB)	0.322	0.306	0.321	0.318	0.259
HDP-HMM (VB)	0.312	0.291	0.309	0.305	0.244

compared to fitting with VB inference or using a parametric HMM is quite significant.

Though computation cost is often a major factor with Bayesian inference, with the optimized PyHSMM package, the models can be fit to the synthetic data within less minutes on an Apple MacBook Air. The runtime necessarily grows the number of neurons and the truncation limit on the number of latent states. As the model complexity grows, we must also run our MCMC algorithm for more iterations, which often motivates the use of variational inference algorithms instead. Given our optimized implementation and the performance improvements yielded by MCMC, we opted for a fully-Bayesian approach using MCMC with HMC for hyperparameter sampling in our subsequent experiments.

Figure 2 shows example traces from the MCMC combined with HMC algorithm for the HDP-HMM running on synthetic dataset 1. This is the same data from which Fig. 1 is generated. The first 5 Markov chain iterations have been omitted to highlight the variation in the latter samples (the first few iterations rapidly move away from the initial conditions). We see that the log likelihood of the data rapidly converges to nearly that of the true model (horizontal dotted line), and the number of states quickly converges to around $m = 35$. Note that the nuisance parameters α_0 and γ do not converge to the true values — this is due to the fact that the solution is insensitive to these parameters or the presence of local optima. However, even the concentration parameters are different from the true values, they are still consistent with the inferred state transition matrix.

Sensitivity of the number of latent states. To test the sensitivity of the number of inferred states to changes in the data, we vary a number of parameters and plotted the number of inferred states in Fig. 3. In all cases, we use synthetic dataset 1, shown in Fig. 1, and HDP-HMMs fit via Gibbs sampling with fixed α_c^0 . First, we vary the number of neurons C , and find that the number of inferred states was relatively stable around the true number of states ($m = 35$). By contrast, as we increase the recording length T , the number of inferred states increases as well. This is because the true underlying data actually does visit more states as we simulate it for longer time. In general, we expect the number of inferred states to grow with the complexity of the data. Next, we vary the scale

Figure 2: MCMC state trajectories for an HDP-HMM fit to the synthetic dataset shown in Fig. 1. True values are shown by the dotted black lines. The first five iterations of the Markov chain are omitted since they differ greatly from the final states. The chain quickly converges to nearly the correct number of states (A) and achieves close to the true log likelihood (B). (C, D) The chain trajectories of hyperparameters α_0 and γ . (E, F) Inferred state transition matrix $\mathbf{P}$ and neuronal firing matrix $\mathbf{\Lambda} = \{\lambda_{c,i}\}$ drawn from the last iteration.

Figure 3: In a synthetic data experiment, we generated a spike train for a population of $C_{\text{true}} = 50$ cells and $T_{\text{true}} = 2000$ time bins. Then we varied the number of neurons C , recording duration T , scale of the firing rate λ , temporal bin size Δt , and inferred the number of inferred latent states from the data. Horizontal dashed lines indicate the ground truth. Box plots are obtained from 10 independent Monte Carlo chains, and each chain was run for 1000 iterations; the number of states in the last iteration is used.

of the firing rate by multiplying the true model’s firing rate by a factor of 0.1, 0.5, 1.0, 2.0, or 10.0, and sampling a new spike count. When the firing rates are very low, most bins do not contain any spikes, and hence it is not possible to resolve as many states. By contrast, when the rate is increased, the number of inferred states is slightly lower than the true number, which is likely the result of a slight mismatch with the prior on the firing rate scale (parameters μ and ν in Section 2.4.1). Finally, we considerate the effect of time bin size by scaling up the bin sizes by factors of 2 through 10. For example, when scaling by a factor of 2, we add the spike counts in each pair of adjacent bins. This has a similar effect to decreasing the recording length by a factor of 2, and hence we see the number of inferred states decrease with bin size.

3.2. Rat hippocampal neuronal ensemble data

Next, we apply the proposed methods to experimental data of the rat hippocampus. Experiments were conducted under the supervision of the Massachusetts Institute of Technology (MIT) Committee on Animal Care and followed the NIH guidelines. The micro-drive arrays containing multiple tetrodes were implanted above the right dorsal hippocampus of male Long-Evans rats. The tetrodes were slowly lowered into the brain reaching the cell layer of CA1 two to four weeks following the date of surgery. Recorded spikes were manually clustered and sorted to obtain single units using a custom software (XClust, M.A.W.).

For demonstration purpose, an ensemble spike train recording of $C = 47$ putative pyramidal neurons was collected from a single rat for a duration of 9.8 minutes. Once stable hippocampal units were obtained, the rat was allowed to freely forage in an approximately circular open field environment (radius: ~ 60 cm). We bin the ensemble spike activity with a bin size of 250 ms and obtain the population vector $\mathbf{y}_t$ in time. To identify the period of rodent locomotion during spatial navigation, we use a velocity threshold (>10 cm/s) to select the RUN epochs and merge them together. One animal’s RUN trajectory and spatial occupancy are shown in Fig. 4A and Fig. 4B, respectively. The empirical probability of a location, $p(\ell)$, is determined by dividing the arena into 220 bins of equal area (11 angular bins and 20 radial bins) and counting the fraction of time points in which the rat is in the corresponding bin.

In experimental data analysis, we focus on nonparametric Bayesian inference for HDP-HMM. For all methods, we increase the truncation level to a large value of $M = 100$. To discover the model order of the variational solutions, we use the number of states visited by the most likely state sequence under the variational posterior. The MCMC algorithms yield samples of state sequences from which the model order can be directly counted.

We perform a quantitative comparison between HMMs, HDP-HMMs, inference algorithms, and hyperparameter setting algorithms, where performance is measured in terms of both predictive log likelihood and decoding error. For both metrics, we train the models on the first 7.8 minutes of data and test on the final two minutes of data for prediction. The results are summarized in

Figure 4: One rat's behavioral trajectory (A) and spatial occupancy (B) in the open field environment.

Table 3: A comparison of HMMs, HDP-HMMs, and inference algorithms on the rat hippocampal data. Performance is measured in predictive log likelihood and mean decoding error on two minutes of held out test data (the best result is marked in bold font).

	Pred. log likelihood (bits/spike)	Decoding error (cm)
HMM ($m = 25$)	0.712	10.85 ± 6.43
HMM ($m = 45$)	0.706	10.71 ± 6.67
HMM ($m = 65$)	0.717	11.01 ± 6.93
HDP-HMM (Gibbs)	0.722	9.56 ± 5.31
HDP-HMM (HMC)	0.646	9.96 ± 6.05
HDP-HMM (EB)	0.579	10.81 ± 6.78
HDP-HMM (VB)	0.602	10.93 ± 6.24

Figure 5: Estimation result from HDP-HMM (Gibbs) for the rat hippocampal ensemble spike data. (A) Estimated state space map, where the mean value of the spatial position for each latent state is shown by a black dot. The size of the dot is proportional to the occupancy of the state. (B) Probability distributions over location corresponding to the top three latent states, measured by state occupancy. The small black dots indicate the location of the animal while in that state, and are used to compute the empirical distribution over location indicated by colored shading. (C) The true (black) and reconstructed (blue) trajectories shown in Cartesian coordinate. For each time bin, we use the mean location of the latent states to determine an estimate of the animal’s location.

Table 3. We find that the HDP-HMM fit by Gibbs sampling with fixed firing rate scale ($\alpha_c^0 = 1$) again outperforms the competing models in both measures.

For the purpose of result assessment, we plot the state-space or state-location map (Fig. 5A), which shows the mean value of the spatial position that each state represented. The size of the black dot is proportional to the occupancy of the state. To compute an “empirical” distribution over locations for a given state, we first compute the posterior distribution over latent states with our inference algorithms. This gives us a set of probabilities $\Pr(S_t = i)$ for all time bins t and states i . Then we compute the average location for each state i by weighting the animal’s location (x_t, y_t) by the probability that the animal was in state i at time t . Summing over time yields a weighted set of locations, which we then bin into equal-area arcs and normalize to get an empirical distribution over locations for each state i .

The empirical location distributions for the top three states as measured by occupancy are shown in Fig. 5B). In Fig. 5C, we show the estimated animal’s spatial trajectories in black, along with the reconstructed location in from the

Figure 6: Estimation result from HDP-HMM (Gibbs) for the rat hippocampal ensemble spike data. (A) The total number of states (solid blue) slowly increases as states are allocated for a small number of time bins. The number of states converges after 2500 iterations. (B) The log likelihood of the training data grows consistently as highly specific states are added. (C, D) The concentration parameters, α_0 and γ also converge after 2500 iterations. (E, F) The inferred state transition matrix $\mathbf{P}$ and neuronal firing matrix $\mathbf{\Lambda} = \{\lambda_{c,i}\}$ drawn from the last iteration.

Figure 7: Measuring the effect of concentration hyperparameters on the number of inferred latent states. We find that the concentration hyperparameters of the **Gamma** priors on the concentration parameters, α_0 and γ , have a minimal effect. Box plots are obtained from 10 independent Monte Carlo chains, and each chain was run for 1000 iterations; the number of states in the last iteration is used.

HDP-HMM with Gibbs sampling in blue. To reconstruct the position, we use the mean of each latent state’s location distribution weighted by the marginal probability of that state under the HDP-HMM. That is,

$$\hat{x}_t = \sum_{i=1}^m \bar{x}_i \Pr(\mathcal{S}_t = i), \quad \hat{y}_t = \sum_{i=1}^m \bar{y}_i \Pr(\mathcal{S}_t = i), \quad (21)$$

where $\bar{x}_i$ and $\bar{y}_i$ denote the average location of the rat while in inferred state i (corresponding to the black dots in Fig. 5A). Note that the animal’s position is not used in model inference, only during result assessment. In the illustrated example (HDP-HMM with MCMC+HMC), the mean reconstruction error in Euclidean distance is 9.07 cm.

As the parameter sample traces in Fig. 6 show, the Markov chain converges in around 2500 iterations. After this point, the total number of states stabilizes to around 65. The concentration parameters α_0 and γ converge within a similar number of iterations. Finally, we show the transition matrix $\mathbf{P}$ and firing rate matrix $\mathbf{A}$ obtained from the final Markov chain sample.

We again evaluate the sensitivity of these model fits to the choice of hyperparameters. For the HDP-HMM fit via Gibbs sampling with fixed α_c^0 , the primary hyperparameters of interest are the concentration hyperparameters, a_{α_0} and a_{γ} in Eq. 5, where we have assumed $\alpha_0 \sim \text{Gamma}(a_{\alpha_0}, 1)$ and $\gamma \sim \text{Gamma}(a_{\gamma}, 1)$. Figure 7 shows the inferred number of states as we vary these two hyperparameters over orders of magnitude. We find that the number of inferred states is stable around 65, indicating the performance robustness to the choice of these hyperparameters.

Looking into the inferred states, we can reconstruct the “place fields” or “state fields” of hippocampal neurons. To do so, we combine the state-location maps (Fig. 5B) with the firing rate of the individual neuron in those states (Fig. 6F) and weight by the marginal probability of the latent state. Together, these give rise to the inferred neuron’s place field. Note that, again, the position

Figure 8: (A) Comparison of inferred and true place fields for four randomly selected hippocampal neurons. The inferred place field (top row) for cell c is a combination of location distributions for each state i weighted by the inferred firing rates $\lambda_{c,i}$, whereas the true place field (bottom row) for cell c is a histogram of locations in which cell c fires. The black dots show the rat's locations used for each histogram. The inferred place fields closely match the true place fields. With adequate spike data recording, we expect a higher latent state dimensionality to yield higher spatial resolution in the inferred place fields. (B) Summary statistics of total variation (TV) distance between the inferred and true place fields for 47 neurons. The TV distance for the four examples in panel (A) are 0.0123, 0.0067, 0.0136, and 0.0091, respectively.

Figure 9: Mutual information of the inferred states and the rat’s position. Latent state are ordered by their occupancy, i.e. the number of times the rat was in that state.

data was only used in reconstruction but not in the inference procedure. Four pairs of inferred and true place fields are shown in Fig. 8A. On the top row is the inferred place field; on the bottom is the true place field computed using the locations of the rat when cell c fired shown by black dots. We further assess the difference between the true and estimated place fields in population. Specifically, we compute the *total variation distance* between the inferred and true place fields for all 47 neurons, and the histogram statistic is shown in Fig. 8B.

In addition, we evaluate the model in terms of the information latent states convey about the rat’s position in the circular environment. To do so, we divide the environment into 121 bins of equal area and treat the rat’s position as a discrete random variable. Likewise, we treat the latent state as a discrete random variable, and we compute the discrete mutual information between these two variables. We investigate the information content of each individual state by constructing a binary random variable indicating whether or not the model is in state i and measuring its mutual information with the rat’s position. The result is shown in Fig. 9, where the latent states are ordered in decreasing order of occupancy. As expected, states that are more frequently occupied carry more information about the rat’s position.

4. Extensions and discussion

4.1. Hidden semi-Markovian models

In experimental data analysis, a striking feature of the inferred state transition matrix (Fig. 6E) is that the first 40 states exhibit strong self-transitions. This is a common feature of time series and has been addressed by a number of augmented Markovian models. In particular, hidden semi-Markovian models (HSMs) explicitly model the duration of time spent in each state separately

from the rest of the state transition matrix (Johnson and Willsky, 2013). Building this into the model allows the Dirichlet or HDP prior over state transition vectors to explain the rest of the transitions, which are often more similar. Alternatively, the “sticky” HMM and HDP-HMM accomplish a similar effect (Fox et al., 2008).

4.2. Statistical and computational considerations

We have seen a great advantage in nonparametric Bayesian formalism (i.e., HDP-HMM vs. HMM) regarding automatic model selection. This is especially important for sparse sample size or short recording in some neuroscience applications, where cross validation on data is infeasible.

For any statistical estimation, we need to consider the “bias vs. variance” problem. In VB inference, there is a potential estimate bias due to bound optimization (since we optimize the lower bound of the marginal likelihood). In addition, because of the mean-field approximation, the parameter’s variance tends to be underestimated. In MCMC inference, the estimate is asymptotically unbiased, however, if the Markov chain mixes slowly, the estimate’s variance can be inaccurate.

Computationally, the fully-Bayesian HDP-HMM inference is the most demanding. In practice, one can choose various inference tools with gradually increased computational resources (VB, empirical Bayes, Gibbs sampling or HMC) depending on the data sample size and complexity. In addition, the convergence of these algorithm may vary according to the choice of hyperparameters.

4.3. Latent state dimensionality and continuous latent state

In experimental data analysis, the number of identified states from HDP-HMM depends on the data as well as the priors of hyperparameters. Given the same size of environment, different numbers of cells or different recording duration may yield different estimation results (Fig. 3), since the nonparametric prior allocates states in accordance with the complexity of the data. We found that the weak priors over the concentration parameters have a minimal effect on the number of inferred states (Fig. 7). Fixing the scale hyperparameter of the firing rate prior distribution and performing Gibbs sampling over the scale of the prior is a simple and robust method.

In our problem, we formulate the latent state is discrete (finite or infinite) and infer the state-transition matrix, from which we can derive the “topology graph” of the unknown environment (Chen et al., 2012a, 2014). In parallel to the discrete-state HMM or HDP-HMM, we can also formulate a continuous state-space model, where the state is Gaussian and the observation is Poisson (Brown et al., 1998; Smith and Brown, 2003; Yu et al., 2009; Buesing et al., 2012). Various inference algorithms (Gaussian or variational approximation) have been developed for such models in the literature. Different from discrete state, the continuous-state has a smoother representation (due to infinite spatial resolution). However, similar to the HMM, we will need to deal with the model

selection (dimensionality of latent state) problem, which is often tackled by cross validation (Yu et al., 2009). In addition, continuous latent state is subject to sign/scale ambiguity. For the purpose of representing space and spatial topology, the discrete-state representation is more appropriate. Provided that the animal’s behavior (spatial location) is available, the continuous representation of space will be more accurate. In general, nonparametric Bayesian inference can be applied to continuous or discrete states, as well as continuous or discrete observations (Teh et al., 2006; Van Gael et al., 2008; Fox et al., 2008, 2010; Chen, 2015).

4.4. Robustness of the population firing model

A key assumption in our probabilistic model is the Poisson likelihood. Although this assumption may not be true in experimental data, our results have showed excellent performance. To further assess the robustness of HDP-HMM-Poisson model in experimental data analysis, at every temporal bin we further add additional homogeneous non-Poissonian noise to the observed population spike counts by drawing from a NB distribution (with varying levels of mean 0.25-1.0 and variance 0.5-2.0), and repeat the decoding error analysis. We have found that, as a general trend, the median decoding error gradually grows as increasing noise mean or variance; yet the decoding performance remains quite satisfactory (results not shown).

4.5. Use of soft-labeled spikes

Thus far, we have assumed that all recorded ensemble spikes are sorted and clustered into single units. Nevertheless, it is known that spike sorting is complex, time-consuming and error-prone (Wood and Black, 2008; Shalchyan and Farina, 2014). On the one hand, sorting error is inevitable when there are overlapping features (such as spike energy, amplitudes or principal components). On the other hand, traditional spike-sorting procedures often throw away considerable non-clusterable “noisy” spikes, which might contain informative tuning information. How to use these noisy spikes and maximize the information efficiency remains an open question. In other words, can we conduct the ensemble spike analysis using unsorted spikes?

Motivated from a sorting-free ensemble decoding analysis (Chen et al., 2012b; Kloosterman et al., 2014), we may use a soft-clustering method based on a Gaussian mixtures model (parameterized by an augmented vector $\xi = \{\ell_c, \mu_c, \Sigma_c\}_{c=1}^K$ that characterizes the weights, mean, and covariance parameters of the Gaussian mixtures). By clustering the spike waveform feature space, we assign each spike with a “soft” class label (about the unit identity) according to the posterior probability within the K -mixtures. In the feature space, the points close to (far away from) the c -th cluster center are associated with a probability assignment value close to (smaller than) 1 in the c -th class. Because of the soft membership of individual spikes, the spike count $y_{c,t}$ ($c = 1, \dots, K$) within a time interval can be a non-integer value. Consequently, we replace

the variable C with K to indicate that the number of neurons is unknown, and rewrite the log likelihood as follows

$$\log p(\mathbf{y}_{1:T}|\mathcal{S}, \boldsymbol{\theta}, \boldsymbol{\xi}) = \sum_{t=1}^T \sum_{c=1}^K \log p(y_{c,t}|S_t, \boldsymbol{\theta}, \boldsymbol{\xi}) \quad (22)$$

In this case, the inference procedure consists of two steps. At the first stage, d -dimensional spike waveform features are clustered using a “constrained” Gaussian mixture model (Zou and Adams, 2012), which can be either finite or infinite. In the case of infinite Gaussian mixtures, we can also resort to the nonparametric Bayesian approach (Rasmussen, 2000; Görür and Rasmussen, 2010; Wood and Black, 2008). Upon completing the inference, each spike will be given a posterior probability of being assigned to each cluster. At the second stage, we sum the soft-labeled spikes to obtain the probabilistic spike count $y_{c,t}$ for all K -clusters, and the remaining nonparametric Bayesian (MCMC or VB) inference procedure remains unchanged. A detailed investigation of this idea will be pursued in future work.

5. Conclusion

In this paper, we have explored the use of HDP-HMMs with Poisson likelihoods to analyze rat hippocampal ensemble spike data during spatial navigation. Compared to the parametric finite-state HMM, the HDP-HMM allows more flexibility to model the experimental data (without relying on time-consuming cross-validation in model selection). We evaluate two nonparametric Bayesian inference algorithms for HDP-HMM, one based on VB and the other based on MCMC. Furthermore, we consider two approaches for hyperparameter selection, an issue that is particularly important for the real-life application. It is found that the MCMC algorithm with HMC updates for the hyperparameters is robust and achieves the best performance in all simulated and experimental data. Our investigation shows a promising direction in applying nonparametric Bayesian methods for ensemble neuronal spike data analysis.

The unsupervised Bayesian inference approach allows us (or hippocampus downstream structures) to read out spatial information from hippocampal neuronal ensembles without *a priori* place receptive field information. One important future research direction is to apply this method to investigate sleep-associated hippocampal ensemble spike activity during either slow wave sleep (SWS) or rapid eye movement (REM) sleep (Louie and Wilson, 2001; Lee and Wilson, 2002). Traditionally, one would rely on place receptive fields estimated from pre-sleep run behavior to infer the content of population spike activity in sleep, but this approach is accompanied with many well-known statistical challenges, such as non-stationarity, firing rate remapping, and timescale warping. Our approach proposed here can provide an effective and complementary paradigm to investigate neural representation of hippocampal population codes without direct measurement of spatial correlate (Chen et al., 2015). The same

principle can also be applied to neocortical ensemble spike data (Ji and Wilson, 2007; Peyrache et al., 2009; Gulati et al., 2014).

Conflict of interest

The authors declare that there is no conflict of commercial interest.

Acknowledgments

Z.C. is supported by an NSF-CRCNS award (No. IIS-1307645) from the National Science Foundation. S.W.L. is supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship and by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216. M.J.J. is supported by the Harvard/MIT Joint Research Grants Program. M.A.W. is supported by the NIH Grant R01-MH06197, TR01-GM10498, ONR-MURI N00014-10-1-0936 grant. This material is also based upon work supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216. We also thank the members of the Wilson laboratory for sharing the experimental data and valuable discussions. S.W.L. and M.J.J. thank Ryan Adams for his support and helpful advice.

References

- Beal, M. J., Ghahramani, Z., Rasmussen, C. E., 2002. The infinite hidden Markov model. In: *Advances in Neural Information Processing Systems 14*. MIT Press, Cambridge, MA, pp. 577–585.
- Brown, E. N., Frank, L. M., Tang, D., c. Quirk, M., Wilson, M. A., 1998. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience* 18 (18), 7411–7425.
- Bryant, M., Sudderth, E. B., 2012. Truly nonparametric online variational inference for hierarchical dirichlet processes. In: *Advances in Neural Information Processing Systems 25*. pp. 2699–2707.
- Buesing, L., Macke, J. H., Sahani, M., 2012. Learning stable, regularised latent models of neural population dynamics. *Network: Computation in Neural Systems* 23, 24–47.
- Chen, Z., 2013. An overview of Bayesian methods for neural spike train analysis. *Computational Intelligence and Neuroscience* 2013, 251905.
- Chen, Z. (Ed.), 2015. *Advanced State Space Methods in Neural and Clinical Data*. Cambridge University Press.

- Chen, Z., Gomperts, S. N., Yamamoto, J., Wilson, M. A., 2014. Neural representation of spatial topology in the rodent hippocampus. *Neural Computation* 26 (1), 1–39.
- Chen, Z., Grosmark, A., Penagos, H., Buzsaki, G., Wilson, M. A., 2015. Statistical analysis towards sleep-associated hippocampal ensemble spike activity. Manuscript under review.
- Chen, Z., Kloosterman, F., Brown, E. N., Wilson, M. A., 2012a. Uncovering spatial topology represented by rat hippocampal population neuronal codes. *Journal of Computational Neuroscience* 33 (2), 227–255.
- Chen, Z., Kloosterman, F., Layton, S., Wilson, M. A., 2012b. Transductive neural decoding for unsorted neuronal spikes of rat hippocampus. In: *Proceedings of IEEE Engineering in Medicine and Biology Conference*. pp. 1310–1313.
- Chen, Z., Putrino, D. F., Ghosh, S., Barbieri, R., Brown, E. N., 2011. Statistical inference for assessing functional connectivity of neuronal ensembles with sparse spiking data. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 19 (2), 121–135.
- Cowles, M. K., Carlin, B. P., 1996. Markov chain Monte Carlo convergence diagnostics: a comparative review. *Journal of the American Statistical Association* 91, 883–904.
- Curto, C., Itskov, V., 2008. Cell groups reveal structure of stimulus space. *PLoS Computational Biology* 4 (10), e1000205.
- Dabaghian, Y., Brandt, V. L., Frank, L., 2014. Reconceiving the hippocampal map as a topological template. *eLife* 3, e03476.
- Dabaghian, Y., Memoli, F., Frank, L., Carlsson, G., 2012. A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Computational Biology* 8 (8), e1002581.
- Davidson, T. J., Kloosterman, F., Wilson, M. A., 2009. Hippocampal replay of extended experience. *Neuron* 63 (4), 497–507.
- Ewens, W. J., 1990. Population genetics theory—the past and the future. In: Lessard, S. (Ed.), *Mathematical and Statistical Developments of Evolutionary Theory*. Springer, pp. 177–227.
- Ferguson, T. S., 1973. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 209–230.
- Fox, E. B., Sudderth, E. B., Jordan, M. I., Willsky, A. S., 2008. An HDP-HMM for systems with state persistence. In: *Proceedings of the 25th International Conference on Machine learning*. pp. 312–319.

- Fox, E. B., Sudderth, E. B., Jordan, M. I., Willsky, A. S., 2010. Bayesian nonparametric methods for learning Markov switching processes. *IEEE Signal Processing Magazine* 27, 43–54.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., Rubin, D. B., 2013. *Bayesian Data Analysis*, 3rd Edition. CRC press.
- Goris, R. L., Movshon, J. A., Simoncelli, E. P., 2014. Partitioning neuronal variability. *Nature Neuroscience* 17, 858–865.
- Görür, D., Rasmussen, C. E., 2010. Dirichlet process gaussian mixture models: Choice of the base distribution. *Journal of Computer Science and Technology* 25 (4), 653–664.
- Gulati, T., Ramanathan, D. S., Wong, C. C., Ganguly, K., 2014. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nature Neuroscience* 17 (8), 1107–1113.
- Hoffman, M. D., Blei, D. M., Wang, C., Paisley, J., 2013. Stochastic variational inference. *Journal of Machine Learning Research* 14 (1), 1303–1347.
- Ishwaran, H., Zarepour, M., 2002. Exact and approximate sum representations for the Dirichlet process. *Canadian Journal of Statistics* 30 (2), 269–283.
- Ji, D., Wilson, M. A., 2007. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience* 10 (1), 100–107.
- Johnson, M. J., June 2014. *Bayesian time series models and scalable inference*. Ph.D. thesis, Massachusetts Institute of Technology.
- Johnson, M. J., Willsky, A. S., 2013. Bayesian nonparametric hidden semi-Markov models. *Journal of Machine Learning Research* 14 (1), 673–701.
- Johnson, M. J., Willsky, A. S., 2014. Stochastic variational inference for Bayesian time series models. *JMLR W&CP* 32, 1854–1862.
- Kloosterman, F., Layton, S. P., Chen, Z., Wilson, M. A., 2014. Bayesian decoding using unsorted spikes in the rat hippocampus. *Journal of Neurophysiology* 111 (1), 217–227.
- Kuhn, H. W., 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly* 2 (1-2), 83–97.
- Lee, A. K., Wilson, M. A., 2002. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36 (6), 1183–1194.
- Liang, P., Petrov, S., Jordan, M. I., Klein, D., 2007. The infinite PCFG using hierarchical Dirichlet processes. In: *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*. pp. 688–697.

- Louie, K., Wilson, M. A., 2001. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29, 145–156.
- Mishchenko, Y., Vogelstein, J. T., Paninski, L., 2011. A bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data. *Annals of Applied Statistics* 5, 1229–1261.
- Neal, R. M., 2010. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 113–162.
- O’Keefe, J., Nadel, L., 1978. *The Hippocampus as a Cognitive Map*. Vol. 3. Clarendon Press.
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S., Battaglia, F., 2009. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuroscience* 12 (7), 919–926.
- Raftery, A. E., Lewis, S., 1992. How many iterations in the Gibbs sampler? In: Bernardo, J. M., Berger, J., Dawid, A. P., Smith, A. F. M. (Eds.), *Bayesian Statistics*. Oxford University Press, pp. 763–773.
- Rasmussen, C. E., 2000. The infinite Gaussian mixture model. In: *Advances in Neural Information Processing Systems 12*. MIT Press, Cambridge, MA, pp. 554–560.
- Scott, S. L., 2002. Bayesian methods for hidden Markov models: Recursive computing in the 21st century. *Journal of the American Statistical Association* 97 (457), 337–351.
- Sethuraman, J., 1994. A constructive definition of Dirichlet priors. *Statistica Sinica* 4, 639–650.
- Shalchyan, V., Farina, D., 2014. A non-parametric bayesian approach for clustering and tracking non-stationarities of neural spikes. *Journal of Neuroscience Methods* 223, 85–91.
- Smith, A. C., Brown, E. N., 2003. Estimating a state-space model from point process observations. *Neural Computation* 15 (5), 965–991.
- Teh, Y. W., Jordan, M. I., 2010. Hierarchical Bayesian nonparametric models with applications. In: Hjort, N. L., Holmes, C., Müller, P., Walker, S. G. (Eds.), *Bayesian Nonparametrics*. Cambridge University Press, pp. 158–207.
- Teh, Y. W., Jordan, M. I., Beal, M. J., Blei, D. M., 2006. Hierarchical Dirichlet processes. *Journal of the American Statistical Association* 101, 1566–1581.
- Van Gael, J., Saatci, Y., Teh, Y. W., Ghahramani, Z., 2008. Beam sampling for the infinite hidden Markov model. In: *Proceedings of the 25th International Conference on Machine Learning*. pp. 1088–1095.

- Wood, F., Black, M. J., 2008. A nonparametric Bayesian alternative to spike sorting. *Journal of Neuroscience Methods* 173 (1), 1–12.
- Yau, C., Papaspiliopoulos, O., Roberts, G. O., Holmes, C., 2011. Bayesian non-parametric hidden Markov models with applications in genomics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73 (1), 37–57.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., Sahani, M., 2009. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology* 102, 614–635.
- Zou, J. Y., Adams, R. P., 2012. Priors for diversity in generative latent variable models. In: *Advances in Neural Information Processing Systems* 24. MIT Press, Cambridge, MA.