Published in final edited form as:

J Am Stat Assoc. 2016 ; 111(514): 459–471. doi:10.1080/01621459.2015.1116988.

A Dynamic Bayesian Model for Characterizing Cross-Neuronal Interactions During Decision-Making

Bo Zhou, David E. Moorman, Sam Behseta, Hernando Ombao, and Babak Shahbaba

Department of Statistics, University of California, Irvine, CA

Babak Shahbaba: babaks@uci.edu

Abstract

The goal of this paper is to develop a novel statistical model for studying cross-neuronal spike train interactions during decision making. For an individual to successfully complete the task of decision-making, a number of temporally-organized events must occur: stimuli must be detected, potential outcomes must be evaluated, behaviors must be executed or inhibited, and outcomes (such as reward or no-reward) must be experienced. Due to the complexity of this process, it is likely the case that decision-making is encoded by the temporally-precise interactions between large populations of neurons. Most existing statistical models, however, are inadequate for analyzing such a phenomenon because they provide only an aggregated measure of interactions over time. To address this considerable limitation, we propose a dynamic Bayesian model which captures the time-varying nature of neuronal activity (such as the time-varying strength of the interactions between neurons). The proposed method yielded results that reveal new insight into the dynamic nature of population coding in the prefrontal cortex during decision making. In our analysis, we note that while some neurons in the prefrontal cortex do not synchronize their firing activity until the presence of a reward, a different set of neurons synchronize their activity shortly after stimulus onset. These differentially synchronizing sub-populations of neurons suggests a continuum of population representation of the reward-seeking task. Secondly, our analyses also suggest that the degree of synchronization differs between the rewarded and non-rewarded conditions. Moreover, the proposed model is scalable to handle data on many simultaneously-recorded neurons and is applicable to analyzing other types of multivariate time series data with latent structure. Supplementary materials (including computer codes) for our paper are available online.

Correspondence to: Babak Shahbaba, babaks@uci.edu.

Supplementary materials

Supplementary materials are provided in four separate files contained in a single archive:

Computational details: This file includes the details of our computational methods along with the Markov Chain Monte Carlo (MCMC) algorithms we used for Bayesian inference.

Negative correlation: This file includes an additional illustrative example showing that our method can easily detect negatively correlated neurons.

Experimental details: This file provides the details of experimental data along with our results for model checking and diagnostics.

Computer codes: This zipped file includes all the MATLAB codes to implement our method and reproduce our results. It also contains the experimental data discussed in our paper.

Keywords

Decision-making; Spike trains; Dynamic synchrony; Gaussian processes

1 Introduction

This paper is motivated by neurophysiological studies on cognitive behaviors such as decision-making. Decision-making, and related behaviors are complex because they involve organized patterns of cognitive and behavioral subcomponents. For an individual to successfully complete a decision-making task, a number of temporally-organized events must occur: stimuli must be detected, potential outcomes must be evaluated, behaviors must be executed and inhibited, and outcomes (such as reward or no-reward) must be experienced. Here, we develop a novel statistical model for studying spike-train interactions between neurons recorded from a rat in a basic decision-making experiment. The potential clinical impact of this study is broad. Disrupted decision-making is recognized to be a common feature across many psychiatric disorders including schizophrenia and ADHD. Understanding how populations of neurons function in decision-making and related behaviors will guide targeting of specific neural systems for development of treatments.

In recent years, there have been a number of studies investigating how the activity of single neurons (see Figure 1) are related to decision-making (Gold and Shadlen 2007; Shadlen and Kiani 2013). However, it is very likely that decision-making is encoded by the temporally-precise interactions between populations of neurons (Buzsáki 2004, 2010). With this information in mind, we propose a statistical model that can capture the dynamics of the cross-neuronal interactions and can be used to test hypotheses on differences in neuronal responses between different experimental conditions. More importantly, our goal is to develop a model that is flexible and powerful and thus can be used by neuroscientists for testing existing theories as well developing new ways of thinking of how populations of neurons relate to behavior with high temporal specificity. Our proposed model characterizes the activity of many neurons in a population by connecting the joint distribution of multiple spike trains to their marginals via a parametric copula process model. In general, models that couple the joint distribution of two (or more) variables to their individual marginal distributions are called copula models. Using such models, we are able to separate the inference on cross-neuronal interactions from the individual marginal firing activity of neurons.

To estimate the underlying time-varying firing rates, we use a flexible Gaussian process model. In addition to allowing the firing rates to change over time, our proposed model can also capture time-varying interactions among neurons. Inference in our method involves a computationally efficient sampling scheme, which can be easily applied to a moderate (tens) to large (hundreds) number of simultaneously recorded neurons.

Early analysis of behavioral neurophysiological data involving multiple simultaneously-recorded neurons focused on correlation of activity across pairs of neurons (Narayanan and Laubach 2009) using a joint peristimulus time histogram (JPSTH), which shows spike intensity over time (Gerstein and Perkel 1969) to identify synchrony between a pair of

neurons. However, to capture complex behavior of neural networks, many new statistical methods have been recently proposed for analyzing several simultaneously recorded neurons in order to extract information contained in their temporal interactions under different experimental conditions. For more discussion on analysis of spike trains, refer to Harrison et al. (2013); Brillinger (1988); Brown et al. (2004); Kass et al. (2005); West (2007); Reich et al. (1998); Barbieri et al. (2001); Kass and Ventura (2001); Grün et al. (2002); Kass et al. (2005); Rigat et al. (2006); Patnaik et al. (2008); Pillow et al. (2008); Jacobs et al. (2009); Diekman et al. (2009); Sastry and Unnikrishnan (2010); Kottas et al. (2012); Kelly and Kass (2012); Shahbaba et al. (2014).

In contrast to many of these existing methods, our approach is based on a dynamic copula process model. Several stationary copula models have been previously proposed for neuroscience problems. Berkes et al. (2009) propose a variety of copula models for capturing neural dependencies and develop an efficient maximum likelihood procedure for inference. Swihart et al. (2010) proposed a unified approach to model multivariate binary data using copulas on partitions. They established a modeling framework which produced likelihood-based inference about the effects of the covariates on marginal success probabilities while accounting for the correlation among multivariate outputs.

Although they did not refer to their model as a copula, in a recent work, Kass et al. (2011); Kelly and Kass (2012) proposed a new method for detecting synchrony by modeling the joint firing probability of two neurons, denoted $\lambda^{AB}(t|H_t^{AB})$, by a factorization of the marginals

$$\lambda^{AB}(t|H_t^{AB}) = \zeta(t) \lambda^A(t|H_t^A) \lambda^B(t|H_t^B)$$

where H_t denotes the history of firing up to time t and the quantity $|\zeta(t) - 1|$ can be interpreted as the deviation of co-firing from what is predicted by independence (i.e., when $\zeta(t) = 1$ for all t). This is an important work, mainly because it offers a coherent approach for the hypothesis testing problem of the independence among multiple spike trains by utilizing generalized linear modeling. Shahbaba et al. (2014) extended this work to multiple neurons by a semi-parametric Bayesian Model. The non-parametric Gaussian process model was proposed to estimate the time-varying firing probabilities while the parametric copula model accounted for the correlation structure among the multiple neurons. The main limitation of these methods is that the cross-dependence structure among neurons is assumed to be constant in time which is unlikely to be true given the complexity of neuronal processes.

As we will illustrate in this paper, these static (stationary) models that aggregate cross-neuronal spike-train interactions over time might provide misleading results. Indeed, to address this issue, there already exist many dynamic methods for modeling brain functional and effective connectivity (Friston et al. 1997; Cribben et al. 2013; Ombao et al. 2005; Ombao and Van Bellegem 2008; Motta and Ombao 2012; Park et al. 2014; Lindquist et al. 2014). However, these approaches are primarily designed for continuous-valued signals such as functional magnetic resonance imaging (fMRI) and electroencephalogram (EEG) data.

Also, some of these methods either assume that the change points are known (Friston et al. 1997), or that there is a finite number of unknown change points between which the model remains approximately static. In contrast, our proposed model in this paper allows for smoothly-varying dynamic spike train interactions among neurons by introducing a set of time-dependent indicators that can adapt to changes in the strength of interactions. Thus, our model not only identifies a subset of functionally connected neurons, but also provides temporally precise connectivity among those neurons. Furthermore, our analysis of simultaneous recorded neurons gives an in-depth cellular-level understanding of the neuronal circuitry while fMRI and EEG studies do not. We propose a semi-parametric Bayesian dynamic model to study the joint activity of neurons in a population. The proposed model captures cross-dependence among neurons while taking into account the marginal time-varying firing probabilities of each neuron. In that sense, our method can be considered as a copula process (Wilson and Ghahramani 2012) and a generalization of the copula method by Swihart et al. (2010). Further, our approach extends the state-of-the-art methods in modeling spike trains by allowing the interaction among neurons to evolve over time. That is, our method does not assume a static (stationary) connectivity among neurons.

In this paper, we evaluate the performance of our proposed model via simulation studies and apply it to neuronal spike train recordings collected during a basic decision-making experiment. The goal of the experiment was to investigate the role of the prefrontal cortex in rats with respect to reward-seeking behaviors and inhibition of reward-seeking in the absence of a rewarded outcome. The activity of the 6 prefrontal neurons was recorded simultaneously. During the experiment, rats chose to either press or withhold presses to the presented levers. A sucrose reward was given when rats pressed one designated lever and received no reward when pressing the other. See Figure 1 for the spike train data of a neuron recorded under various conditions.

The remainder of the paper is organized as follows. In Section 2, we present the proposed dynamic Bayesian model. To set the stage, we begin by describing a stationary version of our model for multiple spike trains and evaluate its performance using simulation studies and then generalize our method to detect time-varying synchronies. In Section 3, we analyze the neuronal spike train data described above and discuss our results and their potential impact on further studies on decision making and related cognitive functions. We conclude the paper with future directions in Section 4.

2 Detecting Time-Varying Synchronies

As discussed above, our goal is to develop a model that captures the time-varying cross-neuronal spike train interactions and to identify differences in the nature of these interactions during the rewarded vs. non-rewarded conditions. To this end, we propose a non-stationary Bayesian copula process model. The details of our model are presented in Section 2.2. Before we describe our dynamic model, however, we first describe a simpler stationary version of our model.

2.1 Stationary copula model for detecting synchrony

To model the cross-dependence among multiple spike trains, we first discretize the time into small bins. Here, we set the intervals to 100 ms. It is possible to use smaller intervals (e.g., 5ms) for the stationary model discussed in this section; however, our proposed non-stationary model requires longer intervals to be effective and computationally efficient. The spike trains represented by point processes can then be transformed to binary time series which consist of 1s and 0s, where 1 indicates presence and 0 indicates absence of spikes in each time bin. Let $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{iT})$ be the transformed spike train of the i^{th} neuron, where $i = 1, \dots, n$ and observations are over time bins $t = 1, \dots, T$. In general, there are R trials for each neuron. Given the model parameters, we treat these trials as iid samples. Therefore, in what follows, we present our model for a single trial for simplicity.

For a given time bin, t , we model the joint firing probability of multiple spike trains at time t , $P_r(Y_{1t} = y_{1t}, \dots, Y_{nt} = y_{nt})$, as a function of the marginal probabilities, $P_r(Y_{it} = y_{it})$. If the neurons are independent, then the joint probability is equal to the product of the marginal probabilities. Note that the joint probability has the following simplex constraint:

$$\sum_{(y_{1t}, \dots, y_{nt}) \in (0,1)^n} P_r(Y_{1t}=y_{1t}, \dots, Y_{nt}=y_{nt}) = 1, \text{ for each } t=1, 2, \dots, T.$$

To preserve the above constraint, we use a continuous latent variable u_{it} and a threshold τ_{it} corresponding to each Y_{it} and model the observed spike trains as follows:

$$Y_{it} = (-\infty, \tau_{it}](u_{it}) = \begin{cases} 1, & \text{if } u_{it} \leq \tau_{it} \\ 0, & \text{otherwise.} \end{cases}$$

Here 1/0 indicates the presence/absence of spikes within that time bin.

It is worth mentioning that in our model τ_{it} represents an aggregate of those external conditions that lead to neuronal spiking. As such, τ_{it} may vary as a function of time, as it will be affected by the biological processes leading to action potential and the external stimuli associated with the experiment. Additionally, the underlying latent variable $u_{\{i\}t} = (u_{1t}, \dots, u_{nt})$ allows for the manifestation of the *network* at time t . The strength and structure of this network can be traced to the components of the covariance matrix Σ for these latent variables.

For any time t , we assume that $u_{\{i\}t}$ follows a multivariate Gaussian distribution with mean zero,

$$u_{\{i\}t} \sim N(0, \Sigma).$$

We use the notation $u_{\{i\}t}$ to indicate that unlike τ_{it} , which is a neuron-specific parameter (i.e., will be estimated for each neuron separately), the latent variables $u_{\{i\}t}$ are defined jointly for all neurons and will be estimated using all spike trains simultaneously.

For the moment, we assume that Σ does not vary over time so $u_{\{j\}t}$ is a stationary copula process. This will be allowed to evolve with time in the dynamic setting. Moreover, the support of the joint distribution of the latent variables is the Cartesian cross of the real lines which can be partitioned into 2^n quadrants by intersecting thresholds, τ_{it} . There are 2^n combination of outputs from n neurons; hence, there are 2^n joint probabilities. Our model guarantees the simplex constraint by mapping the probabilities in the 2^n quadrants (which sum to 1) to the 2^n joint probabilities of n neurons.

2.1.1 On the cross-dependence of the spike trains and the latent processes—

In our method, we use the correlation parameter ρ_{ij} to capture the relationship between neurons i and j . To show the validity of this approach, we examine the pairwise correlation among multiple spike trains using the observed data,

$$\begin{aligned} \text{corr}[Y_{it}, Y_{jt}] &= \frac{E[Y_{it}Y_{jt}] - E[Y_{it}]E[Y_{jt}]}{\sqrt{\text{Var}[Y_{it}]\text{Var}[Y_{jt}]}} \\ E[Y_{it}] &= P_r(u_{it} \leq \tau_{it}) = P_r\left(\frac{u_{it}}{\sqrt{\sigma_{ii}}} \leq \frac{\tau_{it}}{\sqrt{\sigma_{ii}}}\right) = \Phi\left(\frac{\tau_{it}}{\sqrt{\sigma_{ii}}}\right) \\ \text{Var}[Y_{it}] &= \Phi\left(\frac{\tau_{it}}{\sqrt{\sigma_{ii}}}\right)\left(1 - \Phi\left(\frac{\tau_{it}}{\sqrt{\sigma_{ii}}}\right)\right) \\ E[Y_{it}Y_{jt}] &= P_r(u_{it} \leq \tau_{it}, u_{jt} \leq \tau_{jt}) \\ &= \int_{-\infty}^{\tau_{it}} \int_{-\infty}^{\tau_{jt}} \frac{e^{-\frac{1}{2(1-\rho_{ij}^2)}\left(\frac{u_{it}^2}{\sigma_{ii}} - 2\frac{\rho_{ij}u_{it}u_{jt}}{\sqrt{\sigma_{ii}\sigma_{jj}}} + \frac{u_{jt}^2}{\sigma_{jj}}\right)}}{2\pi\sqrt{(1-\rho_{ij}^2)\sigma_{ii}\sigma_{jj}}} du_{it} du_{jt} \\ &\equiv F\left(\frac{\tau_{it}}{\sqrt{\sigma_{ii}}}, \frac{\tau_{jt}}{\sqrt{\sigma_{jj}}}, \rho_{ij}\right). \end{aligned}$$

Here Φ is the cumulative distribution function of standard normal distribution. Nelson

(2006) showed that given $\frac{\tau_{it}}{\sqrt{\sigma_{ii}}}$ and $\frac{\tau_{jt}}{\sqrt{\sigma_{jj}}}$, the cumulative distribution, $f\left(\frac{\tau_{it}}{\sqrt{\sigma_{ii}}}, \frac{\tau_{jt}}{\sqrt{\sigma_{jj}}}, \rho_{ij}\right)$, of bivariate normal distribution provides comprehensive concordance ordering with respect to ρ_{ij} .

$$F(., ., -1) \leq F(., ., \rho) \leq F(., ., \rho') \leq F(., ., 1)$$

where $-1 \leq \rho \leq \rho' \leq 1$. Hence, $\text{corr}[Y_{it}, Y_{jt}]$ is a non-decreasing function with respect to ρ_{ij} . In addition, when $\rho_{ij} = 0$, then $\text{corr}[Y_{it}, Y_{jt}] = 0$.

Remark: The above derivation posits that the spike trains preserve the dependence structure of the latent variables. In other words, if the latent variables are positively correlated (or negatively correlated, or independent) then so are the spike trains. Further, stronger (or weaker) correlation between the latent variables corresponds to stronger (or weaker) correlation between the spike trains. This allows us to make inference about the cross-dependence among spike trains in terms of the correlation parameters ρ_{ij} for the latent variables. Finally, it is worth emphasizing that the perceived unidentifiability of the threshold parameter does not affect our inference since the standardized threshold, $\frac{\tau_{jt}}{\sqrt{\sigma_{jj}}}$ and the correlation parameter (which is used for inference) remain identifiable. Also, note that

σ_{ji} is shared by all the time points; hence, it controls the overall firing rate. In contrast, τ_{it} is specific to the t^{th} time interval.

2.1.2 Gaussian process prior on the thresholds—Note that the joint probabilities of spike trains depend on both the latent variables, $u_{\{i\}t}$, and thresholds, τ_{it} . While the latent variables specify the dependence structure as discussed above, the thresholds determine the marginal probabilities of firing for each neuron. Specifically, the marginal firing probabilities, $P_f(Y_{it} = 1)$, depends on the thresholds τ_{it} as discussed above. The marginal probabilities usually follow nonlinear patterns over time. We can accommodate this in our model using the thresholds. That is, our model adjusts τ_{it} in a smooth way to capture the firing probability of neuron i at time t using the observed spike trains. To this end, we assume that τ_{it} is the t^{th} observation from a function, $\tau_i(t)$ with a Gaussian process prior.

A Gaussian process (GP) on the real line is a random real-valued function $x(t)$, with moments determined by its mean function $\mathbb{E}x(s)$ and kernel $\kappa(s, t) = \text{Cov}(x(s), x(t))$. More precisely, all finite-dimensional distributions $(x(t_1), \dots, x(t_m))$ are multivariate Gaussian with mean $(\mathbb{E}x(t_1), \dots, \mathbb{E}x(t_m))$, and with covariance matrix $(\kappa(t_k, t_\ell))_{k, \ell=1}^m$. Since the latter must be positive semi-definite for every finite collection of inputs $t_1, \dots, t_m$, only certain kernels κ are *valid* (i.e., define Gaussian processes). Thus when using Gaussian processes, a practitioner often chooses from among the few popular classes of kernels, such as the squared exponential (SE), Ornstein-Uhlenbeck (OU), Matérn, polynomial, and linear combinations of these. For example, we can use the following covariance form, which combines a random constant with the SE kernel and iid observation noise (Rasmussen and Williams 2006; Neal 1998):

$$\begin{aligned} C|_{s,t} &= \text{Cov}[x(s), x(t)] \\ &= \lambda^2 + \theta^2 \exp[-\gamma^2(t-s)^2] + \delta_{st}\nu^2 \end{aligned} \quad (1)$$

Here, λ , θ , γ and ν are hyperparameters with their own hyperpriors. In general, the choice of kernel encodes our *qualitative* beliefs about the underlying signal. For instance, samples from a GP with OU kernel are always non-differentiable functions $x(t)$, and the SE kernel generates only infinitely differentiable functions.

Applications of GP models have been explored extensively in spatial modeling (Sampson and Guttorp 1992; Schmidt and O'Hagan 2003; Cressie 1993; Stein 1999; Gelfand et al. 2005; Banerjee et al. 2008b; Duan et al. 2007; Quick et al. 2013). Here, we use GP to develop a flexible model for thresholds as a nonlinear function of time. To this end, for each neuron i , we assume the following Gaussian process prior:

$$\tau_i(t) \sim GP(0, C_i).$$

where C_i has the above SE kernel form (1). Note that a finite-dimensional $(\tau_{i1}, \dots, \tau_{iT})$ has a multivariate normal prior distribution with mean zero and covariance $C_{i|t,s} = \text{Cov}(\tau_i(s), \tau_i(t))$.

Alternative kernels for faster computation: In general, as the number of time bins, T , increases, computing the likelihood of Gaussian process models becomes time-consuming as it involves inversion of large-scale covariance matrices with $\mathcal{O}(T^3)$ complexity. To mitigate this issue, one could use the Wiener process (a.k.a. the Brownian motion) prior, where the covariance function has a simpler form,

$$C_i|_{s,t} = \text{Cov}(\tau_i(s), \tau_i(t)) = \theta_i \min(s, t).$$

This way, we can reduce the computational complexity of the GP model from $\mathcal{O}(T^3)$ for the SE kernel (1) to $\mathcal{O}(T)$. The resulting process, however, is not stationary. Alternatively, one could use the Ornstein-Uhlenbeck (OU) process prior instead. The OU process, $\tau_i(t)$, can be defined in terms of the following stochastic differential equation:

$$d\tau_i(t) = -\theta_i(\tau_i(t) - \mu_i)dt + \gamma_i dW_i(t),$$

where $W_i(t)$ is the Wiener process, μ_i is the long-term mean of the process, γ_i controls the degree of volatility, and θ_i is the rate by which the process reverts to its long-term mean. Similar to the Wiener process, the computational complexity of OU increases linearly with T , but unlike the Wiener process, the OU process is stationary: given $\tau_{i0} \sim N(\mu_i, \gamma_i^2 / 2\theta_i)$, we have $\tau_{it} \sim N(\mu_i, \gamma_i^2 / 2\theta_i)$ for all t .

2.1.3 Summary on the stationary model—The left panel of Figure 2 shows the schematic representation of our proposed stationary model, which uses a set of latent variables and thresholds to model the joint probability of multiple spike trains. (See Section S.1.1 in Supplementary Materials for more details.) The latent variables are modeled by a stationary copula process, and the thresholds are modeled using a Gaussian process model with the SE kernel (alternatively, OU could be used for faster computation). The copula process model connects the marginal time-varying firing probabilities to the joint probabilities of multiple spike trains. The thresholds control the time-varying marginal firing probabilities, while the latent variables determine the cross-dependence among neurons. Overall, our model involves four types of parameters: 1) thresholds, τ_{it} , 2) latent variables, $u_{\{i\}}$, 3) copula parameters, Σ , and 4) hyperparameters of the Gaussian process model (i.e., the parameters that define C).

In our numerical experiments, we used weakly informative priors for all hyperparameters and used Markov Chain Monte Carlo (MCMC) to simulate samples from posterior distributions. More details are provided in Section S.1.1 in Supplementary Materials.

2.1.4 An illustration of the stationary model—To illustrate the stationary model, we use a simple numerical experiment. First, we generate the spike trains of four neurons, where the first pair of neurons are positively correlated, the second pair of neurons are negatively correlated, and that the two pairs are independent of each other. To this end, we generate the time-varying marginal firing probabilities for each neuron and then calculate the joint probabilities of the two neurons by multiplying the product of marginal firing

probabilities by an extra term, ζ , so that $\zeta = 1$ indicates independence while $\zeta < 1$ and $\zeta > 1$ indicate negative correlation and positive correlation between the two neurons respectively. Note that we are using a data generating mechanism that is different from our model in order to test for the robustness of our approach.

In this simulation study, we set $\zeta = 1.3$ for the first pair of neurons and $\zeta = 0.7$ for the second pair of neurons. Table 1 shows the posterior estimates of the correlation parameters ρ_{ij} 's with their corresponding 95% posterior intervals. Figure 3 shows the peri-stimulus time histogram (PSTH) of these neurons with their estimated number of spikes according to our model. The results presented in Table 1 show that our model correctly detects the dependence among the two pairs of neurons.

2.2 Non-stationary copula model for detecting synchrony

As discussed above, our main objective is to develop a flexible non-stationary method that allows for interactions among neuronal spike trains to vary over time. To this end, we use a time-varying covariance matrix for the latent variables. The time-dependent covariance matrix implies a dynamic dependence structure among spike trains. Without imposing a structure, however, the covariance matrix would become too complex, especially when n is large. We impose such a structure by introducing a series of time-dependent binary indicators in the stationary copula process model discussed in the previous section,

$$u_{\{i\}t} \sim N(0, \sum_t^*), \quad \text{where } \sum_t^* = \begin{bmatrix} \sigma_{11} & \sigma_{12} I_1^t I_2^t & \cdots & \sigma_{1n} I_1^t I_n^t \\ \sigma_{12} I_2^t I_1^t & \sigma_{22} & \cdots & \sigma_{2n} I_2^t I_n^t \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{1n} I_n^t I_1^t & \sigma_{2n} I_n^t I_2^t & \cdots & \sigma_{nn} \end{bmatrix}$$

2.2.1 On the indicator I_i^t —Here, $I_i^t = 0$ indicates that the i -th neuron is not correlated with any other neurons at time t . These time-dependent indicators identify neurons that are involved in the network and impose a structure on the covariance matrix. We assume $I_i^t \sim \text{Bernoulli}(p_i(t))$. To ensure $0 < p_i(t) < 1$, we use the probit link function such that $p_i(t) = \Phi(q_i(t))$. A Brownian motion prior is assumed for $q_i(t)$,

$$q_i(t) \sim GP(\eta, K_i),$$

Note that over a finite set of time points, $(q_{i1}, \dots, q_{iT})$ has a multivariate normal prior distribution with mean η and covariance $K_{i,s,t} = \text{Cov}(q_i(s), q_i(t)) = \theta_i \min(s, t)$. In this model, the mean of the Brownian motion is not zero. In general, we could treat the mean, η , as a hyperparameter. For simplicity, however, we fix it to small number (here, $\eta = -1$) in order to encourage sparsity in the network.

By introducing time-dependent binary indicators in to the copula process, our model is capable of identifying time-varying interactions among multiple neurons. Note that in this setting, the correlation between two neurons becomes $\rho_{ij}^t = \rho_{ij} I_i^t I_j^t$. The two neurons are

correlated if both I_i^t and I_j^t are non-zero; ρ_{ij} determines the overall strength of the correlation.

Remark: The time-varying correlation matrix $\sum_t^*$ is positive definite as long as Σ is positive definite.

2.2.2 Biological justification—Many neurons might be involved in a complex process such as decision-making over time. The scientific community has begun to regard the involvement of neurons in the underlying network of complex processes as a dynamic phenomenon, changing with time (Buzsáki 2010). Cortical neurons, particularly neurons in prefrontal cortex, have been characterized as exhibiting mixed selectivity, encoding multiple, high-dimensional representations, i.e., exhibiting changes in firing rate that scale with multiple parameters in experimental tasks (Rigotti et al. 2013). The optimal, and likely way by which mixed-selectivity neurons encode task variables at the level of the ensemble is for them to dynamically participate in different ensembles depending on the context and specific behaviors being performed, with single neurons potentially participating in multiple ensembles depending on the objectives of the experimental task (Barak et al. 2013; Durstewitz et al. 2010; Hyman et al. 2012; Lapish et al. 2008; Sakurai et al. 2013). In our model, Σ captures the overall network: a collection of neurons involved in a dynamic process. The indicator variables on the other hand capture the dynamic membership of these neurons in the network.

2.2.3 Markov network vs. covariance graph—As discussed by Bien and Tibshirani (2011), covariance estimation is one of the most fundamental problems in statistics. In general, this is a difficult problem since the number of parameters grows quadratically as the number of variables (here, neurons) increases. To solve this problem, many researchers have followed Dempster (1972) and developed methods to reduce the effective number of parameters by imposing sparsity on the *inverse* covariance matrix (see for example, Meinshausen and Bühlmann 2006; Yuan and Lin 2007; Banerjee et al. 2008a; Friedman et al. 2008). By setting an element of inverse covariance to zero, we create *conditional* independence between its corresponding variables. The resulting model, known as *Markov network*, can be presented as a graph with variables as nodes. An edge between two nodes represent conditional dependency between the corresponding variables.

Noticeably, many authors have proposed methods based on sparsity in the covariance matrix itself. In this case, zeros in the covariance matrix correspond to *marginal* independence between variables, and the resulting model is called *covariance graph* (Bien and Tibshirani 2011; Drton and Richardson 2008; Rothman et al. 2009). While, a Markov network is a suitable model to infer graphical *structure*, a covariance graph is typically used for *model selection* (Bien and Tibshirani 2011). The latter is a more appropriate modeling choice in our case since we intend to identify neurons that are involved in the decision-making process. That is, we intend to identify how neurons cluster together (a cluster consists of neurons that are directly or indirectly dependent) during a decision-making process. More specifically, because the imposed sparsity can change over time, our method can be regarded as a *dynamic model selection* approach.

2.2.4 Connection to other methods—As mentioned in the introduction, our approach is related to a number of existing methods such as the GLM-based method of Kass et al. (2011) and Kelly and Kass (2012), the semi-parametric model of Shahbaba et al. (2014), and copula-based models of Wilson and Ghahramani (2012) and Swihart et al. (2010). Our method is also related to the model proposed by Cunningham et al. (2007), who assume that the underlying non-negative firing rate for spike train is a draw from a Gaussian process. However, unlike the method proposed in this paper, they assume that the observed spike train is a conditionally inhomogeneous gamma-interval process given the underlying firing rate.

Since Cox (1981) several state-space models for binary time series data have been developed which are closely related to our proposed model. See, for example, Fahrmeir (1992); Carlin and Polson (1992); Song (2000); Czado and Song (2008). A binary state-space model is characterized as follows. Let $\{Y_t\}$ be a binary time series with corresponding sequence of probabilities $\{\mu_t\}$ that is defined via the state variable X_t as follows:

$$\mu_t = h(\mathbf{M}_t' X_t) \quad \text{where } X_t = \mathbf{Q}_t X_{t-1} + \varepsilon_t$$

where h^{-1} is a link function as specified for generalized linear models (GLM, see McCullagh and Nelder (1989)); $\mathbf{M}_t$ is a vector that contains time-varying covariates; $\mathbf{Q}_t$ is the transition matrix; and ε_t is a random variable (or vector) with zero mean.

The state-space model above contains both deterministic and random components which has been developed for modeling longitudinal count data in Zeger (1988); Chan and Ledolter (1995); Fahrmeir and Lang (2001a,b). Our proposed model currently does not include a deterministic component in the probability sequence, but its framework is general and can broadly include a deterministic component. The key distinction between our proposed model and the general state-space models is that our latent equation does not follow the AR(1) structure of a state equation. Our state-equation is broadly defined as a Gaussian process (GP) on the thresholds. In our current implementation, we particularly use a GP model.

Finally, we also point the connection between our model and the so-called threshold model in Albert and Chib (1993), where $Y_t = 1 \Leftrightarrow Z_t < 0$. In our model, the threshold is allowed to vary with time (rather than a fixed 0) and that the threshold process follows a GP process rather than the usual AR(1).

2.2.5 Illustrative example—In this section, we use an example to illustrate our non-stationary copula process method. Here, we generalize ζ_t from the previous example so the dependencies can change over time. As before, we use weakly informative priors for all hyperparameters and used Markov Chain Monte Carlo (MCMC) to simulate samples from posterior distributions. More details are provided in Section S.1.1 in Supplementary Materials.

Our example is designed to highlight the problems with stationary models (and thus emphasize the need for non-stationary approaches such as the one being proposed). More specifically, we demonstrate that our model can identify time-varying dependencies while

methods based on static models such as the methods of Kass et al. (2011) and Shahbaba et al. (2014) provide an incomplete picture of neuronal activity and, moreover, could give misleading results. To compare our proposed approach to those in Kass et al. (2011) and Shahbaba et al. (2014), we first generate spike trains for a pair of neurons for which ζ_t follows the pattern shown in the left panel of Figure 4. Due to the assumption that the correlation structure is stationary, the method of Kass et al. (2011) reports a p-value of 0.555 for the test of $H_0 : \zeta = 1$ (independence) vs. $H_a : \zeta < 1$. The method of Shahbaba et al. (2014) reports a 95% credible interval of [0.875, 1.144] for ζ which covers 1. Both methods incorrectly conclude that the neuronal spike trains to be independent. Neither method is capable of detecting the synchrony between the two neurons when this synchrony is brief relative to the period of non-synchrony. We also applied the static version of our proposed model (described in Section 2.1) to the simulated data. The 95% credible interval for the correlation parameter is [-0.039, 0.066] which covers 0 and thus leading to no significant evidence for dependence between the two neurons. In contrast to these stationary approaches, the proposed dynamic model captures the time-varying pattern of synchrony between the two neurons as shown in Figure 4 (Right). This illustration serves as a reminder that when the true process follows a dependence structure that is time-varying (which is the more likely scenario compared to the simplistic assumption of stationarity), then it would be appropriate to fit a non-stationary model. In Section S.2.1 (Supplementary Materials), we provide another illustrative example to show that our method can easily detect negatively correlated neurons.

2.2.6 Simulation studies—To further illustrate the advantage of our method, we conducted two simulation studies to compare our dynamic model with the method in Kass et al. (2011) which is the state-of-the-art approach for studying cross-neuronal interactions. For the first simulation study, we compared the two methods in terms of their ability to correctly detect the correlation between two neurons. To this end, we simulated data according to the approach we discussed above. Namely, the two neurons are independent for the first 80% of time ($\zeta_t = 1$ for $t \leq 0.8$) and dependent for the last 20% of time ($\zeta_t < 1$ for $t > 0.8$). We conducted simulations for 5 scenarios where we set the extra term ζ_t for the last 20% of time to 1.0, 1.2, 1.4, 1.6, and 1.8 respectively. Corresponding to each pair (which we call “Pair1”) under each of the scenarios, we also simulate an independent pair (which we call “Pair2”) where $\zeta_t = 1$ over the entire time. We then apply both models to each pair of neurons and make inference about their dependence. For scenarios 2–4, the models are expected to identify Pair1 as dependent, while identifying Pair2 as independent. The first scenario was used as the control where the models should have identified both pairs as independent. We repeated each scenario 100 times and use ROC curves to assess each model’s performance. Note that for the first scenario, we expected the ROC curve to be diagonal. For the remaining scenarios, the higher the area under ROC curve, the better the model. Figure 5 shows the ROC curves, which illustrate the performance of the two models (our proposed in color red; Kass et al. (2011) in green) in detecting dependence between neurons in Pair1. The results show that as the strength of correlation increases (ζ_t increases from 1.0 to 1.8), both methods exhibit an increase in power. Moreover, the results demonstrate that our proposed dynamic model always outperforms the method of Kass et al. (2011).

For the second simulation study, we compare the predictive power of the models. We simulate the datasets according to the five scenarios discussed above. This time, however, we randomly select two-thirds of the data to train both models (i.e., estimate parameters), and then test the models on the remaining one-third of the data by using the firing status of one neuron to estimate the firing probability for the second neuron. We derive the estimated firing probability, $\hat{p}_{2k}$ for the k^{th} observation in the test set for the second neuron using parameters estimated using the training data. Models are evaluated based on their log predictive probability (LPP) using the estimated firing probability, $\hat{p}_{2k}$, as follows:

$$LPP = \sum_{k \in \text{test set}} y_{2k} \log(\hat{p}_{2k}) + (1 - y_{2k}) \log(1 - \hat{p}_{2k}).$$

The model with a larger LPP value is considered to have greater predictive power. As before, we repeat each scenario 100 times. Figure 6 shows the 95% intervals of LPP for each scenario using our proposed model (shown as red bars) and the method of Kass et al. (2011) (shown as green bars). Again, the results show that both models display an increasing predictive power as the strength of correlation between two neurons increases. Moreover, consistent with the previous results, our proposed dynamic model outperforms the method in Kass et al. (2011) in general.

3 Analysis of the spike-train data

The ultimate goal of this research is the application to neurophysiological data collected during performance of cognitive behaviors such as decision-making. Disrupted decision-making is a common feature across many psychiatric disorders including, but not limited to, ADHD, bipolar, and compulsive behavior disorders, mood and anxiety disorders, and schizophrenia (Sharp et al. 2012). Decision-making, therefore is a potent behavioral endophenotype that can be used to study disruptions in the neural framework underlying mental disorders (Insel and Cuthbert 2009). To study decision making in animal models we conducted an experiment which we describe below. We then give an overview of the statistical approaches for analyzing neuronal spike train.

3.1 Description of the experiment

In this experiment, we recorded the activity of multiple neurons in the prefrontal cortex of rats while they were presented with stimuli that either predicted or did not predict the availability of a reward. During the recording/test sessions, two different stimuli were presented: tone 1 (10 KHz) or tone 2 (5 KHz) individually and in pseudorandom order. At the same time, one of two levers was presented an active-lever, paired with tone 1 (Rewarded-Stimulus - RS) and an inactive-lever paired with tone 2 (Non-rewarded Stimulus - NS). Pressing the active lever resulted in the offset of tone 1, retraction of the lever, and illumination of the reward receptacle. If the rat then went to the reward receptacle, 0.1 ml of 15% sucrose solution was delivered as a reward. Pressing the inactive lever produced no effect. Tones were played for a maximum of 3 sec, and levers remained extended for a maximum of 10 sec. See Moorman and Aston-Jones (2015) for more details.

We recorded the activity of multiple single neurons simultaneously by sampling from multiple channels of an implanted 16-wire electrode array. We recorded both action potentials (i.e., spikes—the firing activity of single neurons) and local field potentials (averaged population activity). In this report we focus on action potential data. Here, we focused our analysis on a recording session from one rat. We identified six neurons that exhibited consistent firing across the recording session (note that other neurons recorded in this session exhibited low or inconsistent firing across the session). More details about the experiment are provided in Section S.3.1 (Supplementary Materials).

3.2 Results and Interpretation

Throughout this section, we discretize time using 100 ms intervals. As a preliminary step, we first applied the static model to the spike train data of two neurons. The results show that under the rewarded stimulus the 95% credible interval for the correlation parameter was $[0.035, 0.147]$ compared to $[-0.039, 0.049]$ under the non-rewarded stimulus. This suggests that there was some sort of aggregated synchrony between the two neurons under rewarded stimulus but no evidence of aggregated synchrony under the non-rewarded stimulus. The proposed model suggests an interesting result: that the temporal relationship activity between the two neurons differed based on whether the stimulus predicted a rewarded or non-rewarded outcome.

As noted, the static analysis above has limitations because it does not portray a complete picture of the time-varying cross-dependence between neurons. As a next step, we applied the dynamic model to the same pair of neurons. Figure 7 shows the estimated correlation parameters based on our proposed dynamic model. The results suggest that under the rewarded stimulus, the two neurons started to co-fire after 4 seconds. In contrast, for the unrewarded stimulus, the two neurons remained uncorrelated throughout the experiment. As noted above, the strong cross-dependence between these two neurons was seen during the reward-receipt epoch in rewarded trials. This comparison allowed us to reliably say that the onset of correlation in rewarded trials was not simply a by-product of stimulus presentation (as stimuli are presented in both contexts) but has some relationship between the outcomes or behaviors. Thus, these two neurons (Figure 7) are likely members of a population encoding reward or a behavior related to reward-seeking.

There are a number of reasons that explain why neuronal populations may exhibit synchrony. In the periphery, sensory stimuli can simultaneously activate neurons, producing sensory-evoked synchrony. In some cases, synchrony arises from correlated input via, for example, neurons or neural populations that innervate multiple target populations, driving these multiple targets at the same time. For example, in the prefrontal cortex, where neurons studied here were recorded, synchrony is thought to result from strong, correlated, patterned input from the hippocampus and possibly other brain areas (Benchenane et al. 2010). In other cases, notably in cortex, interactions between inhibitory GABA interneurons and pyramidal neurons are responsible for generating synchrony across populations (Fujisawa et al. 2008). Correlated activity can also be seen among populations of neurons when activity increases and is maintained, as has been seen, for example, in synchronous activity in the

prefrontal cortex underlying working memory (Sakurai et al. 2013). Results for model checking and diagnostics are provided in Section S.3.2 (Supplementary Materials).

Next, we examined the cross-dependence between all of the 6 neurons in the data. Section S.3.3 (Supplementary Materials) provides the plots for the time-varying correlation parameters among these neurons under different scenarios. For simplicity, in Figures 8 and 9, we divided the time series into two time intervals: Period 1 covers the first 5 seconds following the stimulation onset (i.e., [0,5] seconds) while Period 2 covers the last 5 seconds after stimulation onset (i.e., (5,10] seconds). Within Period 1 of the rewarded scenario, neurons 2, 5, and 6 seemed to be weakly correlated, while others were not at all implicated in the network. During Period 2, neurons 1, 2, 3, and 4 became strongly correlated (Figure 8). Under the non-rewarded stimulus, neuron 1, 3, and 4 are correlated throughout both epochs. Towards the end of Period 2 (approximately 7–10 seconds), neurons 5 and 6 join the network with moderate correlations (Figure 9). The results suggest that neurons 1, 3, and 4 are involved in the network under both scenarios; whereas, neuron 2 plays a differential role under the two scenarios: it is always involved in the correlation network under rewarded stimulus, while it remains isolated from the network under non-rewarded stimulus. In addition, neurons 5 and 6 vary under different stimuli: they are strongly involved in the network at the beginning following presentation of the rewarded stimulus and weakly involved in the network at the end of session after presentation of the non-rewarded stimulus.

The results described here reveal new information about network encoding of different aspects of the behavioral task, in particular providing us two powerful and intersecting insights into population coding of behavior. First, our model reliably measures correlation among groups of neurons over prolonged periods of time. This is an important advance in understanding how events and behaviors can influence synchronous activity in neural populations. Intriguingly, neurons 2 and 5 exhibit synchronous activity in what appears to be the time between lever-press and reward acquisition. Thus by differentially synchronizing different populations of neurons over time, we see a continuum of population representation of the reward-seeking task. Importantly, the degree of synchronization is different in the non-rewarded condition. In this case, neurons 1, 3, and 4 exhibit synchronous activity rapidly after the onset of the non-reward-predicting tone. This synchronization may serve as a response-inhibition signal that allows the animal to withhold responding to the non-rewarded lever.

Brain regions such as the ventral medial prefrontal cortex have frequently been characterized as playing a role in response inhibition: inactivation of these areas increases behavior and stimulation of them suppresses it (Peters et al. 2009). The neural signaling underlying this type of executive control has remained elusive however, although some neurons have been shown to be activated during behavioral suppression (Milad and Quirk 2002). However, it is unlikely that the small numbers of neurons showing increased activation in these circumstances represents the mechanism by which populations of neurons encode behavior. Rather, population signaling, reflected in correlated firing, is likely to underscore neural coding of these relatively sophisticated behaviors.

4 Conclusions and future directions

We have developed a Bayesian dynamic model for analyzing cross-dependence between neurons in a population. The proposed model has a number of advantages: (1.) It captures the time-varying dependencies between neurons and thus is applicable for analyzing spike train data recorded while a subject is performing a complex cognitive task; (2.) The proposed model is useful for testing differences in cross-neuronal correlations in activity between different experimental conditions; (3.) It can be applied to a moderate to large number of neurons, and the computational complexity increases with the number of interacting neurons only; (4.) The proposed model can be applied to different data from a broad spectrum including continuous-valued time series that have some latent structure; (5.) Finally, using the proposed model, our analysis yielded results that revealed new insight into the dynamic nature of population coding in the prefrontal cortex during decision making.

Our model of course can be generalized in several ways. Although it is already a step forward towards dynamic modeling of neural networks, our current model assumes that if two neurons are related, the relationship is either positive or negative so that the sign remains constant but may vary in strength over time. However, this model can be extended for situations where the direction of the relationship changes (from positive to negative or vice-versa) over time. For example, the model can be generalized by allowing the indicator variables to be between -1 and 1 . This was not necessary in the current situation because we did not observe any evidence of such behavior in our data.

The model produced results that led us to pose new hypotheses related to the prefrontal control of executive function (response execution/inhibition). The large hypothesis resulting from these data is that complex behaviors such as decision-making come about by ensemble coding of discrete behavioral components (stimulus representation, response execution, reward consumption, etc.). These behaviors are encoded by temporally discrete synchronization of specific populations of neurons, although these populations share members. There are a number of key experiments that must be performed to verify this hypothesis and to refine our understanding of encoding along these lines. It is critical to isolate specific subcomponents of behavior that may be encoded. So, for example, studying neural encoding specifically during response execution or specifically during response inhibition (as opposed to in more complex tasks) will allow a precise association between synchronous populations and behavior. In addition, we need to test whether these ensembles reliably form—with the same member neurons—each time a behavior is performed or whether they exhibit subtle differences on a trial-to-trial basis.

The conclusions from this analysis are primary and will need to be studied from further experiments on decision-making in our laboratory and different contexts. Nevertheless, the results of the analysis of even this limited data set are already highly intriguing and already begin to reveal new insight into the dynamic nature of population coding in the prefrontal cortex during basic cognitive tasks. In particular, using these models to characterize how populations of neurons synchronize at precise times across behavior will let investigators focus precisely on the relationship between population coding and discrete components of behavior. For example, it is overly-general (i.e., lacking in temporal precision) to state that

neural populations exhibit synchronous activity during reward-seeking. Instead, we can specify that particular neural populations are correlated during specific subcomponents of these behaviors (stimulus presentation, lever approach, reward-well entry, reward consumption, etc.). As demonstrated in our analysis, by narrowing down our window of focus with respect to population encoding, we can achieve a previously unconsidered combination of spatial breadth and temporal precision. Typically the focus on neural coding of discrete behavioral intervals is limited to analysis of single-neuron activity.

Our proposed method in this paper can be utilized for current behavioral neuroscience studies, which typically involve 10s of simultaneously recorded neurons. (See the review by Buzsáki et al. (2015) on the current state of simultaneous ensemble recording.) For such studies, investigators need adequate analytical tools for understanding the temporal relationship among a relatively large number of neurons. Our method provides a pragmatic approach for responding to that need. However, driven by advances in technology, the number of neurons capable of being recorded simultaneously is increasing. As a result, the computational demand of neuroscience studies is going to increase in future. To address the computational challenges of such studies and to maintain our method as a practical approach, we need to continue working on the scalability and computational efficiency of our proposed model – as truly large-scale recordings (e.g., thousands of neurons) will soon become commonplace in neuroscience.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We would like to thank the Associate Editor and anonymous referees for their insightful suggestions and constructive comments that helped us to improve our paper. Babak Shahbaba was supported by the NIH grant R01 AI107034. David Moorman was supported by the NIH grant R21-DA032005. Hernando Ombao would like to acknowledge support from NSF Division of Mathematical Sciences and NSF Division of Social and Economic Sciences.

References

- Albert JH, Chib S. Bayesian analysis of binary and polychotomous response data. *JASA*. 1993; 88:669–679.
- Banerjee O, Ghaoui LE, d'Aspremont A. Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data. *Journal of Machine Learning Research*. 2008a; 9:485–516.
- Banerjee S, Gelfand AE, Finley AO, Sang H. Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society Series B, Statistical methodology*. 2008b; 70:825–848. [PubMed: 19750209]
- Barak O, Rigotti M, Fusi S. The Sparseness of Mixed Selectivity Neurons Controls the Generalization Discrimination Trade-Off. *The Journal of Neuroscience*. 2013; 33:3844–3856. [PubMed: 23447596]
- Barbieri R, Quirk MC, Frank LM, Wilson MA, Brown EN. Construction and analysis of non-Poisson stimulus-response models of neural spiking activity. *Journal of Neuroscience Methods*. 2001; 105:25–37. [PubMed: 11166363]
- Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, Battaglia FP, Wiener SI. Coherent Theta Oscillations and Reorganization of Spike Timing in the Hippocampal-Prefrontal Network upon Learning. *Neuron*. 2010; 66:921–936. [PubMed: 20620877]

- Berkes, P.; Wood, F.; Pillow, J. Characterizing neural dependencies with copula models. In: Koller, D.; Schuurmans, D.; Bengio, Y.; Bottou, L., editors. *Advances in Neural Information Processing Systems*. Vol. 21. 2009. p. 129–136.
- Bien J, Tibshirani RJ. Sparse estimation of a covariance matrix. *Biometrika*. 2011; 98:807–820. [PubMed: 23049130]
- Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*. 1988; 59:189–200. [PubMed: 3179344]
- Brown EN, Kass RE, Mitra PP. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neuroscience*. 2004; 7:456–461. [PubMed: 15114358]
- Buzsáki G. Large-scale recording of neuronal ensembles. *Nature Neuroscience*. 2004; 7:446–51. [PubMed: 15114356]
- Buzsáki G. Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron*. 2010; 68:362–385. [PubMed: 21040841]
- Buzsáki G, Stark E, Berényi A, Khodagholy D, Kipke DR, Yoon E, Wise KD. Tools for Probing Local Circuits: High-Density Silicon Probes Combined with Optogenetics. *Neuron*. 2015; 86:92–105. [PubMed: 25856489]
- Carlin B, Polson NG. Monte Carlo Bayesian Methods for Discrete Regression Models and Categorical Time Series. *Bayesian Statistics*. 1992; 4:577–586.
- Chan KS, Ledolter J. Monte Carlo EM Estimation for Time Series Models Involving Counts. *Journal of the American Statistical Association*. 1995; 90:242–252.
- Cox DR. Statistical analysis of time series, recent developments. *Scandinavian Journal of Statistics*. 1981; 8:93–115.
- Cressie, N. *Statistics for Spatial Data*. New York: Wiley; 1993.
- Cribben I, Wager T, Lindquist M. Detecting functional connectivity change points for single-subject fMRI data. *Frontiers in Computational Neuroscience*. 2013; 7
- Cunningham, JP.; Yu, BM.; Shenoy, KV.; Sahani, M. Inferring Neural Firing Rates from Spike Trains Using Gaussian Processes. In: Platt, JC.; Koller, D.; Singer, Y.; Roweis, ST., editors. *NIPS*. 2007.
- Czado C, Song P. State space mixed models for longitudinal observations with binary and binomial responses. *Statistical Papers*. 2008; 49:691–714.
- Dempster AP. Covariance selection. *Biometrika*. 1972; 28:157–175.
- Diekman CO, Sastry PS, Unnikrishnan KP. Statistical significance of sequential firing patterns in multi-neuronal spike trains. *Journal of neuroscience methods*. 2009; 182:279–284. [PubMed: 19559053]
- Drton M, Richardson TS. Graphical Methods for Efficient Likelihood Inference in Gaussian Covariance Models. *Journal of Machine Learning Research*. 2008; 9:893–914.
- Duan JA, Guindani M, Gelfand AE. Generalized Spatial Dirichlet Process Models. *Biometrika*. 2007:809–825.
- Durstewitz D, Vitoz NM, Floresco SB, Seamans JK. Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning. *Neuron*. 2010; 66:438–448. [PubMed: 20471356]
- Fahrmeir L. Posterior mode estimation by extended Kalman filtering for multivariate dynamic generalized linear models. *Journal of the American Statistical Association*. 1992:501–509.
- Fahrmeir L, Lang S. Bayesian Inference for Generalized Additive Mixed Models Based on Markov Random Field Priors. *Applied Statistics*. 2001a; 50:201–220.
- Fahrmeir L, Lang S. Bayesian Semiparametric Regression Analysis of Multicategorical Time-Space Data. *Annals of the Institute of Statistical Mathematics*. 2001b; 53:11–30.
- Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*. 2008; 9:432–441. [PubMed: 18079126]
- Friston KJ, Buechel C, Fink GR, Morris J, Rolls E, Dolan RJ. Psychophysiological and Modulatory Interactions in Neuroimaging. *NeuroImage*. 1997; 6:218–229. [PubMed: 9344826]
- Fujisawa S, Amarasingham A, Harrison MTT, Buzsáki G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*. 2008; 11:823–833. [PubMed: 18516033]

- Gelfand AE, Kottas A, MacEachern SN. Bayesian Nonparametric Spatial Modeling with Dirichlet Process Mixing. *Journal of the American Statistical Association*. 2005; 100:1021–1035.
- Gerstein GL, Perkel DH. Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation. *Science*. 1969; 164:828–830. [PubMed: 5767782]
- Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci*. 2007;30.
- Grün S, Diesmann M, Aertsen A. Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Computation*. 2002; 14:43–80. [PubMed: 11747534]
- Harrison, MT.; Amarasingham, A.; Kass, RE. Statistical identification of synchronous spiking. In: Lorenzo, PD.; Victor, J., editors. *Spike Timing: Mechanisms and Function*. Taylor and Francis; 2013.
- Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences*. 2012; 109:5086–5091.
- Insel T, Cuthbert B. Endophenotypes: bridging genomic complexity and disorder heterogeneity. *Biological psychiatry*. 2009; 66:988–989. [PubMed: 19900610]
- Jacobs AL, Fridman G, Douglas RM, Alam NM, Latham PE, Prusky GT, Nirenberg S. Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences of the United States of America*. 2009; 106:5936–5941. [PubMed: 19297621]
- Kass RE, Kelly RC, Loh W-L. Assessment of synchrony in multiple neural spike trains using loglinear point process models. *Annals of Applied Statistics*. 2011;5.
- Kass RE, Ventura V. A spike-train probability model. *Neural Computation*. 2001; 13:1713–1720. [PubMed: 11506667]
- Kass RE, Ventura V, Brown EN. Statistical issues in the analysis of neuronal data. *Journal of Neurophysiology*. 2005; 94:8–25. [PubMed: 15985692]
- Kelly RC, Kass RE. A Framework for Evaluating Pairwise and Multiway Synchrony Among Stimulus-Driven Neurons. *Neural Computation*. 2012;1–26.
- Kottas A, Behseta S, Moorman DE, Poynor V, Olson CR. Bayesian nonparametric analysis of neuronal intensity rates. *Journal of Neuroscience Methods*. 2012;203.
- Lapish CC, Durstewitz D, Chandler JL, Seamans JK. Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proceedings of the National Academy of Sciences*. 2008; 105:11963–11968.
- Lindquist M, Xu Y, Nebel M, Caffo B. Evaluating Dynamic Bivariate Correlations in Resting-state fMRI: A comparison study and a new approach. *NeuroImage*. 2014 in press.
- McCullagh, CE.; Nelder, JA. *Generalized Linear Models*. 2. London: Chapman and Hall; 1989.
- Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics*. 2006; 34:1436–1462.
- Milad MR, Quirk GJ. Neurons in medial prefrontal cortex signal memory for fear extinction. *Nature*. 2002; 420:70–4. [PubMed: 12422216]
- Moorman DE, Aston-Jones G. Prefrontal neurons encode context-based response execution and inhibition in reward seeking and extinction. *Proceedings of the National Academy of Sciences*. 2015; 112:9472–9477.
- Motta G, Ombao H. Evolutionary Factor Analysis of Replicated Time Series. *Biometrics*. 2012; 68:825–836. [PubMed: 22364516]
- Narayanan NS, Laubach M. Methods for studying functional interactions among neuronal populations. *Methods in molecular biology*. 2009; 489:135–165. [PubMed: 18839091]
- Neal RM. Regression and classification using Gaussian process priors. *Bayesian Statistics*. 1998; 6:471–501.
- Ombao H, Van Belleghem S. Coherence Analysis: A Linear Filtering Point Of View. *IEEE Transactions on Signal Processing*. 2008; 56:2259–2266.
- Ombao H, von Sachs R, Guo W. SLEX Analysis of Multivariate Non-Stationary Time Series. *Journal of the American Statistical Association*. 2005; 100:519–531.

- Park T, Eckley I, Ombao H. Estimating the time-evolving partial coherence between signals via multivariate locally stationary wavelet processes. *IEEE Transactions on Signal Processing*. 2014 accepted.
- Patnaik D, Sastry P, Unnikrishnan K. Inferring Neuronal Network Connectivity from Spike Data: A Temporal Data mining Approach. *Scientific Programming*. 2008; 16:49–77.
- Peters J, Kalivas PW, Quirk GJ. Extinction circuits for fear and addiction overlap in prefrontal cortex. *Learning & Memory*. 2009; 16:279–288. [PubMed: 19380710]
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008; 454:995–9. [PubMed: 18650810]
- Quick H, Banerjee S, Carlin BP. Modeling temporal gradients in regionally aggregated California asthma hospitalization data. *Annals of Applied Statistics*. 2013; 7:154–176.
- Rasmussen, CE.; Williams, CKI. *Gaussian Processes for Machine Learning*. 2. MIT Press; 2006.
- Reich DS, Victor JD, Knight BW. The Power Ratio and the Interval Map: Spiking Models and Extracellular Recordings. *Journal of Neuroscience*. 1998; 18:10090–10104. [PubMed: 9822763]
- Rigat F, de Gunst M, van Pelt J. Bayesian Modeling and Analysis of Spatio-Temporal Neuronal Networks. *Bayesian Analysis*. 2006:733–764.
- Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S. The importance of mixed selectivity in complex cognitive tasks. *Nature*. 2013:585–590. advance online publication.
- Rothman AJ, Levina E, Zhu J. Generalized Thresholding of Large Covariance Matrices. *Journal of the American Statistical Association*. 2009; 104:177–186.
- Sakurai Y, Nakazono T, Ishino S, Terada S, Yamaguchi K, Takahashi S. Diverse synchrony of firing reflects diverse cell-assembly coding in the prefrontal cortex. *Journal of Physiology-Paris*. 2013; 107:459–470.
- Sampson PD, Guttorp P. Nonparametric Estimation of Nonstationary Spatial Covariance Structure. *Journal of the American Statistical Association*. 1992; 87:108–119.
- Sastry PS, Unnikrishnan KP. Conditional probability-based significance tests for sequential patterns in multineuronal spike trains. *Neural Comput*. 2010; 22:1025–1059. [PubMed: 19922295]
- Schmidt AM, O'Hagan A. Bayesian inference for non-stationary spatial covariance structure via spatial deformations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2003; 65:743–758.
- Shadlen MN, Kiani R. Decision making as a window on cognition. *Neuron*. 2013; 80:791–806. [PubMed: 24183028]
- Shahbaba B, Zhou B, Ombao H, Moorman D, Behseta S. A Semi-parametric Bayesian Model for Detecting Synchrony Among Multiple Neurons. *Neural Computation*. 2014; 26:2025–51. [PubMed: 24922500]
- Sharp C, Monterosso J, Montague PR. Neuroeconomics: a bridge for translational research. *Biological psychiatry*. 2012; 72:87–92. [PubMed: 22727459]
- Song P. Monte Carlo Kalman filter and smoothing for multivariate discrete state space models. 2000
- Stein, ML. *Statistical Interpolation of Spatial Data: Some Theory for Kriging*. Springer; 1999.
- Swihart, B.; Caffo, B.; Crainiceanu, C. A unified approach to modeling multivariate binary data using copulas over partitions. *Johns Hopkins University, Dept. of Biostatistics Working Papers*; 2010.
- West M. Hierarchical mixture models in neurological transmission analysis. *Journal of the American Statistical Association*. 2007; 92:587–606.
- Wilson, AG.; Ghahramani, Z. In: Flach, P.; De Bie, T.; Bristol, NC., editors. *Modelling Input Dependent Correlations between Multiple Responses*; *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*; UK: Springer; 2012.
- Yuan M, Lin Y. Model selection and estimation in the Gaussian graphical model. *Biometrika*. 2007; 94:19–35.
- Zeger SL. A regression model for time series of counts. *Biometrika*. 1988; 75:621–629.

Figure 1.

The Peri-stimulus time histogram (PSTH) of one neuron under two conditions. The upper plot is the continuous representation (point process) of the spike train data. Each line represents a trial and points on the line represent the firing time of the neuron. The lower plot is a summary of the spike train data by averaging activity across trials.

Figure 2.

A schematic representation of our *stationary* model (a) and *dynamic* model (b). Here, $r = 1, \dots, R$ is the index for trials; $t = 1, \dots, T$ is the index for time; and $i = 1, \dots, n$ is the index for neurons. Note that τ_{it} is a neuron-specific parameter; whereas $u_{\{i\}t}$ is defined jointly for all neurons.

Figure 3.

The Peri-Stimulus Time Histogram (PSTH) of 4 neurons with their estimated number of spikes. The red solid line is the estimated number of spikes, which is calculated by multiplying the estimated marginal firing probabilities by number of trials.

Figure 4.

Left: Time-varying correlation. For the first 80% of time, the two neurons are independent ($\zeta_t = 1$ for $t \leq 0.8$). For the remaining 20% of time, the two neurons become correlated ($\zeta_t > 1$ for $t > 0.8$). **Right:** The posterior estimates of correlation parameters, $\rho_{ij}^t = \rho_{ij} I_i^t I_j^t$, with 95% credible intervals. The solid lines are the posterior estimates and the gray areas are the 95% credible intervals. Note that in our model $\rho = 0$ corresponds to $\zeta = 1$.

Figure 5.

Evaluating our method's performance based on ROC curves for detecting dependence between a pair of neurons with time-varying ζ_r . Our proposed dynamic model is shown as solid red lines; the method of Kass et al. (2011) is shown as dashed green lines.

Figure 6.

95% intervals of the log predictive probability (LPP) on the test sets using our dynamic model (in red) and the method of Kass et al. (2011) (in green) for different values of ζ_t over the last 20% of time interval. Results are based on 100 simulated datasets.

Figure 7.

The estimated correlation parameters among the two neurons. The solid line is the posterior estimate and the gray area is the corresponding 95% credible interval.

Figure 8.

Graphical representation of correlation network under rewarded stimulus. The nodes are the neurons and the edges indicate the connection between neurons. The grayscale indicates the strength of connection.

Figure 9.

Graphical representation of correlation network under non-rewarded stimulus. The nodes are the neurons and the edges indicate the connection between neurons. The grayscale indicates the strength of connection.

Table 1

The posterior estimates of correlation parameters for simulated data with their 95% posterior intervals.

Parameter	Posterior Estimate
ρ_{12}	0.75 (0.66,0.84)
ρ_{13}	0.02 (−0.12,0.15)
ρ_{14}	0.12 (−0.01,0.25)
ρ_{23}	−0.05 (−0.19,0.12)
ρ_{24}	0.03 (−0.13,0.17)
ρ_{34}	−0.46 (−0.61, −0.32)