



HHS Public Access

Author manuscript

Neural Comput. Author manuscript; available in PMC 2015 March 29.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Published in final edited form as:

Neural Comput. 2014 September ; 26(9): 2025–2051. doi:10.1162/NECO_a_00631.

A Semiparametric Bayesian Model for Detecting Synchrony Among Multiple Neurons

Babak Shahbaba¹, Bo Zhou¹, Shiwei Lan¹, Hernando Ombao¹, David Moorman², and Sam Behseta³

¹Department of Statistics, UC Irvine, CA

²Department of Psychology, University of Massachusetts Amherst, MA

³Department of Mathematics, California State University, Fullerton, CA

Abstract

We propose a scalable semiparametric Bayesian model to capture dependencies among multiple neurons by detecting their co-firing (possibly with some lag time) patterns over time. After discretizing time so there is at most one spike at each interval, the resulting sequence of 1's (spike) and 0's (silence) for each neuron is modeled using the logistic function of a continuous latent variable with a Gaussian process prior. For multiple neurons, the corresponding marginal distributions are coupled to their joint probability distribution using a parametric copula model. The advantages of our approach are as follows: the nonparametric component (i.e., the Gaussian process model) provides a flexible framework for modeling the underlying firing rates; the parametric component (i.e., the copula model) allows us to make inference regarding both contemporaneous and lagged relationships among neurons; using the copula model, we construct multivariate probabilistic models by separating the modeling of univariate marginal distributions from the modeling of dependence structure among variables; our method is easy to implement using a computationally efficient sampling algorithm that can be easily extended to high dimensional problems. Using simulated data, we show that our approach could correctly capture temporal dependencies in firing rates and identify synchronous neurons. We also apply our model to spike train data obtained from prefrontal cortical areas.

Keywords

Spike Train; Synchrony; Gaussian Process; Copula

1 Introduction

Neurophysiological studies commonly involve modeling a sequence of spikes (action potentials) over time, known as a *spike train*, for each neuron. However, complex behaviors are driven by networks of neurons instead of a single neuron. In this paper, we propose a flexible, yet robust semiparametric Bayesian method for capturing temporal cross-dependencies among multiple neurons by simultaneous modeling of their spike trains. In contrast to most existing methods, our approach provides a flexible, yet powerful and scalable framework that can be easily extended to high dimensional problems.

For many years preceding ensemble recording, neurons were recorded successively and then combined into synthetic populations based on shared timing. Although this technique continues to produce valuable information to this day (Meyer and Olson, 2011), investigators are gravitating more and more towards simultaneous recording of multiple single neurons (Miller and Wilson, 2008). A major reason that multiple-electrode recording techniques have been embraced is because of the ability to identify the activity and dynamics of populations of neurons simultaneously. It is widely appreciated that groups of neurons encode variables and drive behaviors (Buzsáki, 2010).

Early analysis of simultaneously recorded neurons focused on correlation of activity across pairs of neurons using cross correlation analyses (Narayanan and Laubach, 2009) and analyses of changes in correlation over time, i.e., by using a joint peristimulus time histogram (JPSTH) (Gerstein and Perkel, 1969) or rate correlations (Narayanan and Laubach, 2009). Similar analyses can be performed in the frequency domain by using coherence analysis of neuron pairs using Fourier-transformed neural activity (Brown et al., 2004). These methods attempt to distinguish *exact synchrony* or *lagged synchrony* between a pair of neurons. Subsequently, a class of associated methods were developed for addressing the question of whether exact or lagged synchrony in a pair of neurons is merely due to chance. Later, to test the statistical significance of synchrony, a variety of methods, such as bootstrap confidence intervals, were introduced (Harrison et al., 2013).

To detect the presence of conspicuous spike coincidences in multiple neurons, Grün et al. (2002) proposed a novel method, where such conspicuous coincidences, called *unitary events*, are defined as joint spike constellations that recur more often than what can be explained by chance alone. In their approach, simultaneous spiking events from N neurons are modeled as a joint process composed of N parallel point processes. To test the significance of unitary events, they developed a new method, called joint-surprise, which measures the cumulative probability of finding the same or even larger number of observed coincidences by chance.

Pillow et al. (2008) investigate how correlated spiking activity in complete neural populations depends on the pattern of visual stimulation. They propose to use a generalized linear model to capture the encoding of stimuli in the spike trains of a neural population. In their approach, a cell's input is presented by a set of linear filters and the summed filter responses are exponentiated to obtain an instantaneous spike rate. The set of filters include a stimulus filter, a post-spike filter (to capture dependencies on history), and a set of coupling filter (to capture dependencies on the recent spiking of other cells).

Recent developments in detecting synchrony among neurons include models that account for trial to trial variability and the evolving intensity of firing rates between multiple trials. 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); Rigat et al. (2006); Patnaik et al. (2008); Diekman et al. (2009); Sastry and Unnikrishnan (2010); Kottas et al. (2012).

In a recent work, Kelly and Kass (2012) proposed a new method to quantify synchrony. They argue that separating stimulus effects from history effects would allow for a more

precise estimation of the instantaneous conditional firing rate. Specifically, given the firing history H_t , define $\lambda^A(t|H_t^A)$, $\lambda^B(t|H_t^B)$, and $\lambda^{AB}(t|H_t^{AB})$ to be the conditional firing intensities of neuron A, neuron B, and their synchronous spikes respectively. Independence between the two point processes can be examined by testing the null hypothesis $H_0 : \zeta(t) = 1$,

1, where $\zeta(t) = \frac{\lambda^{AB}(t|H_t^{AB})}{\lambda^A(t|H_t^A)\lambda^B(t|H_t^B)}$. The quantity $[\zeta(t)-1]$ can be interpreted as the deviation of co-firing from what is predicted by independence. Note that we still need to model the marginal probability of firing for each neuron. To do this, one could assume that a spike train follows a Poisson process, which is the simplest form of point processes. The main limitation of this approach is that it assumes that the number of spikes within a particular time frame follows a Poisson distribution. It is, however, very unlikely that actual spike trains follow this assumption (Barbieri et al., 2001; Kass and Ventura, 2001; Reich et al., 1998; Kass et al., 2005; Jacobs et al., 2009). One possible remedy is to use inhomogeneous Poisson process, which assumes time-varying firing rates. See Brillinger (1988); Brown et al. (2004); Kass et al. (2005); West (2007); Rigat et al. (2006); Cunningham et al. (2007); Berkes et al. (2009); Kottas and Behseta (2010); Sacerdote et al. (2012); Kottas et al. (2012) for more alternative methods of modeling spike trains.

In this paper, we propose a new semiparametric method for neural decoding. We first discretize time so that there is at most one spike within each time interval and let the response variable be a binary process comprised of 1s and 0s. We then use a continuous latent variable with Gaussian process prior to model the time-varying and history-dependent firing rate for each neuron. The covariance function for the Gaussian process is specified in a way that it creates prior positive autocorrelation for the latent variable so the firing rate could depend on spiking history. For each neuron, the marginal probability of firing within an interval is modeled by the logistic function of its corresponding latent variable. The main advantage of our model is that it connects the joint distribution of spikes for multiple neurons to their marginals by a parametric copula model in order to capture their cross-dependencies. Another advantage is that our model allows for co-firing of neurons after some lag time.

Cunningham et al. (2007) also assume that the underlying non-negative firing rate 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.

Berkes et al. (2009) also propose a copula model for capturing neural dependencies. They explore a variety of copula models for joint neural response distributions and develop an efficient maximum likelihood procedure for inference. Unlike their method, our proposed copula model in this paper is specified within a semiparametric Bayesian framework that uses Gaussian process model to obtain smooth estimates of firing rates.

Throughout this paper, we study the performance of our proposed method using simulated data and apply it to data from an experiment investigating the role of prefrontal cortical area in rats with respect to reward-seeking behavior and inhibition of reward-seeking in the absence of a rewarded outcome. In this experiment, the activity of 5–25 neurons from

prefrontal cortical area was recorded. During recording, rats chose to either press or withhold presses to presented levers. Pressing lever 1 allowed the rat to acquire a sucrose reward while pressing lever 2 had no effect. (All protocols and procedures followed National Institute of Health guidelines for the care and use of laboratory animals.)

In what follows, Section 2, we first describe our Gaussian process model for the firing rate of a single neuron. In Section 3, we present our method for detecting cofiring (possibly after some lag time) patterns for two neurons. The extension of this method for multiple neurons is presented in Section 4. In Section 5, we provide the details of our sampling algorithms. Finally, in Section 6, we discuss future directions.

2 Gaussian process model of firing rates

To model the underlying firing rate, we use a Gaussian process model. First we discretize time so that there is at most one spike within each time interval. Denote the response variable, y_t , to be a binary time series comprised of 1s (spike) and 0s (silence). The firing rate for each neuron is assumed to depend on an underlying latent variable, $u(t)$, which has a Gaussian process prior. In statistics and machine learning, Gaussian processes are widely used as priors over functions. Similar to the Gaussian distribution, a Gaussian process is defined by its mean (usually set to 0 in prior) and its covariance function $C: f \sim \mathcal{GP}(0, C)$. Here, the function of interest is the underlying latent variable $u(t)$, which is a stochastic process indexed by time t . Hence, the covariance function is defined in terms of t . We use the following covariance form, which includes a wide range of smooth nonlinear functions (Rasmussen and Williams, 2006; Neal, 1998):

$$C_{ij} = \text{Cov}[u(t_i), u(t_j)] = \lambda^2 + \eta^2 \exp[-\rho^2(t_i - t_j)^2] + \delta_{ij}\sigma_\varepsilon^2$$

In this setting, ρ^2 and η^2 control smoothness and the height of oscillations respectively. λ , η , ρ and σ are hyperparameters with their own hyperpriors. Throughout this paper, we put $N(0, 3^2)$ prior on the log of these hyperparameters.

We specify the spike probability, p_t , within time interval t in terms of $u(t)$ through the following transformation:

$$p_t = \frac{1}{1 + \exp[-u(t)]},$$

As $u(t)$ increases, so does p_t .

The prior autocorrelation imposed by this model allows the firing rate to change smoothly over time. Note that this does not mean that we believe the firing patterns over a single trial are smooth. However, over many trials, our method finds a smooth estimate of the firing rate. The dependence on prior firing patterns is through the term $(t_i - t_j)$ in the covariance function. As this term decreases, the correlation between $u(t_i)$ and $u(t_j)$ increases. This is different from other methods (Kass and Ventura, 2001; Kelly and Kass, 2012) that are based on including an explicit term in the model to capture firing history. For our analysis of

experimental data, we discretize the time into 5 ms intervals so there is at most one spike within each interval. Therefore, the temporal correlations in our method are on a slow time scale (Harrison et al., 2013).

When there are R trials (i.e., R spike trains) for each neuron, we model the corresponding spike trains as conditionally independent given the latent variable $u(t)$. Note that we can allow for trial-to-trial variation by including a trial-specific mean parameter such that $[u(t)]^{(r)} \sim \mathcal{GP}(\mu_r, C)$, where $r = 1, \dots, R$, (R = total number of trials or spike trains).

Figure 1 illustrates this method using 40 simulated spike trains for a single neuron. The dashed line shows the true firing rate, $p_t = 5(4+3 \sin(3\pi t))$, for $t = 0, 0.01, \dots, 1$, the solid line shows the posterior expectation of the firing rate, and the gray area shows the corresponding 95% probability interval. The plus signs on the horizontal axis represent spikes over 100 time intervals for one of the 40 trials.

Figure 2 shows the posterior expectation of firing rate (blue curve) overlaid on the PSTH plot of a single neuron with 5 ms bin intervals from the experimental data (discussed above) recorded over 10 seconds.

3 Modeling dependencies between two neurons

Let y_t and z_t be binary data indicating presence or absence of spikes within time interval t for two neurons. Denote p_t to be the spike probability at interval t for the first neuron, and q_t to denote the spike probability at the same interval for the second neuron. Given the corresponding latent variables $u(t)$ and $v(t)$ with Gaussian process priors $\mathcal{GP}(0, C_u)$ and $\mathcal{GP}(0, C_v)$ respectively, we model these probabilities as $p_t = 1/\{1 + \exp[-u(t)]\}$ and $q_t = 1/\{1 + \exp[-v(t)]\}$.

If the the two neurons are independent, the probability of firing at the same time is $P(y_t = 1, z_t = 1) = p_t q_t$. In general, however, we can write the probability of firing simultaneously as the product of their individual probabilities multiplied by a factor, $p_t q_t \zeta$, where ζ represents the excess firing rate ($\zeta > 1$) or the suppression firing rate ($\zeta < 1$) due to dependence between two neurons (Ventura et al., 2005; Kelly and Kass, 2012). That is, ζ accounts for the excess joint spiking beyond what is explained by independence. For independent neurons, $\zeta = 1$. Sometimes, the extra firing can occur after some lag time L . That is, in general, $P(y_t = 1, z_{t+L} = 1) = p_t q_{t+L} \zeta$ for some L . Therefore, the marginal and joint probabilities are

$$P(y_t=1|p, q, \zeta, L)=p_t$$

$$P(y_t=0|p, q, \zeta, L)=1 - p_t$$

$$P(z_t=1|p, q, \zeta, L)=q_t$$

$$P(z_t=0|p, q, \zeta, L)=1 - q_t$$

$$P(y_t=1, z_{t+L}=1|p, q, \zeta, L)=p_t q_{t+L} \zeta$$

$$P(y_t=1, z_{t+L}=0|p, q, \zeta, L)=p_t - p_t q_{t+L} \zeta$$

$$P(y_t=0, z_{t+L}=1|p, q, \zeta, L)=q_{t+L} - p_t q_{t+L} \zeta$$

$$P(y_t=0, z_{t+L}=0|p, q, \zeta, L)=1 - p_t - q_{t+L} + p_t q_{t+L} \zeta$$

where

$$\frac{\max(p_t + q_{t+L} - 1, 0)}{p_t q_{t+L}} \leq \zeta \leq \frac{\min(p_t, q_{t+L})}{p_t q_{t+L}}$$

In this setting, the observed data include two neurons with R trials of spike trains (indexed by $r = 1, 2, \dots, R$) per neuron. Each trial runs for S seconds. We discretize time into T intervals (indexed by $t = 1, 2, \dots, T$) of length S/T such that there are at most 1 spike in within each interval. We assume that the lag L can take a finite set of values from $[-K, K]$ for some biologically meaningful K , and write the likelihood function as follows:

$$\begin{aligned} \mathcal{L}(\zeta, u, v) = & \sum_{k=0}^K 1_{k=L} \prod_{r=1}^R \left[\prod_{t=1}^{T-k} P(y_t^{(r)}, z_{t+k}^{(r)}) \prod_{t=T-k+1}^T P(y_t^{(r)}) \prod_{t=1}^k P(z_t^{(r)}) \right] + \\ & \sum_{k=-K}^{-1} 1_{k=L} \prod_{r=1}^R \left[\prod_{t=1}^{T+k} P(y_{t-k}^{(r)}, z_t^{(r)}) \prod_{t=1}^{-k} P(y_t^{(r)}) \prod_{t=T+k+1}^T P(z_t^{(r)}) \right] \end{aligned}$$

We put uniform priors on ζ and L over the assumed range. As mentioned above, the hyperparameters in the covariance function have weakly informative (i.e., broad) priors: we assume the log of these parameters has a $N(0, 3^2)$ prior. We use Markov Chain Monte Carlo algorithms to simulate samples from the posterior distribution of model parameters given the observed spike trains. See section 5 for more details.

3.1 Illustrative examples

In this section, we use simulated data to illustrate our method. We consider three scenarios: 1) two independent neurons, 2) two dependent neurons with exact synchrony ($L = 0$), and 3) Two dependent neurons with lagged co-firing. In each scenario, we assume a time-varying

firing rate for each neuron and simulate 40 spike trains given the underlying firing rate. For independent neurons, we set $\zeta = 1$, whereas $\zeta > 1$ for dependent neurons.

Two independent neurons—In the first scenario, we consider two independent neurons ($\zeta = 1$). We simulate the spike trains according to our model. The firing probability at time t is set to $0.25 - 0.1 \cos(2\pi t)$ for the first neuron and to $0.15 + 0.2t$ for the second neuron. For each neuron, we generated 40 trials of spike trains and divided each trial into 100 time intervals. The left panel of Figure 3 shows the corresponding Joint Peristimulus Time Histogram (JPSTH). Each cell represents the joint frequency of spikes (darker cells represent higher frequencies) for the two neurons at given times. The marginal distributions of spikes, i.e., Peristimulus Time Histogram (PSTH), for the first neuron is shown along the horizontal axis. The second neuron's PSTH is shown along the vertical axis. The right panel of Figure 3 shows the posterior distributions of ζ and L . For this example, the posterior distribution of ζ is concentrated around 1 with median and 95% posterior probability interval equal to 1.01 and [0.85, 1.12] respectively. This would strongly suggest that the two neurons are independent as expected. Further, the posterior probabilities of all lag values from -10 to 10 are quite small.

Two exact synchronous neurons—For our next example, we simulate data for two dependent neurons with synchrony (i.e., $L = 0$) and we set $\zeta = 1.6$. That is, the probability of co-firing at the same time is 60% higher than that of independent neurons. As before, for each neuron we generate 40 trials of spike trains each discretized into 100 time bins. In this case, the firing probabilities at time t for the two neurons are $0.25 - 0.1 \cos(2\pi t)$. Figure 4 shows their corresponding JPSTH along with the posterior distributions of ζ and L . The posterior median for ζ is 1.598 and the 95% posterior probability interval is [1.548, 1.666]. Therefore, ζ identifies the two neurons in exact synchrony with excess co-firing rate than what is expected by independence. Further, the posterior distribution of L shows that the two neurons are in exact synchrony.

Two dependent neurons with lagged co-firing—Similar to the previous example, we set the probability of co-firing to 60% higher than what we obtain by the independence assumption. Similar to the previous two simulations, we generate 40 trials of spike trains each discretized into 100 time bins. The firing probabilities of the first neurons at time t is set to $0.25 + 0.1 \sin(2\pi t)$. The second neuron has the same firing probability but at time $t + L$. For different trials, we randomly set L to 3, 4, or 5 with probabilities 0.2, 0.5, and 0.3 respectively. Figure 5 shows JPSTH along with the posterior distributions of ζ and L . As before, the posterior distribution of ζ can be used to detect the relationship between the two neurons. For this example, the posterior median and 95% posterior interval for ζ are 1.39 and [1.33, 1.44] respectively. Also, our method could identify the three lag values correctly.

3.2 Power analysis

Next, we evaluate the performance of our proposed approach. More specifically, we compare our approach to the method of Kass et al. (2011), in terms of statistical power for detecting synchronous neurons. To be precise, given the true value of ζ , we compare the ratio of correctly identifying synchrony between two neurons over a large number of

simulated pairs of spike trains. In their approach, Kass et al. (2011) find the marginal firing rate of each neuron using natural cubic splines and then evaluate the amount of excess joint spiking using the bootstrap method. Therefore, for our first simulation study, we generate datasets that conform with the underlying assumptions of both methods. More specifically, we first set the marginal firing rates to $p_t = q_t = 0.2 - 0.1 \cos(12\pi t)$, and then generate the spike trains for the two neurons given ζ (i.e., excess joint firing rate). The left panel of Figure 6 compares the two methods in terms of statistical power for different values of ζ and different number of trials (20, 30, and 40) each with 20 time intervals. For each simulation setting, we generate 240 datasets. In our method, we call the relationship between two neurons significant if the corresponding 95% posterior probability does not include 1. For the method proposed by Kass et al. (2011), we use the 95% bootstrap confidence intervals instead. As we can see, our method (solid curve) has substantially higher power compared to the method of Kass et al. (2011) (dashed curve). Additionally, our method correctly achieves 0.05 level (dotted line) when $\zeta = 1$ (i.e., the two neurons are independent).

For our second simulation, we generate datasets that do not conform with the underlying assumptions of the two methods. Let $Y = (y_1, \dots, y_T)$ and $Z = (z_1, \dots, z_T)$ denote the spike trains for two neurons. We first simulate y_t , i.e., absence or presence of spikes for the first neuron at time t , from $\text{Bernoulli}(p_t)$, where $p_t = 0.25 - 0.1 \cos(12\pi t)$ for $t \in [0, 0.2]$. Then, we simulate z_t for the second neuron from $\text{Bernoulli}(b_0 + b_1 y_t)$ for given values of b_0 and b_1 . We set b_0 (i.e., the baseline probability of firing for the second neuron) to 0.2. When $b_1 = 0$, the two neurons are independent. Positive values of b_1 leads to higher rates of co-firing between the two neurons. When b_1 is negative, the first neuron has an inhibitory effect on the second neuron. For given values of b_1 and number of trials (20, 30, and 40), we generate 240 datasets where each trial has 20 time intervals. The right panel of Figure 6 compares the two methods in terms of statistical power under different settings. As before, our method (solid curves) has higher statistical power compared to the method of Kass et al. (2011) (dashed curves).

3.3 Sensitivity analysis for trial-to-trial variability

As mentioned above, our method can be easily extended to allow for trial-to-trial variability. To examine how such variability can affect our current model, we conduct a sensitivity analysis. Similar to the procedure discussed in the previous section, we start by setting the underlying firing probabilities to $p_t = 0.4 + 0.1 \cos(12t)$ and $\zeta = 1.2$. For each simulated dataset, we set the number of trials to 20, 30, 40, and 50. We found that shifting the firing rate of each trial by a uniformly sampled constant around the true firing rate does not substantially affect our method's power since the Gaussian process model is still capable of estimating the underlying firing rate by averaging over trials. However, adding independent random noise to each trial (i.e., flipping a fraction of time bins from zero to one or from one to zero) could affect performance, especially if the noise rate (i.e., proportion of flips) is high and the number of trials is low. Figure 7 shows the power for different number of trials and varying noise rate from 0 to 10%. As we can see, the power of our method drops slowly as the percentage of noise increases. The drop is more substantial when the number of trials is small (i.e., 20). However, for a reasonable number of trials (e.g., 40 or 50) and a reasonable noise rate (e.g., about 5%) the drop in power is quite small.

3.4 Results for experimental data

We now use our method for analyzing a pair of neurons selected from the experiment discussed in the introduction. (We will apply our method to multiple neurons in the next section.) Although we applied our method to several pairs with different patterns, for brevity we present the results for two pair of neurons; for one pair, the relationship changes under different scenarios; for the other pair, the relationship remains the same under the two scenarios. Our data include 51 spike trains for each neuron under different scenario (rewarded vs. non-rewarded). Each trial runs for 10 seconds. We discretize the time into 5 ms intervals.

Case 1: Two neurons with synchrony under both scenarios—We first present our model’s results for a pair of neurons that appear to be in exact synchrony under both scenarios. Figure 8 shows the posterior distributions of ζ and L under different scenarios. As we can see, the posterior distributions of ζ in both cases are away from 1, and $L = 0$ has the highest posterior probability. These results are further confirmed by empirical results, namely, the number of co-firings, correlation coefficients, and the sample estimates of conditional probabilities presented in Figure 8.

Using the method of Kass et al. (2011), the p-values under the two scenarios are $3.2E-11$ and $1.4E-13$ respectively. While both methods provide similar conclusions, their method is limited to detecting exact synchrony only.

Case 2: Two neurons with synchrony under the rewarded scenario only—Next, we present our model’s results for a pair of neurons appear to be in a moderate synchrony under the rewarded scenario only. Figure 9 shows the posterior distributions of ζ and L under different scenarios. In this case, the posterior distributions of ζ is slightly away from 1 in the first scenario; however, under the second scenario, the tail probability of 1 is not negligible. These results are further confirmed by empirical results presented in Figure 9: only in the first scenario we observe a moderate difference between the conditional probabilities.

Using the method of Kass et al. (2011), the p-values under the two scenarios are $2E-4$ and 0.144 respectively. As discussed above, although for these data the two methods provide similar results in terms of synchrony, our method can be used to make inference regarding possible lag values. Moreover, as we will show in the next section, our method provides a hierarchical Bayesian framework that can be easily extended to multiple neurons.

4 Modeling dependencies among multiple neurons

Temporal relationships among neurons, particularly those that change across different contexts, can provide additional information beyond basic firing rates. Because it is possible to record spike trains from multiple neurons simultaneously, and because network encoding likely spans more than pairs of neurons, we now turn our attention to calculating temporally-related activities among multiple (> 2) simultaneously-recorded neurons.

At lag zero (i.e., $L = 0$), we can rewrite our model for the joint distribution of two neurons in terms of their individual cumulative distributions as follows (we have dropped the index t for simplicity):

$$H(y, z) = \mathcal{C}(F_1(y), F_2(z)) = [1 + \beta \prod_{i=1}^2 (1 - F_i)] \prod_{i=1}^2 F_i$$

where $F_1 = F_1(y) = P(Y = y)$, $F_2 = F_2(z) = P(Z = z)$, and $\beta = \frac{\zeta - 1}{(1 - p)(1 - q)}$. Note that in this case, $\beta = 0$ indicates that the two neurons are independent. In general, models that couple the joint distribution of two (or more) variables to their individual marginal distributions are called copula models. See Nelsen (1998) for detailed discussion of copula models. Let H be n -dimensional distribution functions with marginals F_1, \dots, F_n . Then, an n -dimensional copula is a function of the following form:

$$H(y_1, \dots, y_n) = \mathcal{C}(F_1(y_1), \dots, F_n(y_n)), \text{ for all } y_1, \dots, y_n$$

Here, \mathcal{C} defines the dependence structure between the marginals. Our model for two neurons is in fact a special case of the Farlie-Gumbel-Morgenstern (FGM) copula family (Farlie, 1960; Gumbel, 1960; Morgenstern, 1956; Nelsen, 1998). For n random variables Y_1, Y_2, \dots, Y_n , the FGM copula, \mathcal{C} , has the following form:

$$\mathcal{C} = \left[1 + \sum_{k=2}^n \sum_{1 \leq j_1 < \dots < j_k \leq n} \beta_{j_1 j_2 \dots j_k} \prod_{l=1}^k (1 - F_{j_l}) \right] \prod_{i=1}^n F_i \quad (2)$$

where $F_i = F_i(y_i)$. Restricting our model to second-order interactions, we can generalize our approach for two neurons to a copula-based model for multiple neurons using the FGM copula family,

$$H(y_1, \dots, y_n) = [1 + \sum_{1 \leq j_1 < j_2 \leq n} \beta_{j_1 j_2} \prod_{l=1}^2 (1 - F_{j_l})] \prod_{i=1}^n F_i$$

where $F_i = P(Y_i = y_i)$. Here, we use y_1, \dots, y_n to denote the firing status of n neurons at time t ; $\beta_{j_1 j_2}$ captures the relationship between the j_1^{th} and j_2^{th} neurons. To ensure that probability distribution functions remain within $[0, 1]$, the following constraints on all $\binom{n}{2}$ parameters $\beta_{j_1 j_2}$ are imposed:

$$1 + \sum_{1 \leq j_1 < j_2 \leq n} \beta_{j_1 j_2} \prod_{l=1}^2 \varepsilon_{j_l} \geq 0, \varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$$

Considering all possible combinations of ε_{j_1} and ε_{j_2} in the above condition, there are $n(n - 1)$ linear inequalities, which can be combined into the following inequality:

$$\sum_{1 \leq j_1 < j_2 \leq n} |\beta_{j_1 j_2}| \leq 1$$

4.1 Illustrative example

To illustrate this method, we follow a similar procedure as Section 3.4 and simulate spike trains for three neurons such that neurons 1 and 2 are in exact synchrony, but they are independent from neuron 3. Table 4 shows the estimated β 's along with their corresponding 95% posterior probability intervals using posterior samples from Spherical Hamiltonian Monte Carlo (Spherical HMC). Our method correctly detects the relationship among the neurons: for synchronous neurons, the corresponding β 's are significantly larger than 0 (i.e., 95% posterior probability intervals do not include 0), whereas the remaining β 's are close to 0 (i.e., 95% posterior probability intervals include 0).

4.2 Results for experimental data

We now use our copula-based method for analyzing the experimental data discussed earlier. As mentioned, during task performance the activity of multiple neurons was recorded under two conditions: rewarded stimulus (lever 1) and non-rewarded stimulus (lever 2). Here, we focus on 5 simultaneously recorded neurons. There are 51 trials per neuron under each scenario. We set the time intervals to 5 ms.

Tables 4.2 and 4.2 show the estimates of $\beta_{i,j}$, which capture the association between the i^{th} and j^{th} neurons, under the two scenarios. Figure 10 shows the schematic representation of these results under the two experimental conditions. The solid line indicates significant association.

Our results show that neurons recorded simultaneously in the same brain area are correlated in some conditions and not others. This strongly supports the hypothesis that population coding among neurons (here though correlated activity) is a meaningful way of signaling differences in the environment (rewarded or non-rewarded stimulus) or behavior (going to press the rewarded lever or not pressing) (Buzsáki, 2010). It also shows that neurons in the same brain region are differentially involved in different tasks, an intuitive perspective but one that is neglected by much of behavioral neuroscience. Finally, our results indicate that network correlation is dynamic and that functional pairs—again, even within the same brain area—can appear and disappear depending on the environment or behavior. This suggests (but does not confirm) that correlated activity across separate populations within a single brain region can encode multiple aspects of the task. For example, the pairs that are correlated in reward and not in non-reward could be related to reward-seeking whereas pairs that are correlated in non-reward could be related to response inhibition. Characterizing neural populations within a single brain region based on task-dependent differences in correlated firing is a less-frequently studied phenomenon compared to the frequently pursued goal of identifying the overall function of the brain region based on individual neural firing (Stokes et al., 2013). While our data only begin to address this important

question, the developed model will be critical in application to larger neural populations across multiple tasks in our future research.

5 Computation

We use Markov Chain Monte Carlo (MCMC) algorithms to sample from posterior distribution. The typical number of MCMC iterations is 3000 after discarding pre-convergence samples. Algorithm 1 in Appendix shows the overall sampling procedure. We use the slice sampler (Neal, 2003) for the hyperparameters controlling the covariance function of the Gaussian process model. More specifically, we use the “stepping out” procedure to find an interval around the current state, and then the “shrinkage” procedure to sample from this interval. For latent variables with Gaussian process priors, we use the elliptical slice sampling algorithm proposed by Murray et al. (2010). The details are provided in Algorithm 2 in the appendix.

Sampling from the posterior distribution of β 's in the copula model is quite challenging. As the number of neurons increases, simulating samples from the posterior distribution these parameters becomes difficult because of the imposed constraints (Neal and Roberts, 2008; Sherlock and Roberts, 2009; Neal et al., 2012; Brubaker et al., 2012; Pakman and Paninski, 2012). We have recently developed a new Markov Chain Monte Carlo algorithm for constrained target distributions (Lan et al., 2014) based on Hamiltonian Monte Carlo (HMC) (Duane et al., 1987; Neal, 2011).

In many cases, bounded connected constrained D -dimensional parameter spaces can be bijectively mapped on to the D -dimensional unit ball

$\mathbf{B}_0^D(1) := \{\theta \in \mathbb{R}^D : \|\theta\|_2 = \sqrt{\sum_{i=1}^D \theta_i^2} \leq 1\}$, where θ are parameters. Therefore, our method first maps the D -dimensional constrained domain of parameters to the unit ball. We then augment the original D -dimensional parameter θ with an extra auxiliary variable $\tilde{\theta}_{D+1}$ to form an extended $(D + 1)$ -dimensional parameter $\tilde{\theta} = (\theta, \tilde{\theta}_{D+1})$ such that $\|\tilde{\theta}\|_2 = 1$ so

$\theta_{D+1} = \pm \sqrt{1 - \|\theta\|_2^2}$. This way, the domain of the target distribution is changed from the unit ball $\mathbf{B}_0^D(1)$ to the D -dimensional sphere, $\mathbf{S}^D := \{\tilde{\theta} \in \mathbb{R}^{D+1} : \|\tilde{\theta}\|_2 = 1\}$, through the following transformation:

$$T_{\mathbf{B} \rightarrow \mathbf{S}} : \mathbf{B}_0^D(1) \rightarrow \mathbf{S}^D, \theta \mapsto \tilde{\theta} = (\theta, \pm \sqrt{1 - \|\theta\|_2^2})$$

Note that although $\tilde{\theta}_{D+1}$ can be either positive or negative, its sign does not affect our Monte Carlo estimates since after applying the above transformation, we need to adjust our estimates according to the change of variable theorem as follows:

$$\int_{\mathbf{B}_0^D(1)} f(\theta) d\theta_{\mathbf{B}} = \int_{\mathbf{S}^D_+} f(\tilde{\theta}) \left| \frac{d\theta_{\mathbf{B}}}{d\tilde{\theta}_{\mathbf{S}}} \right| d\tilde{\theta}_{\mathbf{S}}$$

where $\left| \frac{d\theta_B}{d\theta_S} \right| = |\theta_{D+1}|$. Here, $d\theta_B$ and $d\theta_S$ are under Euclidean measure and spherical measure respectively.

Using the above transformation, we define the dynamics on the sphere. This way, the resulting HMC sampler can move freely on S^D while implicitly handling the constraints imposed on the original parameters. As illustrated in Figure 11, the boundary of the constraint, i.e., $\|\theta\|_2 = 1$, corresponds to the equator on the sphere S^D . Therefore, as the sampler moves on the sphere, passing across the equator from one hemisphere to the other translates to “bouncing back” off the the boundary in the original parameter space.

We have shown that by defining HMC on the sphere, besides handling the constraints implicitly, the computational efficiency of the sampling algorithm could be improved since the resulting dynamics has a partial analytical solution (geodesic flow on the sphere). We use this approach, called Spherical HMC, for sampling from the posterior distribution of β 's in the copula model. See Algorithm 3 in Appendix for more details.

Using parallelization (i.e., assigning each neuron to a server), our computational method can handle relatively a large number of neurons. The MATLAB implementation of our method runs on a HPC (High Performance Computing) Beowulf cluster with a total of 64 CPUs (CentOS) and equipped with a GPU. For 10 neurons, 20 trials, and 50 time bins, each iteration of MCMC takes 8.4 seconds with acceptance probability of 0.72. For 50 neurons, the time per iteration increases to 24.5 with similar acceptance probability.

While our proposed method takes advantage of several advanced computational techniques, the current implementation of our method is only useful for tens of neurons. This can be improved in future by reducing the computational complexity of our sampling algorithms. For the GP model, the computational complexity is $\mathcal{O}(T^3)$, where T is the number of time bins. This increases linearly with the number of neurons. Note that we can reduce the computational complexity of the GP model to $\mathcal{O}(T)$ by using a Brownian motion instead. To calculate the joint pdf within each time bin of a trial (i.e., using the copula model), the computational complexity increases exponentially by the number of coffering neurons. The overall complexity increases linearly by the number of trials and the number of time bins.

The space complexity depends on the number of neurons, number of time bins, and number of trials. For a simulation study with 25 neurons, 25 trials, and 50 time bins, the RAM usage is around 4.3GB. Increasing the number of neurons to 50 results in a substantially higher RAM usage close to 7.2GB. If we further increase the number of trials to 50, the RAM usage increases to 10.5GB. If we also increase the number of time bins to 100, the RAM usage increases to 12.2GB.

All computer programs and simulated datasets discussed in this paper are available online at <http://www.ics.uci.edu/~babaks/Site/Codes.html>.

6 Discussion

The method we proposed in this paper benefits from multiple aspects, including flexibility, computational efficiency, interpretability, and generalizability. The latter is especially important because the model offered in this work can be adopted for other computationally intensive biological problems.

We believe that a forte of our proposed method is that it offers a modeling approach to the problems of identification and calibration of co-firings at various lag times. Consequently, from the statistical inferential perspective, our hybrid modeling approach would fair better when compared with methods such as cross-correlation analysis, mainly because not only it sheds light on how signals in the network are communicating, but also it informs the scientist of the sharpness of those cross relationships through posterior confidence intervals. Also, note that although it is possible to run other methods, e.g., the method of Kass et al. (2011), for different lags and perform multiple hypothesis testing, our method offers a modeling paradigm for measuring lagged and exact synchrony at the same time. This is important because it avoids complex and multistage testing procedures for pairs of neurons.

The sampling algorithm proposed for detecting synchrony among multiple neurons is advantageous over commonly used MCMC techniques such as Metropolis-Hastings. This fact becomes even more salient especially considering that the current technology provides high-dimensional data by allowing the simultaneous recording of many neurons. Developing efficient sampling algorithms for such problems has been discussed by Ahmadian et al. (2011).

The analysis presented offers a number of ways in which examining the temporal relationship of activity in multiple neurons can reveal information about population dynamics of neuronal circuits. These kinds of data are critical in going beyond treating populations as averages of single neurons, as is commonly done in physiological studies. They also allow us to ask the question of whether neuronal firing heterogeneity contributes towards a unified whole (Stokes et al., 2013), or whether separates populations in a brain area are differentially involved in different aspects of behavior (Buschman et al., 2012).

In our current model, β are fixed. However, dependencies among neurons could in fact change over time. To address this issue, we could allow β 's to be piecewise constant over time to capture non-stationary neural connections. Change-point detection would of course remain a challenge.

Acknowledgments

This work was supported by NSF grant IIS-1216045 and NIH grant R01-AI107034. DEM was supported by NIH grants DA032005, DA015369, and MH092868. HO was supported by NSF grants from SES and DMS. The first and last authors would like to acknowledge Professor Touraj Daryae for many a thought-provoking conversation.

Appendix

Algorithm 1

Sampling latent variables, copula parameters, and hyperparameters

Initialize the matrix of latent variables, $U|_{(n,D)}$, where the i^{th} column corresponds to the latent variables of the i^{th} neuron, n is the number of time bins, and D is the number of neurons.

Initialize the hyperparameters, θ , which specify the Gaussian process priors for the latent variables.

Initialize the copula model parameters, β , as a $D(D - 1)/2$ vector.

for $i = 1, \dots, B$ **do**

 Sample $U^{(i+1)}$ from posterior distribution conditional on $U^{(i)}$, $\theta^{(i)}$ and $\beta^{(i)}$, $P(U^{(i+1)}|Y, U^{(i)}, \theta^{(i)}, \beta^{(i)})$, using the elliptical slice sampler (Algorithm 2).

for $j = 1, \dots, D$ **do**

 Sample $\theta_j^{(i+1)}$ from the posterior distribution of the hyperparameters of the j^{th} latent variable conditional on the latent variables, $P(\theta_j^{(i+1)}|U_j^{(i+1)}, \theta_j^{(i)})$, using the slice sampler (Neal, 2003).

end for

 Sample $\beta^{(i+1)}$ from the posterior distribution conditional on the latent variables, $P(\beta^{(i+1)}|Y, U^{(i+1)}, \theta^{(i)})$, using Spherical HMC presented (Algorithm 3).

end for

Algorithm 2

Elliptical slice sampler for latent variables

Let U be the current state of the latent variables.

Sample $U^* \sim N(0, \Sigma)$, where Σ is the covariance matrix of the Gaussian process.

Calculate the log-likelihood threshold for the elliptical slice sampler,

$v \sim \text{Uniform}[0, 1]$

$\log y \leftarrow \log(L(U)) + \log(v)$

Let α be the angle for the slice.

Draw a proposal and define the corresponding bracket,

$\alpha \sim \text{Uniform}[0, 2\pi]$

$(\alpha_{\min}, \alpha_{\max}) \leftarrow (\alpha - 2\pi, \alpha)$

Set $U' \leftarrow U \cos(\alpha) + U^* \sin(\alpha)$

while $\log(L(U')) < \log y$ **do**

if $\alpha < 0$ **then**

$\alpha_{\min} \leftarrow \alpha$

else

$\alpha_{\max} \leftarrow \alpha$

end if

$\alpha \sim \text{Uniform}[\alpha_{\min}, \alpha_{\max}]$

$U' \leftarrow U \cos(\alpha) + U^* \sin(\alpha)$

end while

Return U' as the new state.

Algorithm 3

Spherical HMC for copula parameters

Initialize the copula parameters, $\beta^{(1)}$, along with their appropriate transformation, $\tilde{\beta}^{(1)}$, at the current state.

Sample a new momentum value $v^{(1)} \sim \mathcal{N}(0, I_{D+1})$

Define the potential energy, U , as minus log density of $\tilde{\beta}$ and the kinetic energy, K , as minus log density of v .

Set $\tilde{v}^{(1)} \leftarrow v^{(1)} - \tilde{\beta}^{(1)}(\beta^{(1)})^T v^{(1)}$

Calculate the Hamiltonian function: $H(\beta^{(1)}, v^{(1)}) = U(\beta^{(1)}) + K(v^{(1)})$

for $\ell = 1$ to L **do**

$$\tilde{v}^{(\ell+\frac{1}{2})} = \tilde{v}^{(\ell)} - \frac{\varepsilon}{2} \left(\begin{bmatrix} I_D \\ 0 \end{bmatrix} - \tilde{\beta}^{(\ell)} (\beta^{(\ell)})^T \right) \nabla U(\beta^{(\ell)})$$

$$\tilde{\beta}^{(\ell+1)} = \tilde{\beta}^{(\ell)} \cos(\|\tilde{v}^{(\ell+\frac{1}{2})}\| \varepsilon) + \frac{\tilde{v}^{(\ell+\frac{1}{2})}}{\|\tilde{v}^{(\ell+\frac{1}{2})}\|} \sin(\|\tilde{v}^{(\ell+\frac{1}{2})}\| \varepsilon)$$

$$\tilde{v}^{(\ell+\frac{1}{2})} \leftarrow -\tilde{\beta}^{(\ell)} \left\| \tilde{v}^{(\ell+\frac{1}{2})} \right\| \sin(\left\| \tilde{v}^{(\ell+\frac{1}{2})} \right\| \varepsilon) + \tilde{v}^{(\ell+\frac{1}{2})} \cos(\left\| \tilde{v}^{(\ell+\frac{1}{2})} \right\| \varepsilon)$$

$$\tilde{v}^{(\ell+1)} = \tilde{v}^{(\ell+\frac{1}{2})} - \frac{\varepsilon}{2} \left(\begin{bmatrix} I_D \\ 0 \end{bmatrix} - \tilde{\beta}^{(\ell+1)} (\beta^{(\ell+1)})^T \right) \nabla U(\beta^{(\ell+1)})$$

end for

Calculate $H(\tilde{\beta}^{(L+1)}, \tilde{v}^{(L+1)}) = U(\tilde{\beta}^{(L+1)}) + K(\tilde{v}^{(L+1)})$

Calculate the acceptance probability

$\alpha = \min(1, \exp\{-H(\tilde{\beta}^{(L+1)}, \tilde{v}^{(L+1)}) + H(\beta^{(1)}, v^{(1)})\})$

Accept or reject the proposal according to α

References

- Ahmadian Y, Pillow JW, Paninski L. Efficient Markov Chain Monte Carlo methods for decoding neural spike trains. *Neural Computation*. 2011; 23(1):46–96. [PubMed: 20964539]
- 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]
- 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.
- 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(5):456–461.
- Brubaker, M.A.; Salzmann, M.; Urtasun, R. A family of mcmc methods on implicitly defined manifolds. In: Lawrence, N.D.; Girolami, M.A., editors. *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS-12)*; 2012. p. 161-172.
- Buschman TJ, Denovellis EL, Diogo C, Bullock D, Miller EK. Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron*. 2012; 76(4):838–846. [PubMed: 23177967]
- Buzsáki G. Neural syntax: Cell assemblies, synapsemsembles, and readers. *Neuron*. 2010; 68(3):362–385. [PubMed: 21040841]

- 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.
- Diekman CO, Sastry PS, Unnikrishnan KP. Statistical significance of sequential firing patterns in multi-neuronal spike trains. *Journal of neuroscience methods*. 2009; 182(2):279–284. [PubMed: 19559053]
- Duane S, Kennedy A, Pendleton BJ, Roweth D. Hybrid monte carlo. *Physics Letters B*. 1987; 195(2): 216–222.
- Farlie DJG. The performance of some correlation coefficients for a general bivariate distribution. *Biometrika*. 1960; 47(3/4)
- Gerstein GL, Perkel DH. Simultaneously recorded trains of action potentials: Analysis and functional interpretation. *Science*. 1969; 164(3881):828–830. [PubMed: 5767782]
- Grün S, Diesmann M, Aertsen A. Unitary events in multiple single-neuron spiking activity: I. detection and significance. *Neural Computation*. 2002; 14(1):43–80. [PubMed: 11747534]
- Gumbel EJ. Bivariate exponential distributions. *Journal of the American Statistical Association*. 1960; 55:698–707.
- 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.
- 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(14):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. Bayesian nonparametric modeling for comparison of single-neuron firing intensities. *Biometrics*. 2010;277–286. [PubMed: 19432774]
- Kottas A, Behseta S, Moorman DE, Poynor V, Olson CR. Bayesian nonparametric analysis of neuronal intensity rates. *Journal of Neuroscience Methods*. 2012; 203(1)
- Lan, S.; Zhou, B.; Shahbaba, B. Spherical Hamiltonian Monte Carlo for constrained target distributions; *Proceedings of the 31th International Conference on Machine Learning (ICML)*; 2014.
- Meyer T, Olson C. Statistical learning of visual transitions in monkey in-ferotemporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*. 2011; 108:19401–19406. [PubMed: 22084090]
- Miller E, Wilson M. All my circuits: Using multiple electrodes to understand functioning neural networks. *Neuron*. 2008; 60(3):483–488. [PubMed: 18995823]
- Morgenstern D. Einfache beispiele zweidimensionaler verteilungen. *Mitteilungsblatt für Mathematische Statistik*. 1956; 8:234–235.
- Murray I, Adams RP, MacKay DJ. Elliptical slice sampling. *JMLR: W&CP*. 2010; 9:541–548.
- Narayanan NS, Laubach M. Methods for studying functional interactions among neuronal populations. *Methods in molecular biology*. 2009; 489:135–165. [PubMed: 18839091]
- Neal P, Roberts GO. Optimal scaling for random walk metropolis on spherically constrained target densities. *Methodology and Computing in Applied Probability*. 2008; 10(2):277–297.
- Neal P, Roberts GO, Yuen WK. Optimal scaling of random walk metropolis algorithms with discontinuous target densities. *Annals of Applied Probability*. 2012; 22(5):1880–1927.
- Neal RM. Regression and classification using Gaussian process priors. *Bayesian Statistics*. 1998; 6:471–501.
- Neal RM. Slice sampling. *Annals of Statistics*. 2003; 31(3):705–767.

- Neal, RM. MCMC using Hamiltonian dynamics. In: Brooks, S.; Gelman, A.; Jones, G.; Meng, XL., editors. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC; 2011. p. 113-162.
- Nelsen, RB. *An Introduction to Copulas (Lecture Notes in Statistics)*. 1 edition. Springer; 1998.
- Pakman A, Paninski L. Exact Hamiltonian Monte Carlo for Truncated Multivariate Gaussians. ArXiv e-prints. 2012
- Patnaik D, Sastry P, Unnikrishnan K. Inferring neuronal network connectivity from spike data: A temporal data mining approach. *Scientific Programming*. 2008; 16:49–77.
- 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(7207): 995–999. [PubMed: 18650810]
- Rasmussen, CE.; Williams, CKI. *Gaussian Processes for Machine Learning*. 2nd edition. 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(23):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.
- Sacerdote L, Tamborrino M, Zucca C. Detecting dependencies between spike trains of pairs of neurons through copulas. *Brain Res*. 2012; 1434:243–256. [PubMed: 21981802]
- Sastry PS, Unnikrishnan KP. Conditional probability-based significance tests for sequential patterns in multineuronal spike trains. *Neural Comput*. 2010; 22(4):1025–1059. [PubMed: 19922295]
- Sherlock C, Roberts GO. Optimal scaling of the random walk metropolis on elliptically symmetric unimodal targets. *Bernoulli*. 2009; 15(3):774–798.
- Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J. Dynamic coding for cognitive control in prefrontal cortex. *Neuron*. 2013; 78(2):364–375. [PubMed: 23562541]
- Ventura V, Cai C, Kass RE. Statistical assessment of time-varying dependency between two neurons. *J Neurophysiol*. 2005; 94(4):2940–2947. [PubMed: 16160097]
- West M. Hierarchical mixture models in neurological transmission analysis. *Journal of the American Statistical Association*. 2007; 92:587–606.

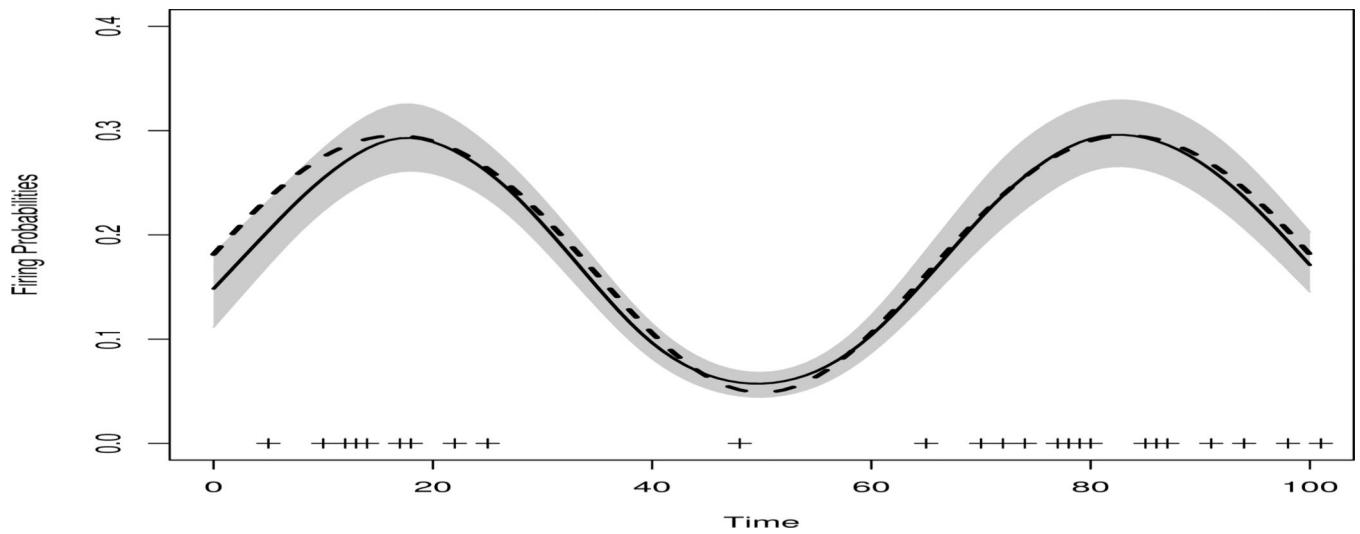


Figure 1.

An illustrative example for using a Gaussian process model for a neuron with 40 trials. The dashed line shows the true firing rate, the solid line shows the posterior expectation of the firing rate, and the gray area shows the corresponding 95% probability interval. The plus signs on the horizontal axis represent spikes over 100 time intervals for one of the 40 trials.

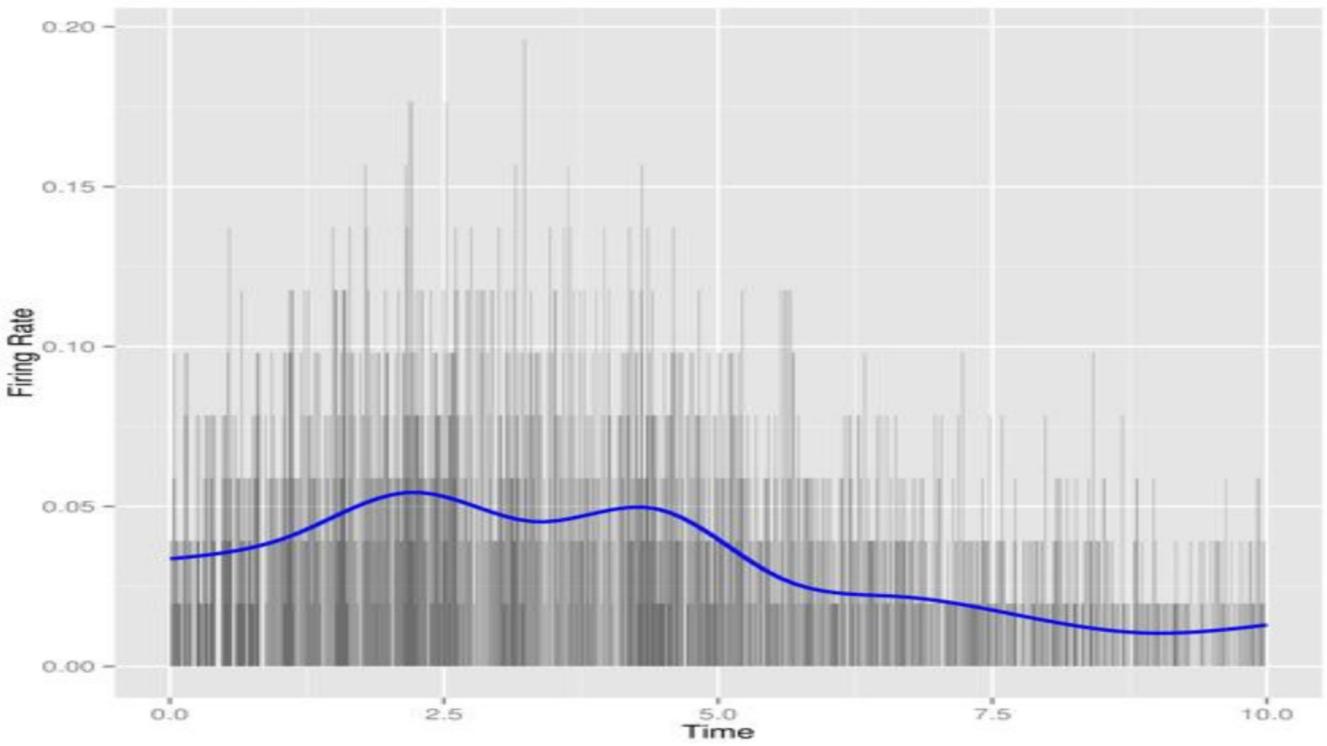


Figure 2.

Using our Gaussian process model to capture the underlying firing rate of a single neuron from prefrontal cortical areas in rat's brain. There are 51 spike trains recorded over 10 seconds. The PSTH plot is generated by creating 5 ms intervals. The curve shows the estimated firing rate (posterior expectation).

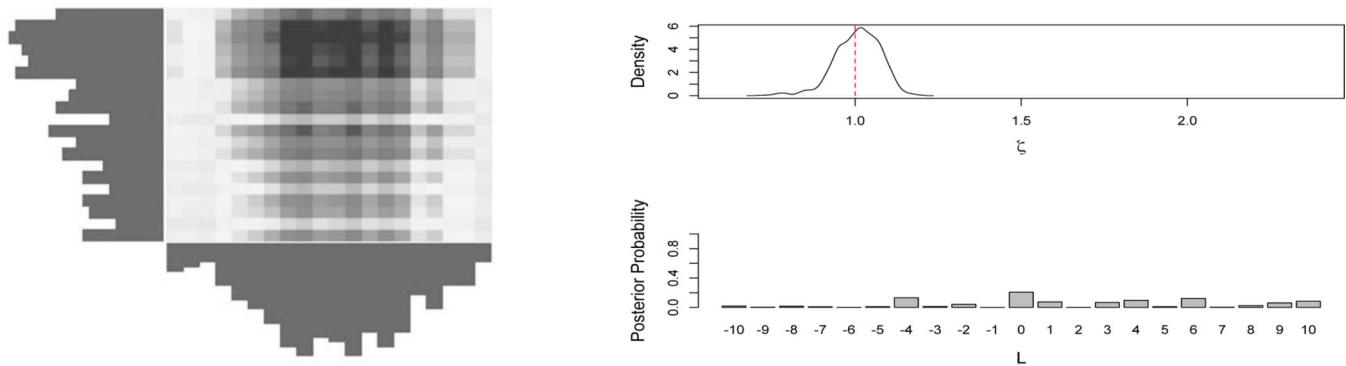


Figure 3. Two independent neurons

The left panel shows the corresponding Joint Peri-Stimulus Time Histogram (JPSTH). The right panel shows the posterior distributions of ζ and L . Darker cells represent higher frequencies.

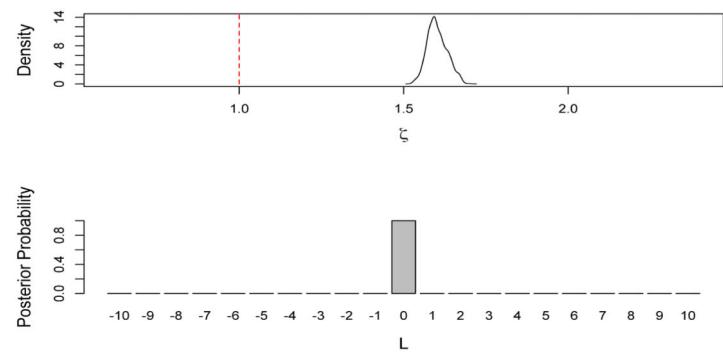
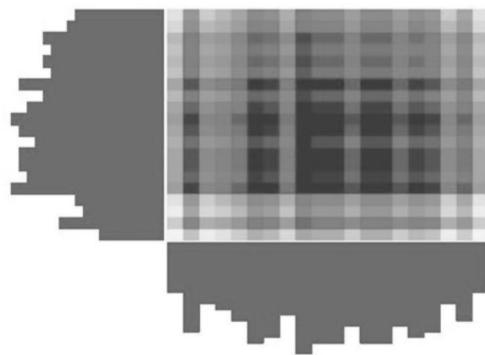


Figure 4. Two dependent neurons in exact synchrony

The left panel show the frequency of spikes over time. The right panel shows the posterior distribution of ζ and L . Darker cells represent higher frequencies.

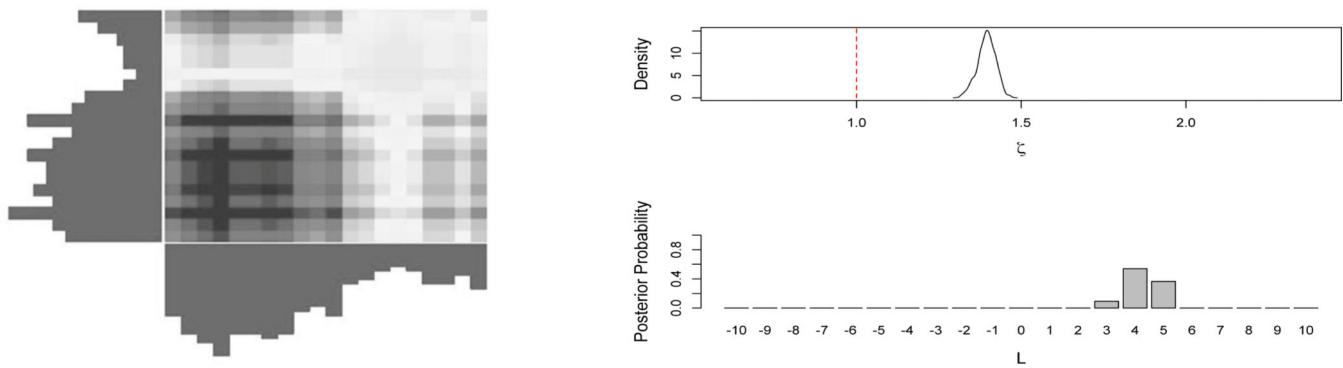
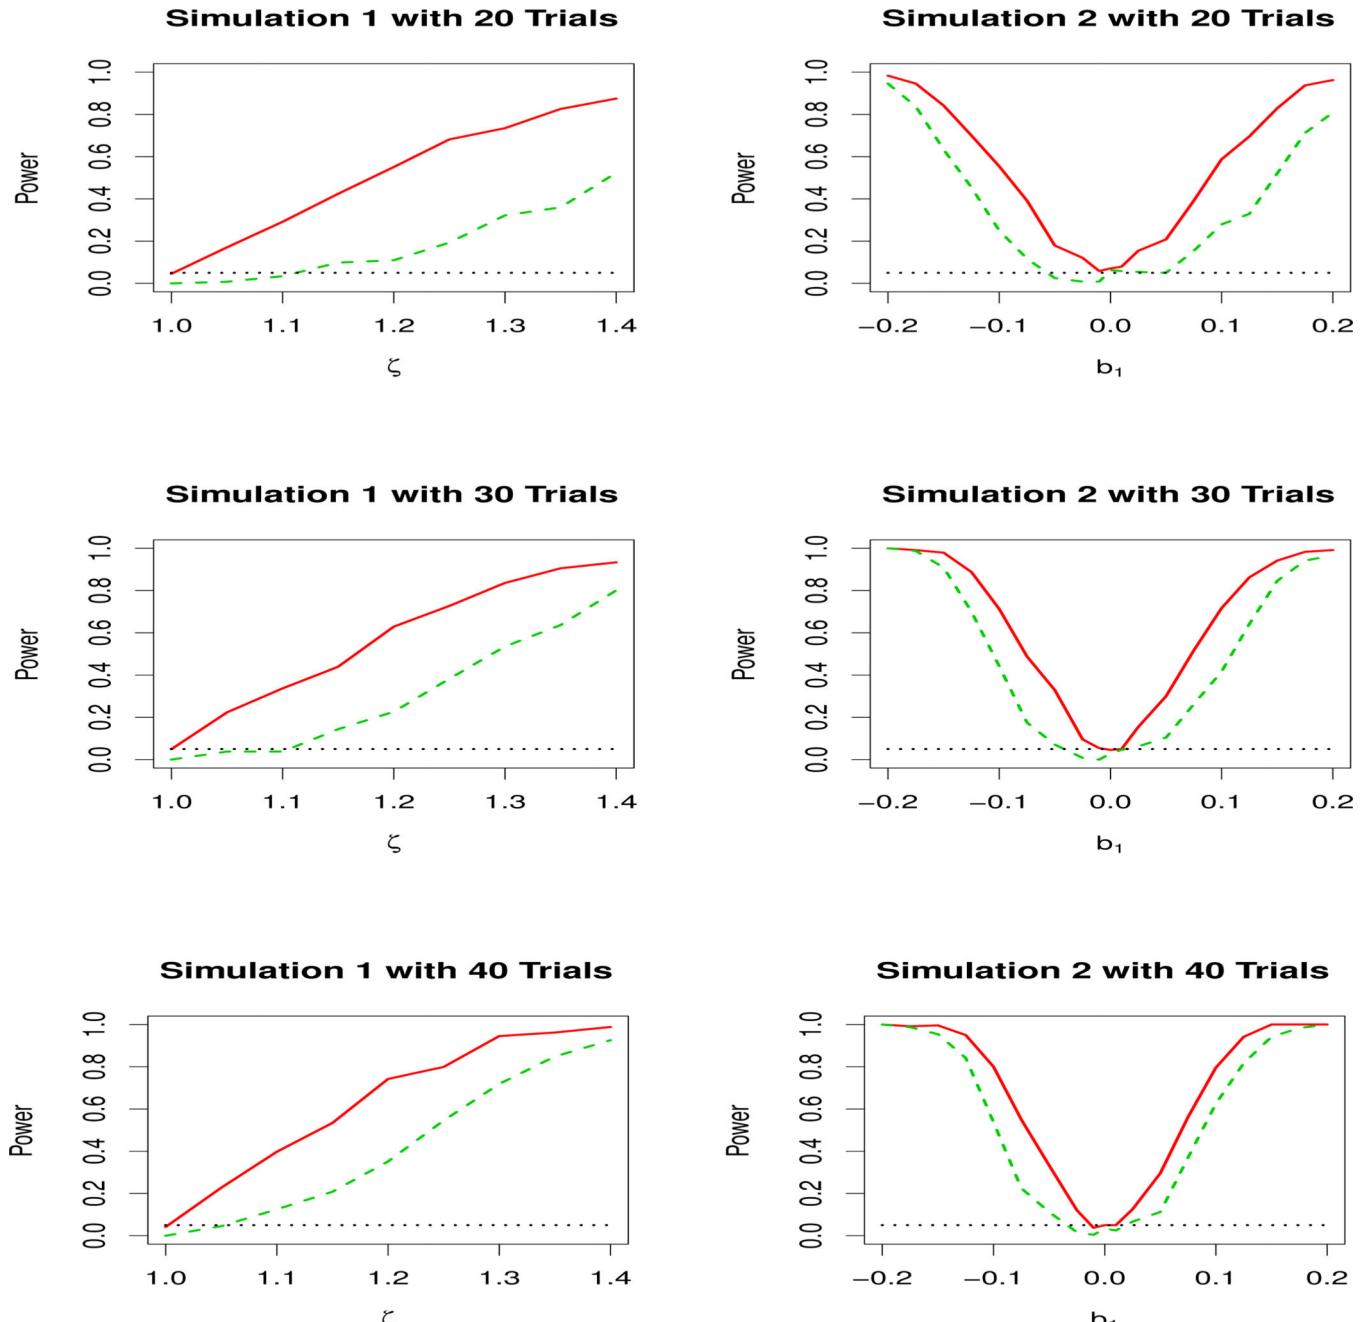


Figure 5. Two dependent neurons in lagged synchrony

The lag values are set to 3, 4, or 5 with probabilities 0.2, 0.5, and 0.3 respectively. The left panel shows the frequency of spikes over time. The right panel shows the posterior distribution of ζ and L . Darker cells represent higher frequencies.

**Figure 6. Power analysis**

Comparing our proposed method (solid curves) to the method of Kass et al. (2011) (dashed curves) based on statistical power using two simulation studies. Here the dotted lines indicate the 0.05 level.

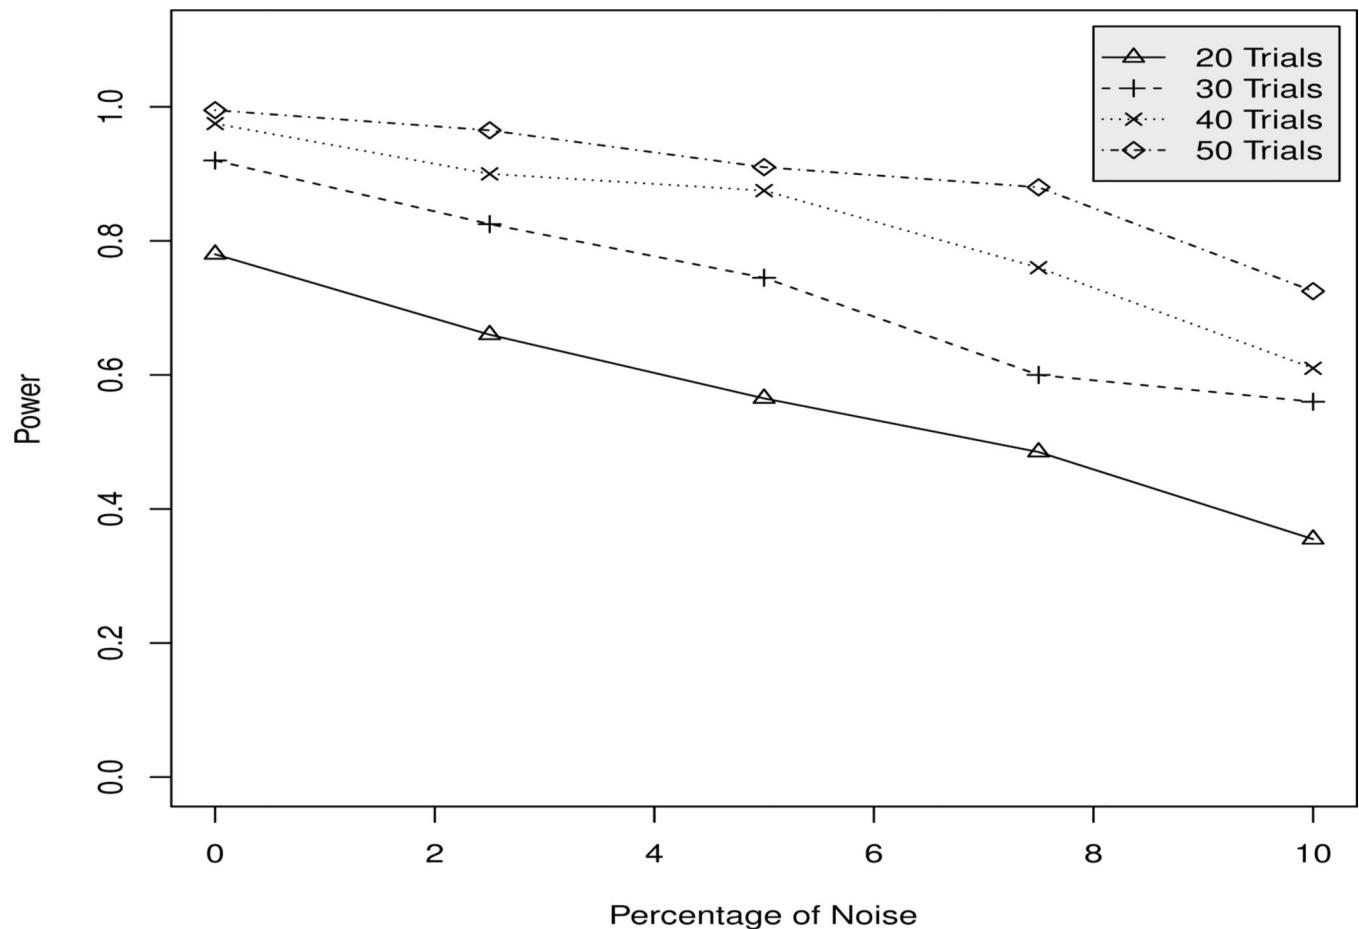
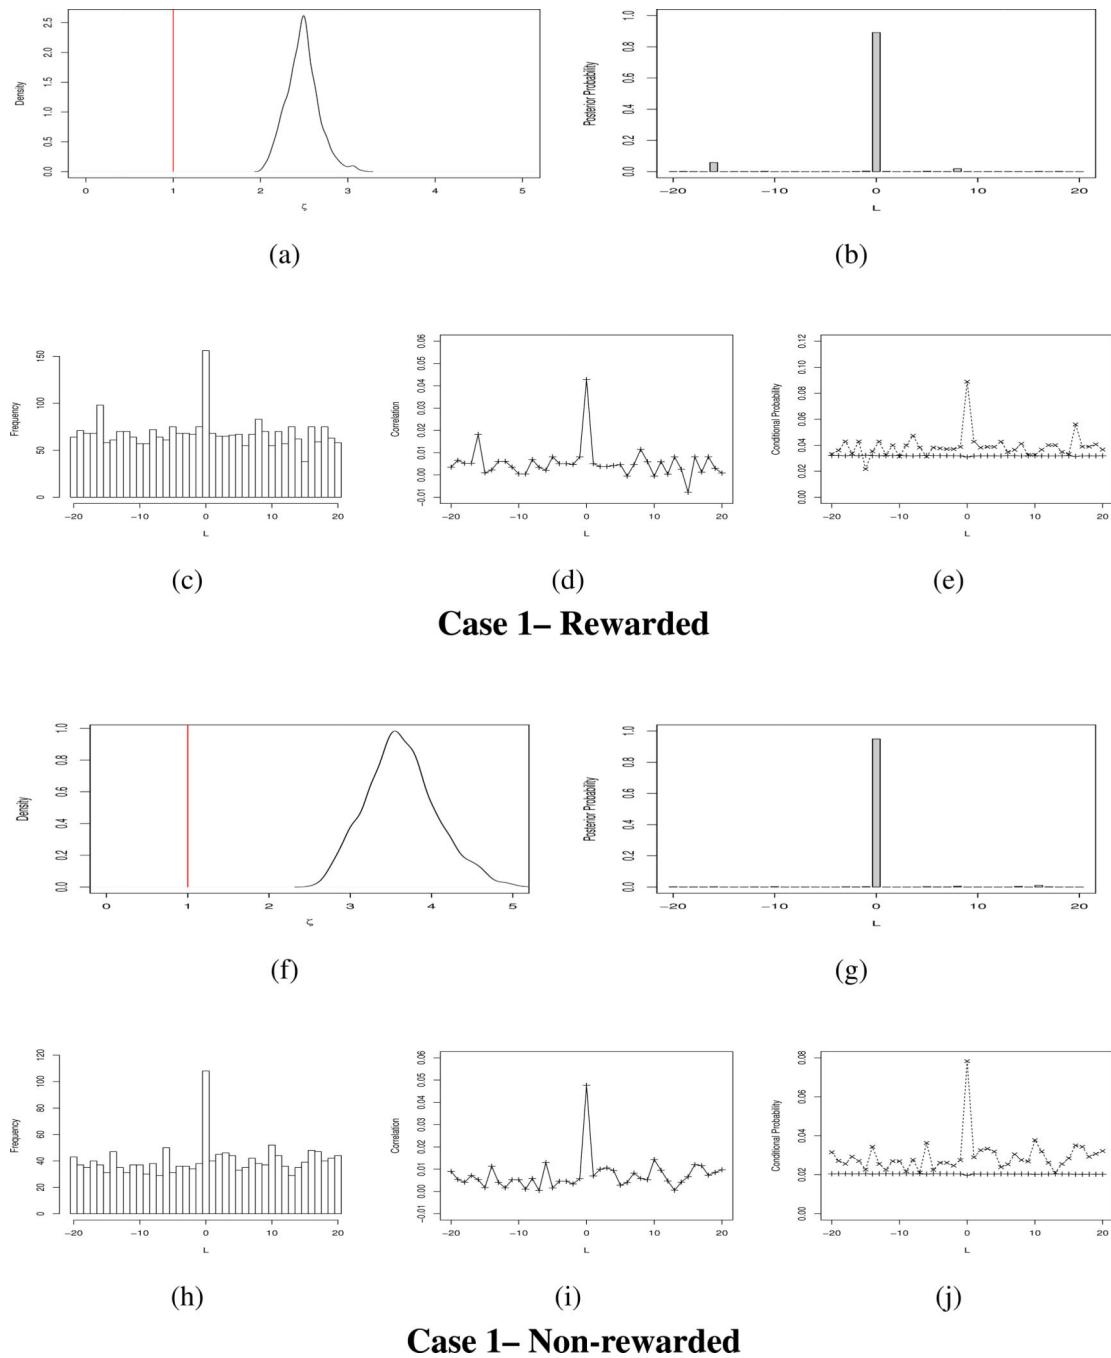
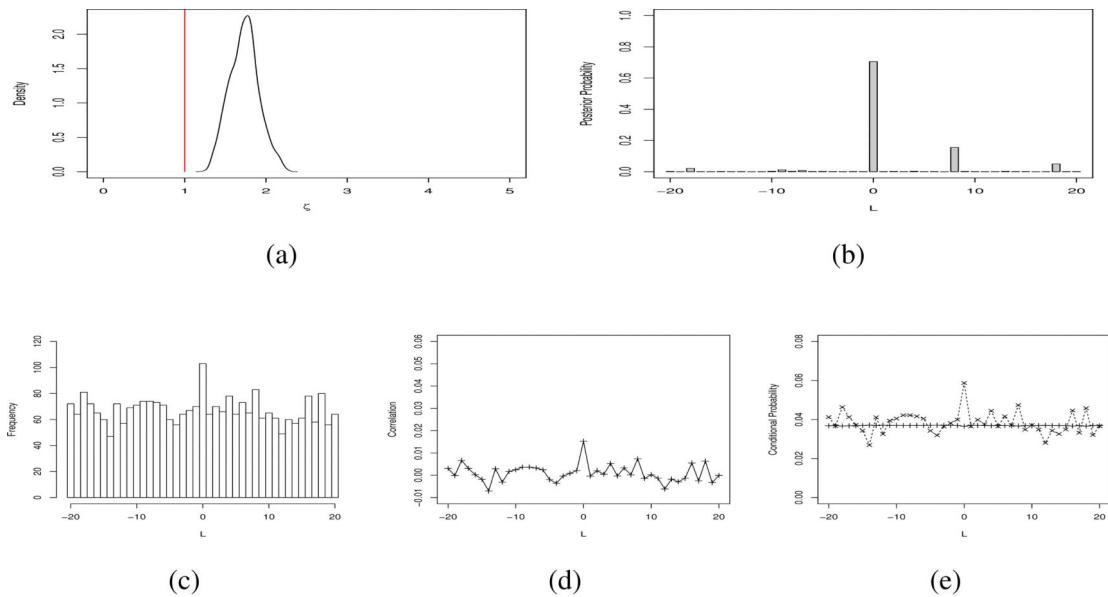


Figure 7. Sensitive analysis for trial-to-trial variability

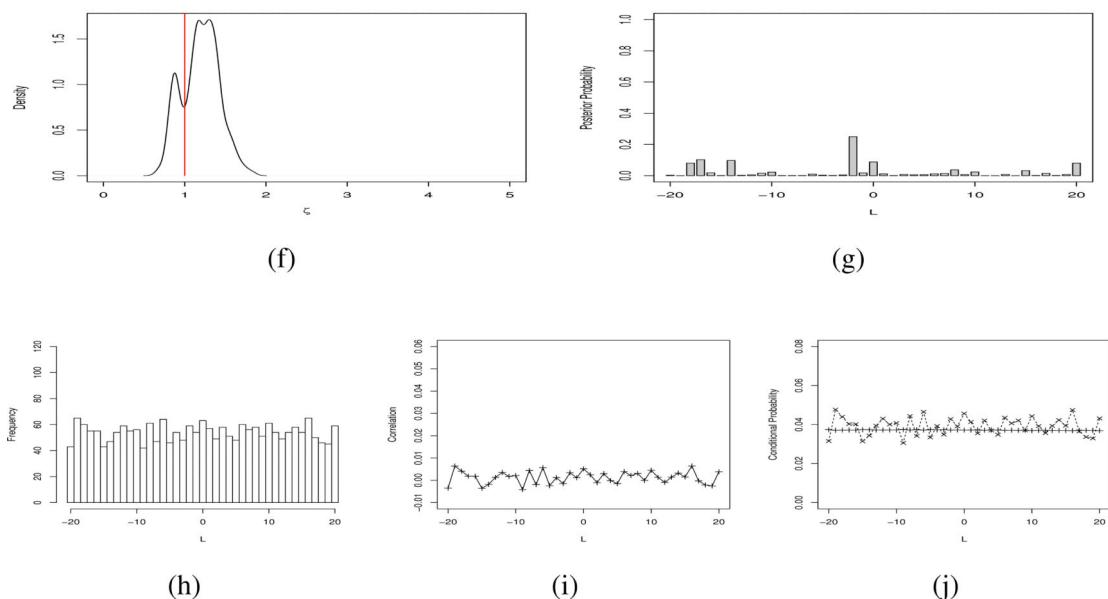
Comparing power for varying number of trials and noise rate (i.e., fraction of time bins in a trial flipped from zero to one or from one to zero).

**Figure 8. Case 1**

a) Posterior distribution of ζ , b) posterior distribution of lag, c) co-firing frequencies, d) correlation coefficients, and e) estimated conditional probabilities of firing for the second neuron given the firing status (0: solid line, 1: dashed line) of the first neuron over different lag values for the rewarded scenario; (f)–(j) are the corresponding plots for the non-rewarded scenario.



Case 2- Rewarded



Case 2- Non-rewarded

Figure 9. Case 2

a) Posterior distribution of ζ , b) posterior distribution of lag, c) co-firing frequencies, d) correlation coefficients, and e) estimated conditional probabilities of firing for the second neuron given the firing status (0: solid line, 1: dashed line) of the first neuron over different lag values for the rewarded scenario; f)–j) are the corresponding plots for the non-rewarded scenario.

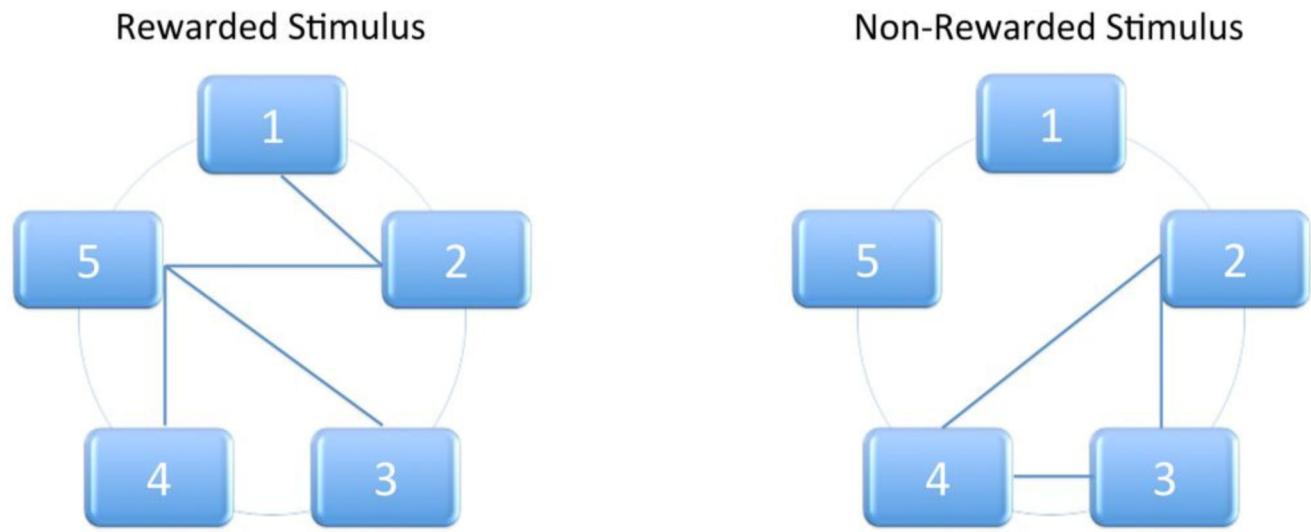


Figure 10.

A schematic representation of connections between five neurons under two experimental conditions. The solid line indicates significant association.

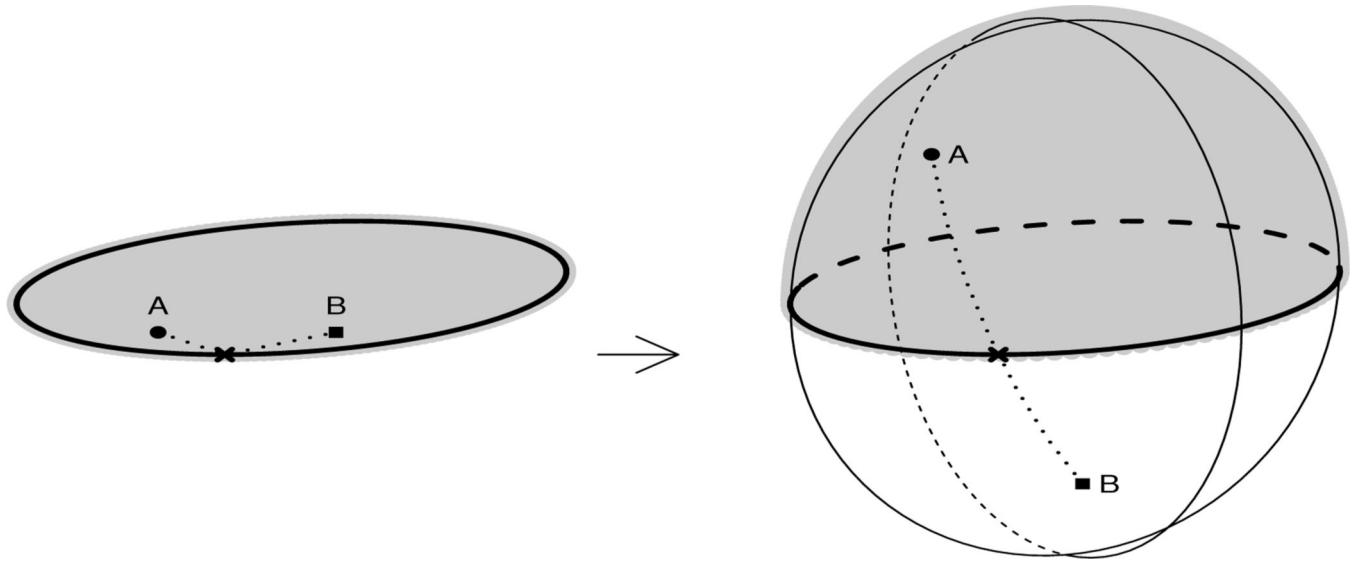


Figure 11.

Transforming unit ball $B_0^D(1)$ to sphere S^D .

Table 1

Estimates of β 's along with their 95% posterior probability intervals for simulated data based on our copula-based model. Here, row i and column j shows the estimate of β_{ij} , which captures the relationship between the i^{th} and j^{th} neurons.

β	2	3
1	0.66 (0.30,0.94)	0.02 (-0.26,0.27)
2		-0.05 (-0.33,0.19)

Table 2

Estimates of β 's along with their 95% probability intervals for the first scenario (Rewarded) based on our copula model.

β	2	3	4	5
1	0.22(0.07,0.39)	0.00(-0.07,0.04)	0.03(-0.02,0.15)	0.01(-0.04,0.08)
2		0.03(-0.02,0.18)	0.06(-0.02,0.22)	0.07(0.00,0.25)
3			0.08(-0.01,0.26)	0.21(0.04,0.38)
4				0.23(0.09,0.40)

Table 3

Estimates of β 's along with their 95% probability intervals for the second scenario (Non-rewarded) based on our copula model.

β	2	3	4	5
1	0.05(-0.02,0.25)	-0.01(-0.09,0.04)	0.15(-0.01,0.37)	0.05(-0.03,0.22)
2		0.21(0.03,0.41)	0.18(0.00,0.37)	0.03(-0.02,0.19)
3			0.17(0.00,0.34)	0.03(-0.02,0.19)
4				0.07(-0.01,0.24)