Trial-to-Trial Variability and its Effect on Time-Varying Dependence Between Two Neurons

Can Cai, Robert E. Kass, and Valérie Ventura

Department of Statistics

and

Center for the Neural Basis of Cognition
Carnegie Mellon University

Running Head: **Synchrony and trial to trial variability**

ABSTRACT

The joint peristimulus time histogram (JPSTH) and cross-correlogram provide a visual representation of correlated activity for a pair of neurons, and the way this activity may increase or decrease over time. In a companion paper (Cai *et al.* 2004a) we showed how a Bootstrap evaluation of the peaks in the smoothed diagonals of the JPSTH may be used to establish the likely validity of apparent time-varying correlation. As noted by Brody (1999a,b) and Ben-Shaul *et al.* (2001), trial-to-trial variation can confound correlation and synchrony effects. In this paper we elaborate on that observation, and present a method of estimating the time-dependent trial-to-trial variation in spike trains that may exceed the natural variation displayed by Poisson and non-Poisson point processes. The statistical problem is somewhat subtle because relatively few spikes per trial are available for estimating a firing-rate function that fluctuates over time. The method developed here uses principal components of the trial-to-trial variability in firing rate functions to obtain a small number of parameters (typically two or three) that characterize the deviation of each trial's firing rate function from the across-trial average firing rate, represented by the smoothed PSTH. The Bootstrap significance test of Cai *et al.* (2004a) is then modified to accommodate these general excitability effects. This methodology allows an investigator to assess whether excitability effects are constant or time-varying, and whether they are shared by two neurons. It is shown that trial-to-trial variation can, in the absence of synchrony, lead to an increase in correlation in spike counts between two neurons as the length of the interval over which spike counts are computed is increased. In data from two V1 neurons we find that highly statistically significant evidence of dependence disappears after adjustment for time-varying trial-to-trial variation.

Keywords: Bootstrap, Correlation, Cross-correlogram, Excitability Effects, Firing Rate, Latency Effects, Non-Poisson Spiking, Peri-Stimulus Time Histogram, Point Process, Principal Components, Smoothing, Spike Train Analysis, Trial-to-Trial Variability

1 Introduction

Spike trains recorded from behaving animals display variation in spike timing both within and across repeated trials. In some cases, variation within and across repeated trials is consistent with the random variation ("noise variation") observed in repeated sequences of events that follow Poisson or non-Poisson point process models (see Barbieri *et al.* 2001, Johnson 1996, Kass and Ventura 2001, and the references therein). In many contexts, however, the conditions of the experiment or the subject may vary across repeated trials enough to produce discernible trial-to-trial spike train variation beyond that predicted by Poisson or other point processes. Such trial-to-trial variation may be of interest not only for its physiological significance (Azouz and Gray 1999; Hanes and Schall 1996) but also because of its effects on statistical procedures. In particular, as observed by Brody (1999a,b), Ben-Shaul *et al.* (2001), and Grün *et al.* (2003), trial-to-trial variation can affect the assessment of correlated firing in a pair of simultaneously-recorded neurons. In this paper we present a statistical procedure for testing and estimating trial-to-trial variation in time-varying firing rate, and apply it to the problem of assessing time-varying dependence between spike trains from two neurons. We also demonstrate that trial-to-trial variation can lead to an increase in correlation in spike counts between two neurons as the length of the interval over which spike counts are computed is increased (as appears, for example, in Figure 2B in Reich *et al.* 2001).

One aspect of trial-to-trial variation is the tendency for neuronal response to shift in time. That is, a neuron may tend to fire earlier or later on some trials than on others, so that re-alignment of trials becomes desirable (Baker and Gerstein 2001; Ventura 2004). It is useful to distinguish such latency effects from variation in the amplitude of firing rate, which is sometimes called "excitability." It is possible to describe trial-to-trial amplitude variation by applying a kernel smoother to each trial's spike train (using a Gaussian filter or something similar; Nawrot *et al.* 1999). That method, however, implicitly treats the spike train as Poisson and it ignores a special feature of this situation: while there is

substantial information on the general shape of the firing-rate function obtained from the aggregated trials (i.e, the smoothed PSTH), there is relatively little information per trial from which to estimate a time-varying function, at least for smaller firing rates (e.g., less than 40 Hz). This creates an especially difficult problem for statistical inference procedures because standard errors and confidence intervals become very wide and statistical tests have little power. To improve estimation and inference, simple representations that take account of the aggregate pattern are needed. One such statistical model was discussed by Brody (1999a,b). He took the firing rate on each trial to be the sum of two components: a background constant firing rate multiplied by a trial-dependent coefficient and a stimulus-induced time-varying function multiplied by a second trial-dependent coefficient. While likely to capture some dominant features of trial-to-trial variation, this model may be too simple for many situations: it requires an experimental period during which the neuron fires at a constant background rate, and it assumes a single multiplicative constant describes the fluctuation in stimulus-induced firing rate. Additional statistical issues concerned the identification of the end of the background period and beginning of the stimulus-induced period, the use of the raw PSTH rather than a smoothed version of it to estimate the stimulus-induced firing rate, and the assumption of Poisson spiking. The method described here avoids the use of a background period, allows two or more coefficients to describe the fluctuation in firing rate, and may be applied with non-Poisson spiking. It also may be used to assess whether excitability effects are shared by two or more neurons.

2 Methods

The phenomenon we wish to describe is illustrated in Figure 1A, which displays five trials of spike trains from a simulated neuron. In this case the spike trains are all assumed to follow inhomogeneous Poisson processes, but the firing-rate intensity functions (shown

on the right-hand side of the figure) are distinct for the different trials. The firing-rate intensity function for the rth trial may be written $\lambda_r(t)$, so that the probability the neuron will fire in the interval $(t, t + dt)$ on trial r is $\lambda_r(t)dt$. In Figure 1B, the five functions $\lambda_r(t)$, are related to the firing-rate intensity averaged across trials, $\lambda(t)$, according to

$$\lambda_r(t) = g_r(t)\lambda(t) \tag{1}$$

where $g_r(t)$ represents the deviation in time-varying amplitude from the common rate on trial r. Thus, the variability of the spike trains in Figure 1A is due partly to the stochastic variation of Poisson processes and partly to the extra variation caused by the factor $g_r(t)$, which would represent trial-dependent changes in the state of the subject (due to such things as attention or motivation effects).

In Sections 2.1 and 2.2 we address the statistical problem of recovering the factors $g_r(t)$. We will use a somewhat more elaborate framework, but to explain the approach let us begin with (1) and assume spike times are recorded at a resolution, such as 1 ms, such that it is acceptable to represent them as a binary sequence across very small time intervals. We write the time resolution as Δt, the sequence for trial r (with $r = 1, \ldots, R$), as $X_r(t_1), X_r(t_2), \ldots, X_r(t_{max})$, (i.e., $X_r(t) = 1$ if there is a spike at time t and $X_r(t) = 0$ otherwise), and also set

$$p_r(t) = P(X_r(t) = 1) = \lambda(t)g_r(t)\Delta t. \tag{2}$$

If instead of binary data we were using quantitative data treated using regression methods, to fit such a deviation it would be reasonable to examine the residual obtained after first removing $\lambda(t)$. Therefore, we first fit $\lambda(t)$ and then use a binary-data analogue to fitting the residuals. This is achieved by introducing the fit $\hat{\lambda}(t)$ (a smoothed version of the PSTH) as an "offset." In the terminology of generalized linear modeling, an offset is an explanatory variable whose coefficient is set equal to 1. The offset replaces the intercept in a linear model or generalized linear model (see, for example, McCullagh and Nelder,

1990). We thus fit the statistical model

$$\begin{aligned} X_r(t) &\sim Ber(p_r(t)) \\ \log p_r(t) &= \log \hat{\lambda}(t) + f_r(t) \end{aligned} \tag{3}$$

where $Ber(p)$ stands for Bernoulli with probability p, i.e., $X_r(t) = 1$ with probability p and is 0 otherwise. Here, $f_r(t)$ will capture the deviation of $\log p_r(t)$ from $\log \hat{\lambda}(t)$.

The necessary elaborations of this simple idea are as follows. First, because the spike train for a single trial is relatively sparse, it is important to reduce the dimensionality of the functional representation. As explained in detail in Section 2.1, we reduce dimensionality in two ways. We begin by fitting (3) with splines having a small number of knots (e.g., 1 to 3 knots). We then introduce a principal component decomposition of the functional variation. This serves in part to reduce dimensionality further, and in part to provide additional interpretation of the variability. The second elaboration, provided in Section 2.2, is to extend the method to non-Poisson data. In Section 2.3 we consider the joint spiking activity of a pair of neurons and show how to adjust the time-varying dependence functions $\zeta_\delta(t)$ used by Cai *et al.* (2004a), and also the cross-correlogram, for trial-to-trial variation. In Section 2.4 we extend the Bootstrap excursion test of Cai *et al.* (2004a) to the case where there exists trial-to-trial variation. Finally, in Section 2.5 we show how to assess whether the trial-to-trial variation is shared across two neurons; the extension to more than two neurons is immediate.

2.1 Time-varying Excitability

For inhomogeneous Poisson processes the trial-to-trial fluctuation in firing-rate intensity will be assumed to have the following form:

$$\lambda_r(t) = g_r(t - \tau_r)\lambda(t - \tau_r) \tag{4}$$

$$\log g_r(t) = w_{0r} + \sum_{j=1}^{J} w_{jr}\phi_j(t) \tag{5}$$

where τ_r is a shift (latency) parameter for trial r and the $\phi_j(t)$ functions are suitably chosen curves. We have introduced the shift parameters τ_r for completeness, due to the availability of methods for estimating these parameters (see Ventura 2004) but otherwise (4) is the same as (1). By restricting the form and number of the $\phi_j(t)$ functions we obtain an interpretable low-dimensional representation of the trial-to-trial variation, in which the coefficients w_{jr} may be estimated from the limited data (relatively small number of spikes) available per trial. As discussed below, the $\phi_j(t)$ functions will be taken to be principal components of the trial-to-trial variation. A special case of (5) is the *constant excitability* model, in which $\log g_r(t) = w_{0r}$; this case was illustrated in Figure 1.

Using the statistical model specified by equations (4) and (5), estimation may be carried out in several steps: first, as already indicated above, we estimate the common rate $\lambda(t)$ and, second, estimate the deviations $g_r(t)$. The third step is to identify a suitable low-dimensional representation of these deviations, and finally in the fourth step each trial's data is fitted as a function of this low dimensional representation. We do not discuss here the estimates of latency effects, which may be obtained by the procedure of Ventura (2004) as a preliminary step.

To obtain the smoothed estimate $\hat{\lambda}(t)$ of $\lambda(t)$ we prefer a spline-based method called BARS (DiMatteo *et al.* 2001; Kass *et al.* 2003) for this purpose, but many alternative smoothing techniques would work well for typical data sets. This accomplishes the first step in the procedure. The second step is to apply binary nonparametric regression to the data from each trial, using (3). Note, however, that in using (4) and (5), $f_r(t)$ is not the same as $\log g_r(t)$ because in (3) the principal component representation in equation (5) has not yet been imposed.

There are again many possible ways to perform the nonparametric regression in (3), but for the sake of statistical efficiency the number of parameters (the effective degrees of freedom) defining $f_r(t)$ should be kept relatively small. We therefore have in our applications assumed $f_r(t)$ may be adequately represented as a natural cubic spline with a

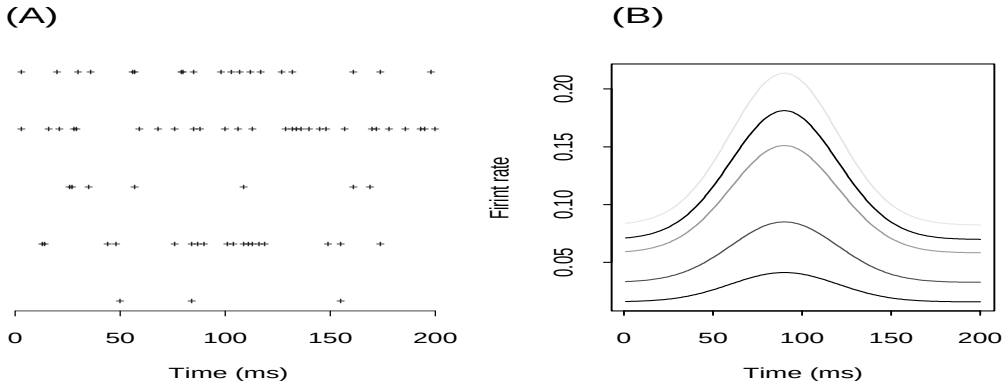

Figure 1: *Simulated spike trains with time-varying excitability effects. (A) The raster plot of 5 trials of a simulated neuron having an excitability effect described by model (4) and (5). (B) The firing rate functions $\lambda_r(t)$ corresponding to the trials in (A). Note that both the magnitude and shape of the firing rate function vary from trial to trial.*

very small number of knots. Fits $\hat{f}_r(t)$ again may be obtained from BARS, which attempts to find optimal locations for the knots. In the applications reported here, however, we have simplified by using a single knot and fixing its location. The fitting of $\hat{f}_r(t)$ is then obtained from generalized linear model software, available in most commercial statistical analysis packages. These regression and smoothing methods are discussed in many sources (e.g., McCullagh and Nelder 1990, and Hastie and Tibshirani 1990).

These first two steps have obtained a fit $\hat{\lambda}(t)$ for the trial-averaged intensity and a collection of fitted deviations $\hat{f}_r(t)$ for each trial. The third and fourth steps replace $\hat{f}_r(t)$ with $\log \hat{g}_r(t)$, thereby representing the variability in terms of the coefficients of principal components. The third step is to compute the covariance matrix $\hat{\Sigma}_f$ of the T-dimensional vectors $(\hat{f}_r(t_1), \ldots, \hat{f}_r(t_{max}))$, where t_{max} is the total number of time points, and to obtain the first few principal components of $\hat{\Sigma}_f$. These principal components are taken to be the functions $\phi_j(t)$ in (5). As we indicated earlier, only a few of these principal components are likely to be useful. A natural cubic spline with k knots has $k + 2$ free parameters.

Therefore, in principle, at most $k+2$ eigenvalues should be nonzero and at most $k+2$ principal components will be meaningful.

The fourth step is to determine J and obtain the coefficients $\{w_{jr}\}$ in (5). This will effectively project each trial's deviation from the PSTH onto the space of principal components. To carry this out, a binary regression is performed as in step 2, except that in equation (5) $f_r(t)$ is replaced by $w_{0r} + \sum_{j=1}^{J} w_{jr}\phi_j(t)$. Assuming there are J^* principal components, there are J^*+1 coefficients w_{jr} including the constant w_{0r}. To determine whether the contribution of some of the principal components is negligible, J^*+2 distinct regression models should be considered: the regression model involving only the offset $\hat{\lambda}(t)$, that involving the offset and the trial-dependent constants $e^{w_{0r}}$, that involving the offset, the constants, and the first principal component, etc., up to the model involving the onset, the constants, and all J^* principal components. These models can be fit sequentially and the deviance difference compared to a chi-squared distribution with R degrees of freedom, where R is the number of trials. Principal components may be included sequentially until the deviance difference is no longer statistically significant.

This four-step procedure will determine an expression for the trial-to-trial variability of the form (5). The resulting set of weights $w_r = (w_{0r}, w_{1r}, w_{2r}, \ldots, w_{Jr})$ describes the variation specific to trial r.

2.2 Non-Poisson Data

For general point process data that are not necessarily Poisson the intensity must be replaced by a conditional intensity, $\lambda_r(t \mid \mathcal{H}_r)$ where $\mathcal{H}_r$ is the spiking history of trial r up to time t. That is, the probability the neuron will fire in the interval $(t, t+dt)$ on trial r is $\lambda_r(t \mid \mathcal{H}_r)dt$, and this probability depends on the spike times prior to time t (e.g., Barbieri *et al.*, 2001; Daley and Vere-Jones, 2003). In addition, the firing-rate intensity averaged across trials $\lambda(t)$ must be replaced by its history-dependent counterpart, which we write as

$\lambda(t \mid \mathcal{H}_r)$ (now omitting the subscript r on λ). This function is the conditional intensity that would occur on trial r if there were no excess trial-to-trial variation. Following Brillinger (1988), Kass and Ventura (2001) have shown how standard software may be used to fit a particular class of non-Poisson models they called inhomogeneous Markov interval (IMI) models. These models assume

$$\lambda(t \mid \mathcal{H}_r) = \lambda(t, s_r^*(t)) \tag{6}$$

where $s_r^*(t)$ is the time of the last spike previous to time t on trial r. IMI models can also be extended to include spike times going further into the past (Kass and Ventura 2001). In practice, the function $\lambda(t \mid \mathcal{H}_r)$ specified by equation (6) is fitted using all available trials, but the dependence on the trial is of a highly specific form: the only way the firing-rates at time t differ on different trials is from the timing of the last spike previous to time t. Similarly, the notation $\lambda(t \mid \mathcal{H}_r)$ without a subscript r on λ is intended to denote dependence on trial only through the spiking history $\mathcal{H}_r$.

With this replacement the statistical model for trial-to-trial variability now must specify the relationship between $\lambda_r(t \mid \mathcal{H}_r)$ and $\lambda(t \mid \mathcal{H}_r)$. The model (4) becomes

$$\lambda_r(t \mid \mathcal{H}_r) = g_r(t - \tau_r)\lambda(t - \tau_r \mid \mathcal{H}_r) \tag{7}$$

and $\lambda(t \mid \mathcal{H}_r)$ must be substituted for $\lambda(t)$ in Equation (2), so that the fitting of the nonparametric regression model (3) is altered. By combining models (7) and (6) we have

$$\lambda_r(t \mid \mathcal{H}_r) = g_r(t)\lambda(t, s_r^*(t))$$

so that generalized linear modeling software may be used (as in Kass and Ventura 2001) and the four-step procedure presented in Section 2.1 remains straightforward.

2.3 Adjustment of Joint Spiking Activity for Trial-to-Trial Variability

Suppose two neurons are recorded simultaneously across multiple trials, and the task is to assess the time-varying dependence between their spike trains. Writing the joint firing rate for neuron 1 at time t and neuron 2 at time $t+\delta$ as $\lambda^{12}(t, t+\delta)$, and the individual firing rate functions as $\lambda^i(t)$ for $i = 1, 2$, using the definition

$$\zeta_\delta(t) = \frac{\lambda^{12}(t, t+\delta)}{\lambda^1(t)\lambda^2(t+\delta)}, \tag{8}$$

Cai *et al.* (2004a) presented and studied a Bootstrap significance test for time-varying dependence in the absence of trial-to-trial variability. The form (8) applies to Poisson data. It is a special case of the more general form

$$\zeta_\delta(t) = \frac{\lambda^{12}(t, t+\delta \mid \mathcal{H})}{\lambda^1(t \mid \mathcal{H}^1)\lambda^2(t+\delta \mid \mathcal{H}^2)}, \tag{9}$$

where $\mathcal{H}^1$ and $\mathcal{H}^2$ are the spiking history of neuron 1 and 2 up to time t and $t+\delta$ respectively, and $\mathcal{H}$ is the combined history, may be applied to non-Poisson data. Cai, *et al.* (2004a) showed how the Bootstrap significance test could be applied to data that follow non-Poisson models.

To take account of trial-to-trial variability, the formulas (4) and (5) must be used to replace $\lambda^i(t)$ with its trial-dependent counterpart $\lambda^i_r(t)$, for $i = 1, 2$, in (8) when applying the Bootstrap procedure to Poisson data, and the sampling scheme must be modified accordingly. Similarly (7) and (5) must be used in (9) to accommodate non-Poisson data. In the general case (which applies to both Poisson and non-Poisson data), we will assume

$$\lambda^{12}_r(t, t+\delta \mid \mathcal{H}_r) = \lambda^1_r(t \mid \mathcal{H}^1_r)\, \lambda^2_r(t+\delta \mid \mathcal{H}^2_r)\, \zeta_\delta(t), \tag{10}$$

that is, $\zeta_\delta(t)$, the time-varying excess joint firing of the two neurons, above that predicted by independence after allowing for trial-to-trial variation, does not itself depend on the trial. To estimate $\zeta_\delta(t)$, we begin with the joint spiking data. Let us set $X^{12}_{\delta,r}(t) = 1$ if on

trial r neuron 1 spikes at time t and neuron 2 spikes at time $t+\delta$, so that we obtain a binary sequence $X^{12}_{\delta,r}(t_1),\ldots,X^{12}_{\delta,r}(t_{max,\delta})$ for each δ and r. For a given δ we use the data from all R trials to fit the nonparametric regression model

$$\begin{aligned} X^{12}_{\delta,r}(t) &\sim Ber\left(\lambda^{12}_r(t,t+\delta \mid \mathcal{H}_r)\right) \\ \log\lambda^{12}_r(t,t+\delta \mid \mathcal{H}_r) &= \log\left(\hat{\lambda}^1_r(t \mid \mathcal{H}^1_r)\cdot\hat{\lambda}^2_r(t+\delta \mid \mathcal{H}^2_r)\right)+\log\zeta_\delta(t) \end{aligned} \tag{11}$$

which is similar to the model in (3) except that now $\zeta_\delta(t)$ does not vary with the trial (as did $f_r(t)$ in (3)). That is, the trial-to-trial variation is already captured in the factors $\hat{\lambda}^1_r(t \mid \mathcal{H}^1_r)$ and $\hat{\lambda}^2_r(t+\delta \mid \mathcal{H}^2_r)$, which are here treated as offsets in the generalized regression. The result of this fitting process is an estimate $\hat{\zeta}_\delta(t)$.

Once we have estimates of the trial-specific firing rates $\lambda^i_r(t \mid \mathcal{H}^i_r)$ it is also straightforward to adjust the JPSTH and cross-correlogram to remove the general excitability effects. Let $Y_{\delta,r}(t)$ be 1 if neuron 1 fires at time t and neuron 2 fires at time $t+\delta$, on trial r, and 0 otherwise. Then the corrected JPSTH at times t and $t+\delta$ is

$$J_C(t,t+\delta)=\frac{1}{R}\sum_r\left(Y_{\delta,r}(t)-\hat{\lambda}^1_r(t \mid \mathcal{H}^1_r)\cdot\hat{\lambda}^2_r(t \mid \mathcal{H}^2_r)\right) \tag{12}$$

and the corrected cross-correlogram at lag δ may be defined as

$$C(\delta)=\sum_t J_C(t,t+\delta). \tag{13}$$

Note that (13) does not include a divisor to produce an average value of $J_C(t,t+\delta)$, nor is a normalization used, as in Aertsen *et al.* (1989). These could be included without substantially changing the methodology.

2.4 Bootstrap Excursion Test

As in Cai, *et al.* (2004a), the Bootstrap significance test uses null bands, obtained under the null hypothesis of independence. The following algorithm generates null bands:

1. Obtain fits of the firing rates $\hat{\lambda}_r^i(t)$, for $i = 1, 2$ and for all trials $r = 1, \ldots, R$.

2. Sample at random and with replacement R numbers from the integers $1, \ldots, R$, to identify a set of R trials that will be used to create the bands.

3. Simulate pairs of spike trains with firing-rate functions $\hat{\lambda}_r^1(t)$ and $\hat{\lambda}_r^2(t)$, where r runs through the set of trials sampled in step 2.

4. For the sample of R spike trains obtained in step 3, obtain new fits of the firing rates. For the sampled trial k, the new fitting firing rates may be written $\hat{\lambda}_{(k)}^i(t)$, for $i = 1, 2$ to distinguish it from $\hat{\lambda}_k^i(t)$ obtained in step 1. The procedure used to obtain $\hat{\lambda}_{(k)}^i(t)$ from sampled trial k in step 3 is the same as that used to obtain $\hat{\lambda}_k^i(t)$ from the actual data.

5. Obtain a fitted function $\hat{\zeta}_\delta(t)$ by fitting (11) for the simulated sample.

6. Repeat step 2-5 N times to get N estimates $\hat{\zeta}_\delta(t)$.

7. Calculate .025 and .975 quantiles for each t to obtain 95% null bands for $\zeta_\delta(t)$.

The remainder of the Bootstrap testing procedure is the same as the original procedure of Cai *et al.* (2004a): we compute

$$G_{obs} = max|\int_{t_s}^{t_e}(\hat{\zeta}_\delta(t) - 1)dt|, \tag{14}$$

where t_s and t_e are the starting and ending time of the period during which $\hat{\zeta}_\delta(t)$ is outside the null bands, and the maximum is taken across all t_s and t_e. Then G_{obs} is the largest area of $\hat{\zeta}_\delta(t)$ exceeding the band. To calculate the p-value, let $\hat{\zeta}_\delta^{(n)}(t)$, $n = 1, \ldots, N$ stand for the estimate of $\zeta_\delta(t)$ obtained from from the nth bootstrap sample. For each Bootstrap sample we compute

$$G_{boot}^{(n)} = max|\int_{t_s}^{t_e}(\zeta_\delta^{(n)}(t) - 1)dt| \tag{15}$$

and then we calculate the p-value of G_{obs} as

$$p = \frac{\text{number of Bootstrap samples for which } \; G_{boot}^{(n)} > G_{obs}}{N+1}.$$

2.5 Trial-to-trial Variation Shared Across Neurons

In Section 2.3 the fitted firing-rate functions $\hat{\lambda}_r^1(t)$ and $\hat{\lambda}_r^2(t)$ were not restricted to have any relationship to each other. It is, however, plausible that the two neurons might either have the same trial-to-trial effects or proportional trial-to-trial effects. These possibilities may be also examined with standard generalized linear model methodology.

We will write the equality and proportionality cases as $g_r^1(t) = g_r^2(t)$ and $g_r^1(t) = \alpha \cdot g_r^2(t)$, where α is a scalar constant. A further special case is the constant excitability model $g_r^i(t) = w_{0r}$, for which equality and proportionality become $w_{0r}^1 = w_{0r}^2$ and $w_{0r}^1 = \alpha \cdot w_{0r}^2$. As explained in many texts on linear regression (the generalized linear model context being analogous), fitting these models is straightforward: in the case of equality, the data and explanatory variables for the two models are concatenated, and then the fitting proceeds as before; in the case of proportionality, indicator variables u^i, $i = 1, 2$, are introduced for which $u = 0$ when the data come from neuron 1 and $u_i = 1$ when the data come from neuron 2, and these multiply the relevant explanatory variables. In each case, the appropriate offset is the concatenation of the individual offsets used for the two neurons separately. Because the degrees of freedom equals the number of free parameters, separate effects contribute double the number of degrees of freedom that shared effects contribute. This is illustrated in Table 4.

Neuron	Excitability effect $g_r(t)$	Global rate $\lambda(t)$
A	$g_r(t) = 1$ for all t and all r	$\lambda(t) = 0.02 + 4f(t, 90, 20)$
B	$g_r(t) = e^{w_{0r}} = a_r \sim Gamma(0.5, 0.5)$	$\lambda(t) = 0.05 + 6f(t, 90, 30)$
C	$g_r(t) = 1 + c_r f(t, 100, 25)$ $c_r = b_r - \bar{b}$, $b_r \sim Gamma(1, 0.025)$	$\lambda(t) = 0.05 + 6f(t, 90, 30)$
D	$g_r(t) = 1 + 60c_r f(t, 150, 70)$ $c_r = b_r - \bar{b}$, $b_r \sim Gamma(1, 0.5)$	$\lambda(t, s) = \lambda_1(t)\lambda_2(s)$ $\lambda_1(t) = 0.05 + 7f(t, 135, 40)$ $\lambda_2(s) = f(s, 5.5, 1.5)$

Table 1: *Firing rates $\lambda_r(t)$ given by (1) for neurons A (no trial-to-trial variation), B (constant trial-to-trial variation), and C (time-varying trial-to-trial variation). $f(t, a, b)$ denotes the normal density function with mean a and standard deviation b. " $\sim$ " means "randomly generated from". The firing rates of all trials are shown in Figure 2.*

3 Results

In Sections 3.1 and 3.2 we illustrate our methods for fitting excitability effects based on simulated Poisson and non-Poisson spike trains. In Section 3.3, we illustrate how to correct synchrony detection plots and measures, and how to carry bootstrap inferences.

3.1 Simulated Data: Poisson Spike Trains

We simulate Poisson spike train data with firing rate $\lambda_r(t)$ given by (1) under three scenarios: no trial-to-trial variation (neuron A), constant trial-to-trial variation (neuron

Poisson neurons A, B, and C

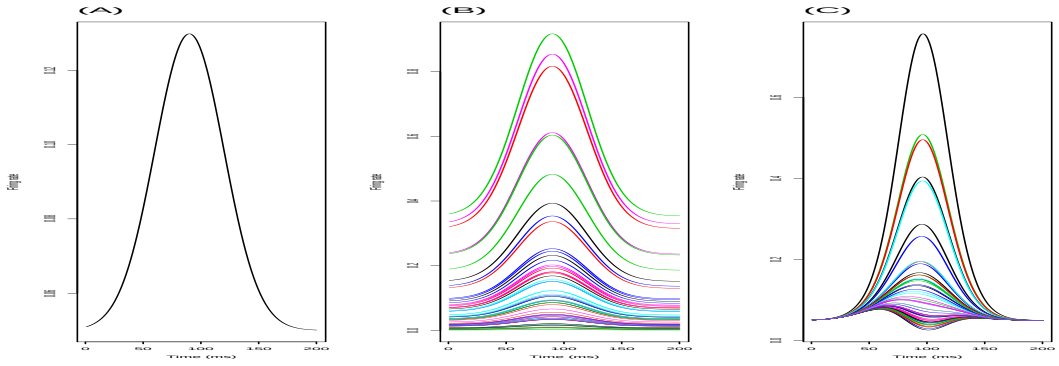

Non-Poisson neuron D

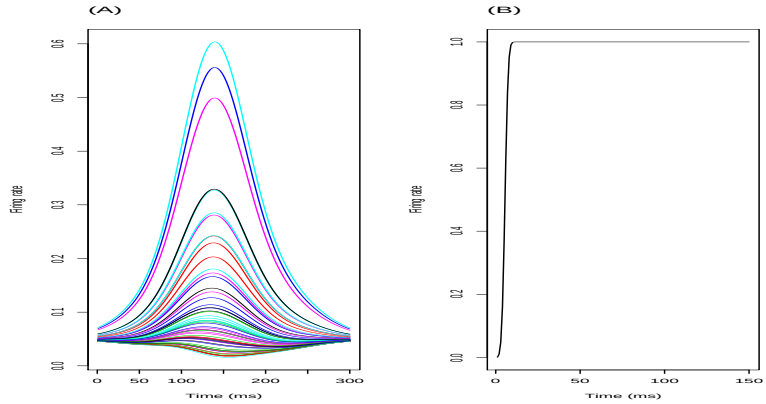

Figure 2: *(A), (B) and (C): true firing rates* $\lambda_r(t) = g_r(t)\lambda(t)$ *for each trials of simulated neurons A, B, and C described in Table 1. For neuron D, we show* $g_r(t)\lambda_1(t)$ *for all trials* $r = 1 \dots R$ *in (A), and* $\lambda_2(s)$ *in (B).*

B), and time-varying trial-to-trial variation (neuron C). Specifics of the firing rates are summarized in Table 1, and shown in Figure 2. All three simulated neurons have 60 trials of spike trains, and there are 200 time bins for each trial.

Table 2 lists the deviance of the fits to the simulated data for Neurons A-C. For neuron A using, first, only the offset (i.e., the smoothed PSTH $\hat{\lambda}(t)$, setting $f_r(t) = 0$, in (3)) and, second, both the offset and the trial-dependent constants $e^{w_{0r}}$. The deviance difference $(1525 - 1482 = 43)$ is compared to a χ^2 distribution with degrees of freedom equal to the difference between the degrees of freedom in the two models $(60 - 0 = 60)$. Because the p-value is much larger than .05 we would conclude, correctly, that no trial-to-trial variation is present for Neuron A. For Neuron B the results indicate extremely strong evidence for trial-to-trial variability ($p < 10^{-5}$), but no evidence for time-dependent trial-to-trial variability ($p = .59 > .05$), and we would therefore correctly conclude that there is only constant trial-to-trial variability for Neuron B. Figure 3 displays the true and the fitted firing rate functions for a representative subset of trials, graphically demonstrating the good fit of the constant excitability model for Neuron B. The fitted rates for all 60 trials can be found in the appendix, Figure 9. Note that the no trial-to-trial variation model would have fitted the same firing rate for all trials. For neuron C the results indicate extremely strong evidence for time-dependent trial-to-trial variability that is captured by a single principal component ($p < 10^{-5}$) but no evidence that a second principal component is required to describe this variability ($p > .05$). Figure 3 displays, for a representative subset of trials, the true firing rate functions, the incorrect constant-excitability fits, and the correct time-varying fits. The fitted rates for all 60 trials can be found in the appendix, Figure 10.

Neuron B

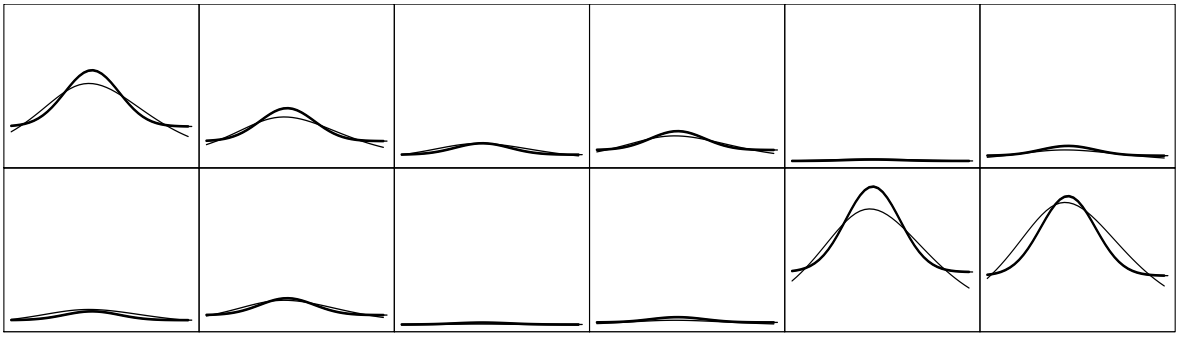

Neuron C

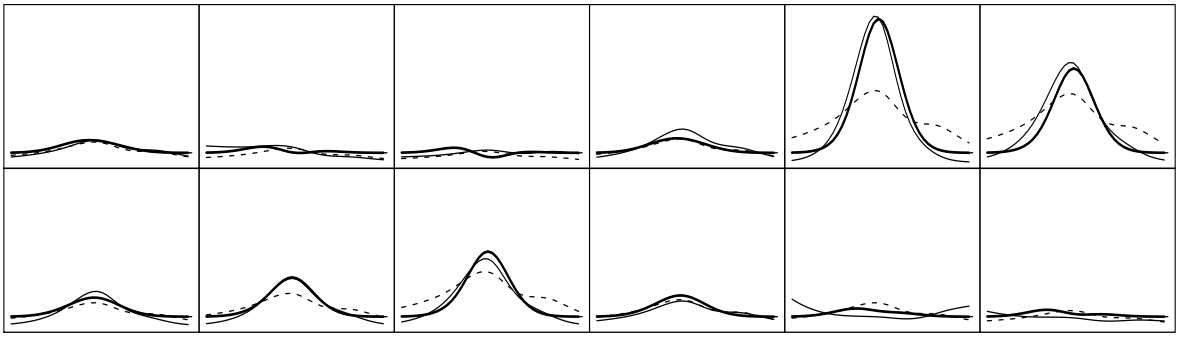

Neuron D

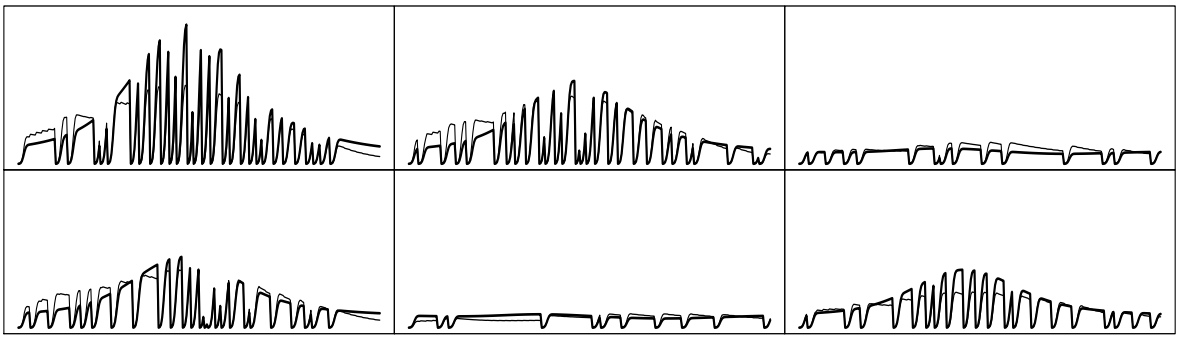

Figure 3: *True and fitted firing rates for Neurons B, C and D for a representative sample of trials. For each neuron and trial, the bold and thin lines are, respectively, the true and fitted firing rates. For Neuron C, we also plotted (dashed curves) the fits of the constant excitability model, which was rejected by the deviance test in Table 2. The fits for all trials can be found in the appendix.*

3.2 Simulated Data: Non-Poisson Spike Trains

Data were also simulated from a non-Poisson neuron (Neuron D) that followed an IMI model as in (6), with the multiplicative form $\lambda(t, s_*(t)) = \beta_1(t)\beta_2(s_*(t))$. Neuron D had 32 trials and 300 time bins. Figure 2 shows the functions $g_r(t)\beta_1(t)$ for all 32 trials, and $\lambda_2(s_*(t))$; specifics are in Table 1.

The four-step estimation procedure was followed (as in the Poisson case), with deviances and p-values recorded in Table 2, and the resulting fitted firing rate functions for six trials are displayed in Figure 3. The fits follow the true firing rate functions on each trial quite well. The fitted rates for all 32 trials can be found in the appendix, Figure 11.

3.3 Simulated Data: Adjustment of Dependence Assessment

We simulated data from pairs of neurons in two situations. In the first scenario the two neurons (Neurons E1 and E2) had independent spike trains, and both had trial-to-trial variation and latency effects; details are given in Table 3. In the second scenario the two neurons (Neurons F1 and F2) were similar to Neurons E1 and E2, but they also had an excess of synchronous spikes, described by $\zeta_0(t)$ in Table 3 and shown in Figure 4. The firing rates of these four neurons are not shown because they are similar to neuron C in Figure 2.

Before we can assess whether synchrony is present, we must fit an appropriate firing rate model to each of the four neurons. First, the latency test of Ventura (2004) detected the presence of latency effects in the firing rates of the four neurons, each with p-values less than 10^{-5}. We therefore estimated the latency effects as in Ventura (2004), and shifted the trials according to the latency estimates.

The excitability tests described in Section 2.1 and illustrated in Section 3.1 were applied

to the four neurons separately. We found that both neurons in the pair E1/E2 had firing rates suitably described by two PCAs, whereas both neurons in the pair F1/F2 had firing rates suitably described by just one PCA. To assess whether the excitability effects are shared by the two neurons, the test introduced in Section 2.5 was used. The p-values given in Table 4 lead to the correct conclusion that the same latency and excitability effects should be fitted to each pair of simulated neurons.

Figure 4 shows the cross-correlogram and the fitted $\hat{\zeta}_0(t)$ for pairs of neurons E1/E2 and F1/F2, with Bootstrap bands, before and after correction for latency and excitability effects.

For neurons E1/E2 (no synchrony), after the latency effect is corrected in Figure 4(B), when there still exists the excitability effect, most of the cross-correlogram values are still outside the 95% bootstrap band. After both the latency and excitability effect is corrected in Figure 4(C), the cross-correlogram indicates that the two neurons are independent. Figure 4(E) shows the estimated $\zeta_0(t)$ of the joint spiking model for the diagonal of the JPSTH. Before being corrected for the latency effect, the estimated ζ_0 has a peak in the middle beyond the Bootstrap band. The area outside the band has p-value < 0.0001. After the latency effect is screened out but still with the excitability effect, the magnitude of the $\hat{\zeta}$ is reduced but still outside the band with p-value of the largest area outside band equal to 0.001; all p-values are recorded in Table 5. After both the latency and excitability effect are corrected, the magnitude of $\hat{\zeta}$ is further reduced and the largest area outside the bands has p-value 0.63. This leads correctly to the conclusion that the two neurons are independent.

The same procedure also yields the correct conclusion for the correlated pair F1/F2.

Neurons E1/E2 : no synchrony

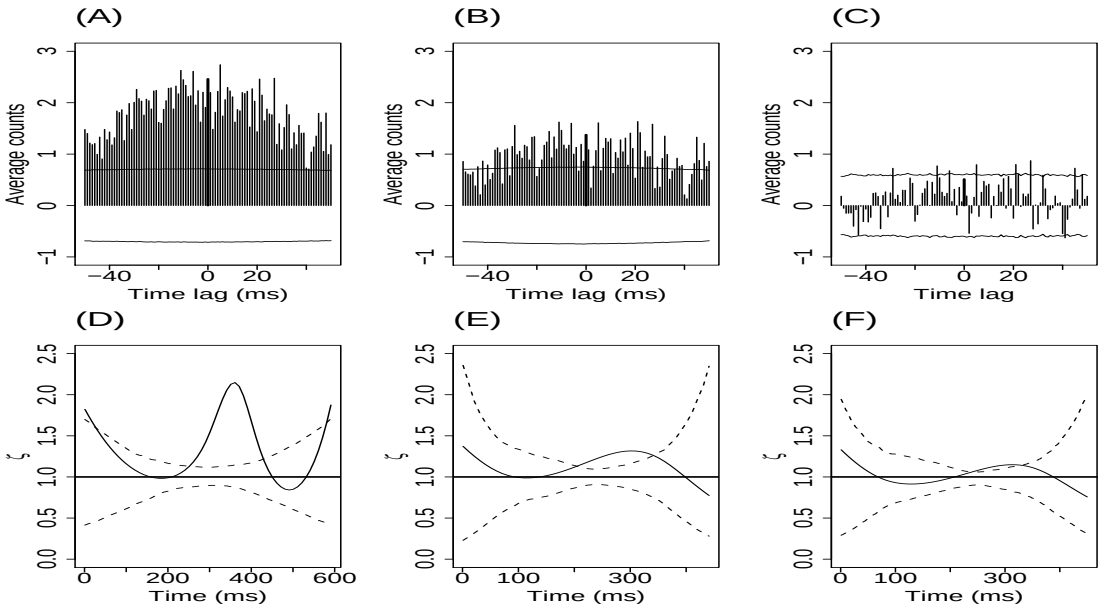

Neurons F1/F2 : synchrony

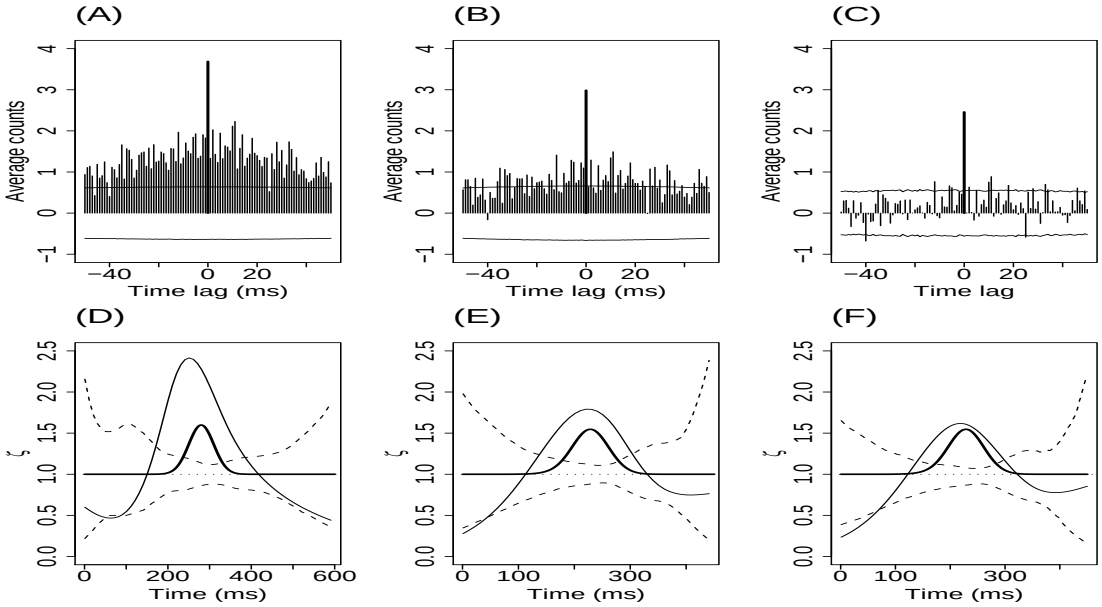

Figure 4: *Neurons pairs with latency and excitability effects, and no synchrony (E1/E2) or synchrony (F1/F2). (A) Shuffle corrected cross-correlogram before correction for latency and excitability, (B) after correction for latency, and (C) after correction for latency and excitability, with 95% Bootstrap bands. (D), (E) and (F) Corresponding panels for* $\hat{\zeta}_0(t)$ *of the joint spiking model, with 95% bootstrap bands; the p-values for the excursion of* $\hat{\zeta}_0(t)$ *outside the bands are in Table 5.*

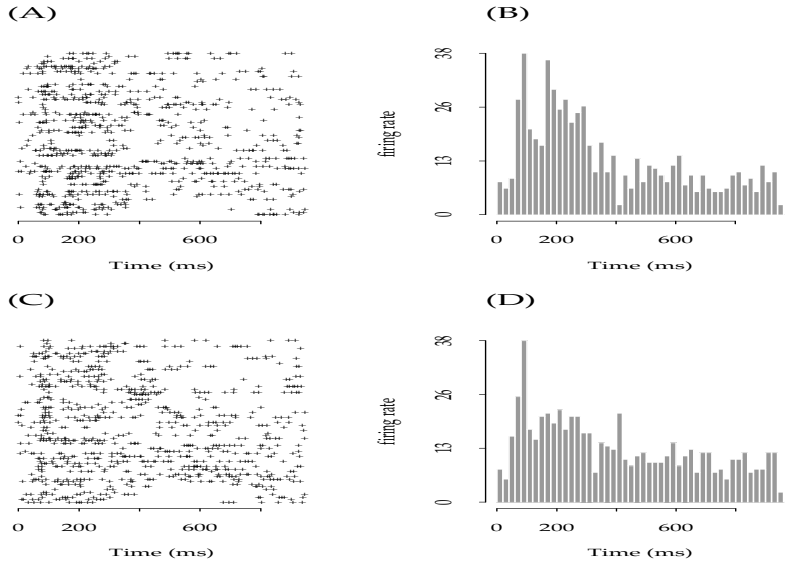

Figure 5: *Raster plot and PSTH for (A–B) neuron 380506.s and (C–D) neuron 380506.u. The time bins are 2.8 msec wide.*

3.4 Application to a Pair of V1 Neurons

We assessed the effect of trial-to-trial variability on the correlation of two neurons recorded simultaneously in the primary visual cortex of an anesthetized macaque monkey (Aronov *et al.* 2003, units 380506.st, 90 deg spatial phase). The data consist of 64 trials shown in Figure 5. The stimulus was identical for all trials, and consisted of a standing sinusoidal grating that appeared at time 0, and disappeared at 237 ms.

Figure 6(A), which displays the rate adjusted CC with bins 2.8 ms wide[1], suggests that spike time synchronization may occur at many time lags, but it may be masked by effects of trial-to-trial variation. The estimated time course synchrony at lag $2 \times 2.8 = 5.6$ msec, $\hat{\zeta}_{5.6}(t)$ is also displayed in Figure 6(C), along with 95% Bootstrap bands; $\hat{\zeta}_{5.6}(t)$ exceeds

[1]The choice of 2.8 ms is due to a limitation in spike recording during the experiment: if the spike of one neuron occurred immediately (less than roughly 1.3 - 2 ms) after the firing of the other neuron, it could not be recorded. Because it is impossible to detect a joint spike for for a time lag less than 2.8 we analyze time lags greater than 2.8 ms or less than -2.8 ms.

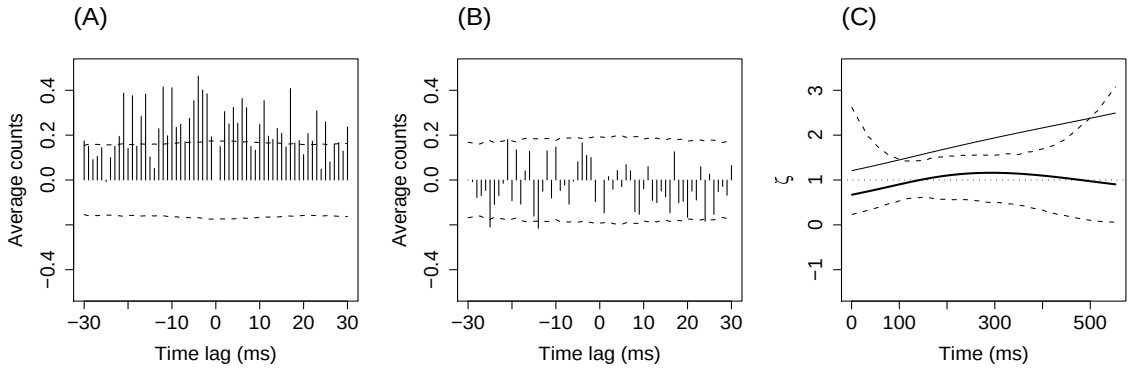

Figure 6: *Cross-correlograms for two neurons in Aronov* et al. *(2003), (A) adjusted for firing rate modulations only and (B) adjusted for both firing rate modulation and time-varying excitability effects. The dashed curves are pointwise 95% probability bands computed under the null hypothesis of independence. (C) The time-varying synchrony curve* $\hat{\zeta}_{5.6}(t)$ *before (thin solid curve) and after (bold curve) the adjustment for time-varying excitability effects. The dashed curves are pointwise 95% Bootstrap bands computed under the null hypothesis* $\zeta_{5.6}(t) = 1$ *for all* t.

the upper band for most values of t, suggesting that lagged synchrony between the two neurons is significant for almost for the entire duration of the experiment.

To adjust for trial-to-trial variability we applied the latency test of Ventura (2004) and the tests for excitability effects described above. Before doing so a preliminary check on the Poisson assumption was carried out. Based on a mean versus variance plot, there appeared to be a small amount of extra variability above that predicted by the Poisson assumption. The Poisson test of Brown *et al.* (2002), based on the time rescaling theorem for Poisson processes, also confirmed a very slight departure from Poisson. These effects may have been due to trial-to-trial variation and were not large enough to warrant the more complicated procedures required for non-Poisson spike trains. When these analyses were repeated after removal of the trial-to-trial effects (described below), the Poisson assumption was judged to be adequate.

The test for latency effects in Ventura (2004) suggested that no latency effects were present, with p-values $P = 0.096$ and $P = 0.174$ for the two neurons respectively. However, the test for excitability effects were highly statistically significant ($P = 0.00014$ and $P = 0.00019$ respectively). Consideration of shared excitability models showed that the most appropriate model was the shared non-constant excitability model ($P < 0.00001$ against the constant excitability model, $P = 0.752$ for different versus same excitability effects). Applying the sequential fitting procedure we then determined that two principal components should be used to capture the non-constant excitability effects, with each principal component explaining respectively 83% and 14% of the variability in the observed pairs of spike trains. Therefore, the firing rate of neuron $i = 1, 2$ on trial r, $\lambda_r^i(t)$, can be summarized as

$$\lambda_r^i(t) = \lambda^i(t) \exp\left(w_{r0} + w_{r1}\phi_1(t) + w_{r2}\phi_2(t)\right) \tag{16}$$

where the two principal components $\phi_1(t)$ and $\phi_2(t)$ are displayed in Figure 7.

An interesting additional interpretation is obtained by examining the relationship of the coefficients w_{r0}, w_{r1}, w_{r2} across trials. Figure 7, suggests that w_1 and w_2 are almost equal across trials, and that $(w_1 + w_2)/2$ is close to $-w_0$, so that the firing rate model in (16) is close to

$$\lambda_r^i(t) = \lambda^i(t)\, \alpha_{r0}\, \exp(\phi(t)), \tag{17}$$

where $\alpha_0^r = \exp(w_0^r)$ and $\phi(t) = 1 - \phi_1(t) - \phi_2(t)$ is also plotted in Figure 7. That is, the trial-to-trial variability is reasonably well represented by a function that itself resembles the individual firing rate intensities (but is statistically distinguishable from them).

Figure 6 (B) shows the corrected cross-correlogram, with Bootstrap bands obtained as in Section 2.3. We conclude that the correlation indicated by the uncorrected cross-correlogram (in Part (A)) is due entirely to excitability effects. Figure 6(C) displays the estimates of $\zeta_{5.6}(t)$ (corresponding to the second bin to the right of 0 in the cross-correlograms) before and after the adjustment for excitability. Before excitability adjust-

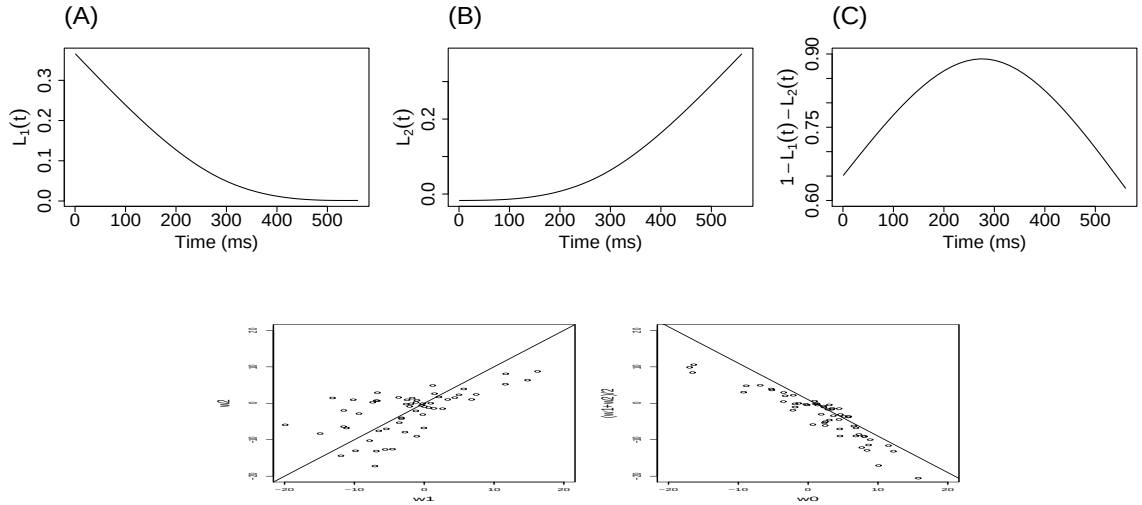

Figure 7: *Shared excitability effects for two neurons from Aronov* et al. *(2003). (A) First and (B) second principal components, $\phi_1(t)$ and $\phi_2(t)$. (C) The function $\phi(t) = 1 - \phi_1(t) - \phi_2(t)$ in (17). Bottom panel, left: Plot of the coefficients w_1 against w_2. Each point in the plot represents a coefficient pair on a particular trial. Bottom panel, right: Plot of $\frac{w_1+w_2}{2}$ against w_0.*

ment, the p-value for synchrony at lag 5.6 ms is less than 10^{-5} while after the adjustment $\hat{\zeta}(t)$ lies within the Bootstrap bands. There is, therefore, no evidence of synchrony at that lag, beyond the correlations induced by rate and excitability modulations. We reached similar conclusion at all lags; Figure 12 in the appendix display the same as Figure 6(C) for all lags.

Based all the analysis in this section, we conclude that (1) there is very strong evidence of time-varying excitability for these two neurons, (2) the excitability effects appear to be shared by the two neurons, and (3) there is no evidence of spike timing synchronization between these neurons.

3.5 The Effect of Trial-to-trial Variation on Spike Count Correlation as the Measurement Interval Length Increases

We now investigate the way excitability effects influence spike count correlation, as a function of the length of interval over which the count is recorded. Consider random variables Y_r^1 and Y_r^2 representing theoretical spike counts over an interval of length c for two neurons, and for simplicity assume they follow the same probability distribution. We will also assume the following:

1. The spike counts tend to increase proportionally to c in the sense that the expectation of Y_r^i $(i = 1, 2)$ is proportional to c.

2. The constant excitability model holds, so that the expected spike count for each neuron is multiplied by a random variable e^{X_r}. We may then write the expectation conditionally on the trial effect X_r in the form

$$E(Y_r^i|X_r) = c \cdot \beta \cdot e^{X_r}$$

   where β is the expected spike count when $c = 1$ and $X_r = 0$.

3. After conditioning on the trial r, Y_r^1 and Y_r^2 are independent.

4. The conditional variance is proportional to the conditional expectation:

$$V(Y_r^i|X_r) = k \cdot E(Y_r^i|X_r).$$

Under these assumptions the correlation of Y_r^1 and Y_r^2 may be computed in terms of c. Note that in this situation, because of Assumption 3, the correlation is due entirely to trial-to-trial variation. Also, in the case of Poisson counts we would have $k = 1$. Under-dispersion occurs when $k < 1$ and over-dispersion when $k > 1$.

We now compute the correlation of Y_r^1 and Y_r^2 under these assumptions. We omit the subscript r and use standard (elementary) formulas for the variances and covariance in terms of the conditional variances, covariance, and expectations. For $i = 1, 2$ we have

$$\begin{aligned} V(Y^i) &= E(V(Y^i|X)) + V(E(Y^i|X)) \\ &= E(kc\beta e^X) + V(c\beta e^{+X}) \\ &= kc\beta E(e^X) + c^2\beta^2 V(e^X) \end{aligned}$$

and

$$\begin{aligned} COV(Y_1, Y_2) &= E(COV(Y^1, Y^2|X)) + COV(E(Y^1|X), E(Y^2|X)) \\ &= 0 + c^2\beta^2 V(e^X) \end{aligned}$$

and, writing $\mu = E(e^X)$ and $\sigma^2 = V(e^X)$ we therefore obtain

$$\begin{aligned} COR(Y^1, Y^2) &= \frac{c^2\beta^2\sigma^2}{kc\beta\mu + c^2\beta^2\sigma^2} \\ &= \frac{c\beta}{\frac{k}{\gamma} + c\beta} \end{aligned} \tag{18}$$

where $\gamma = \sigma^2/\mu$. This shows that the correlation will increase as c increases.

To see the implication of Equation (18), suppose that the constant excitability e^{X_r} has a log-Normal distribution, i.e., $\log X_r$ has a Normal distribution. It is easily verified that if Z has a Normal distribution with mean 0 and variance 1 then

$$E(e^{a+bZ}) = e^{a+b^2/2}$$

and

$$V(e^{a+bZ}) = e^{2a+b^2}(e^{b^2} - 1).$$

These give the ratio

$$\gamma = e^{a+b^2/2}(e^{b^2} - 1).$$

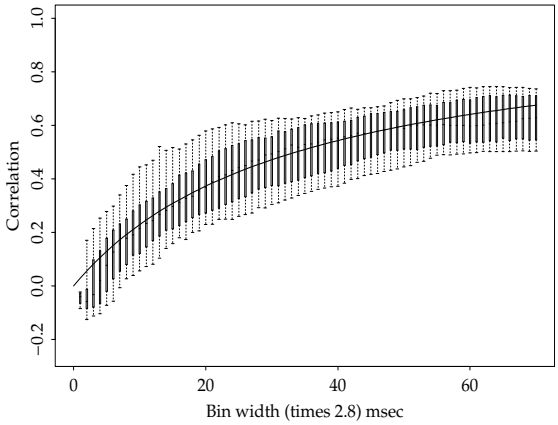

Figure 8: *Correlation across trials between the spike counts for the two neurons from Aronov* et al. *(2003), as a function of the bin size (*$\times 2.8$ *msec). For each bin size, we plotted the box plot (10th, 90th quantiles and quartiles) of the correlations obtained by sliding the bin along experimental time. The bold curve is (18) with* $\frac{k}{\beta\gamma} = 34$.

Using this formula we may compute the correlation for various scenarios. For example, with a mean firing rate of 50 spikes per second ($\beta = 50$), for $k = 1$, in the presence of 25% trial-to-trial variation (X_r has mean $a = 0$ and standard deviation $b = .25$), the correlation of two spike counts in intervals of length 2ms, 20ms, and 200ms ($c = .002, .02, .2$) will be .007, .06, and .4. With a firing rate of 20 spikes per second and 12.5% trial-to-trial variation, correlations for counts in intervals of length 2ms, 100ms, and 1000ms become .0006, .03, and .24, which are roughly consistent with those reported by Reich *et al.* (2001). Figure 8 displays the correlation of spike counts as a function of bin size for the pair of V1 neurons analyzed above (in Figure 5). Overlayed is a plot of the function in Equation (18), which agrees well with the data.

4 Discussion

The statistical methodology presented here uses a small number of parameters to describe time-varying trial-to-trial variation, or "excitability" effects. It provides assessments of whether these effects (i) are present, (ii) are constant or time-varying, (iii) are shared by two neurons, and (iv) are adequately described by one or two parameters. In addition, the methodology allows the trial-to-trial to be removed, so that assessments of correlation and synchrony, or lagged time-locked firing, may be suitably adjusted. The fitting procedures may be implemented with widely-available software for generalized linear models, and the Bootstrap is easily implemented in high-level programming languages such as MATLAB or R. The use of model (3), with insertion of the appropriate offset for each neuron, allows the assessment of whether excitability effects are shared across neurons, which may be of substantial interest. The models (4) (and their non-Poisson generalization) allows excitability effects to be removed so that correlation and synchrony assessments may be adjusted, and the restriction in (5) reduces the noise in the fitting, making the inference procedures more powerful.

One limitation of the method used here is that we have ignored variability introduced in steps 1 and 2 of the fitting process when we make statistical inferences, subsequently. More complicated Bayesian or Bootstrap methods could provide more comprehensive inference procedures. Similarly, it would be possible to incorporate estimation of latency into a general procedure, rather than to rely on a preliminary assessment using the method of Ventura (2004).

We showed that multiplicative excitability effects produce spike count correlations that grow as the spike count interval increases. This result does not hold for additive excitability effects (following a similar approach). For realistic firing rates spike count correlations that are very small on time scales of tens of milliseconds become non-negligible on scales of hundreds of milliseconds. This may help explain correlations that have been reported

and discussed (Shadlen and Newsome 1998). The method presented here has been applied to two neurons, but the general approach may, in principle, be extended to allow consideration of shared effects among many neurons recorded simultaneously.

Excess trial-to-trial variability, beyond that predicted by stochastic variation of point processes that have no latency or excitability effects, may be common in brain regions involved in higher-order processing. One way to relate trial-to-trial variation in neuronal activity to behavioral measures, such as reaction time, is to record large numbers of trials and then aggregate them into groups according to levels of the behavioral response, as in Hanes and Schall (1996). The methods used here, however, are applicable with much smaller numbers of trials.

5 Appendix: Supplementary Figures

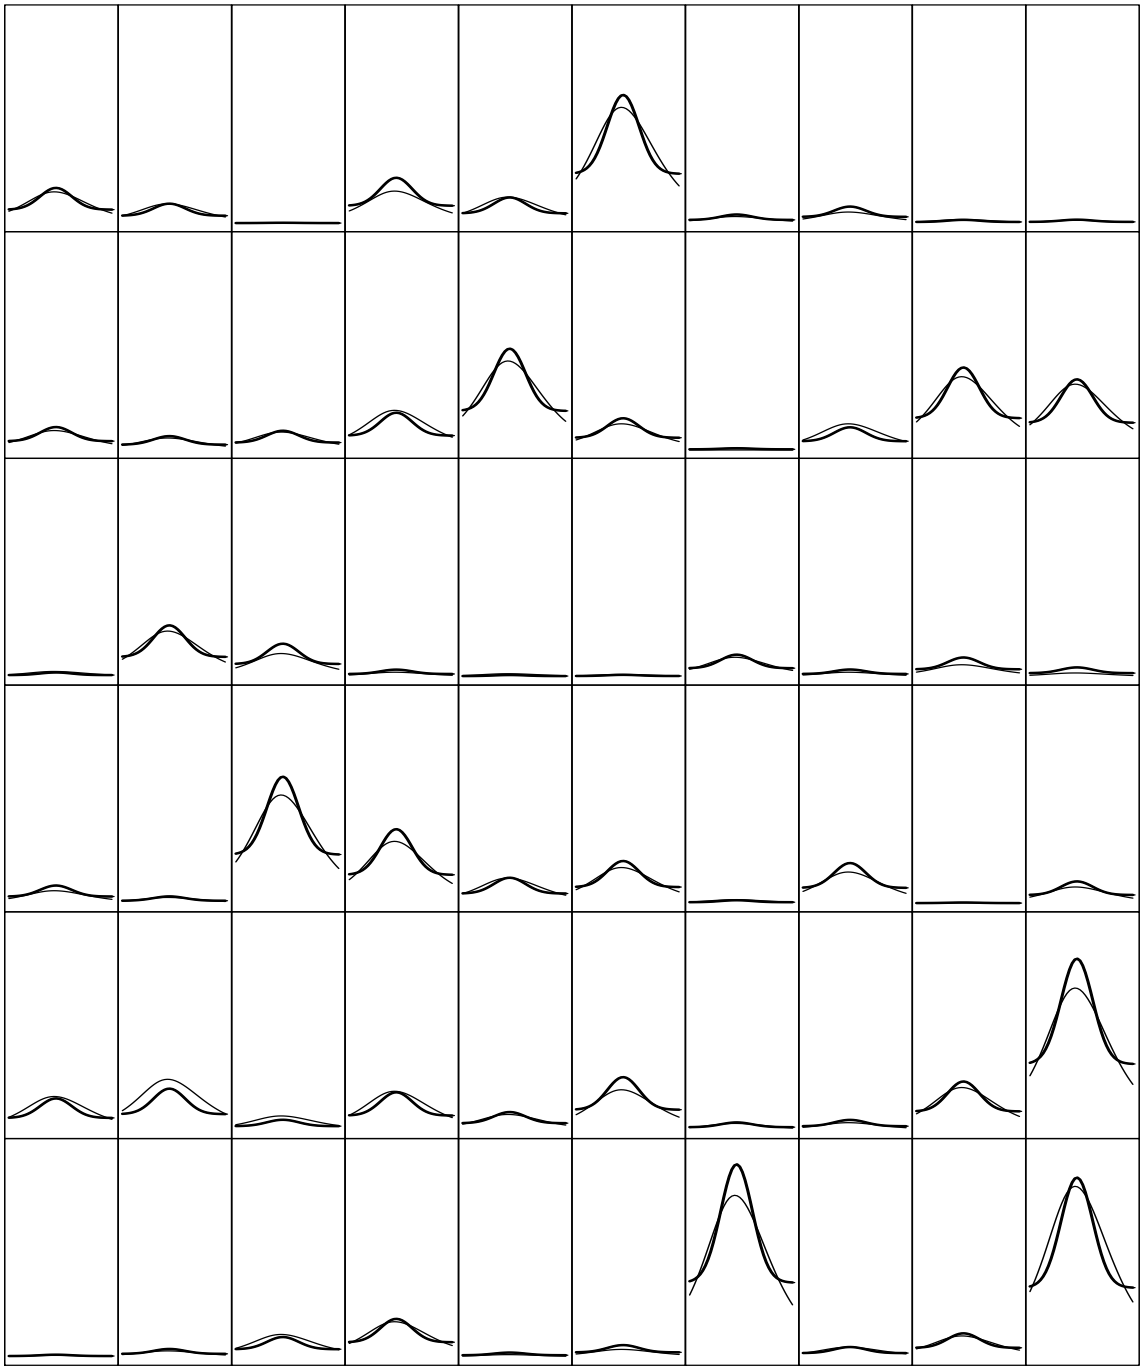

Figure 9: *Neuron B: true (bold curves) and fitted (thin curves) rates for all simulated trials.*

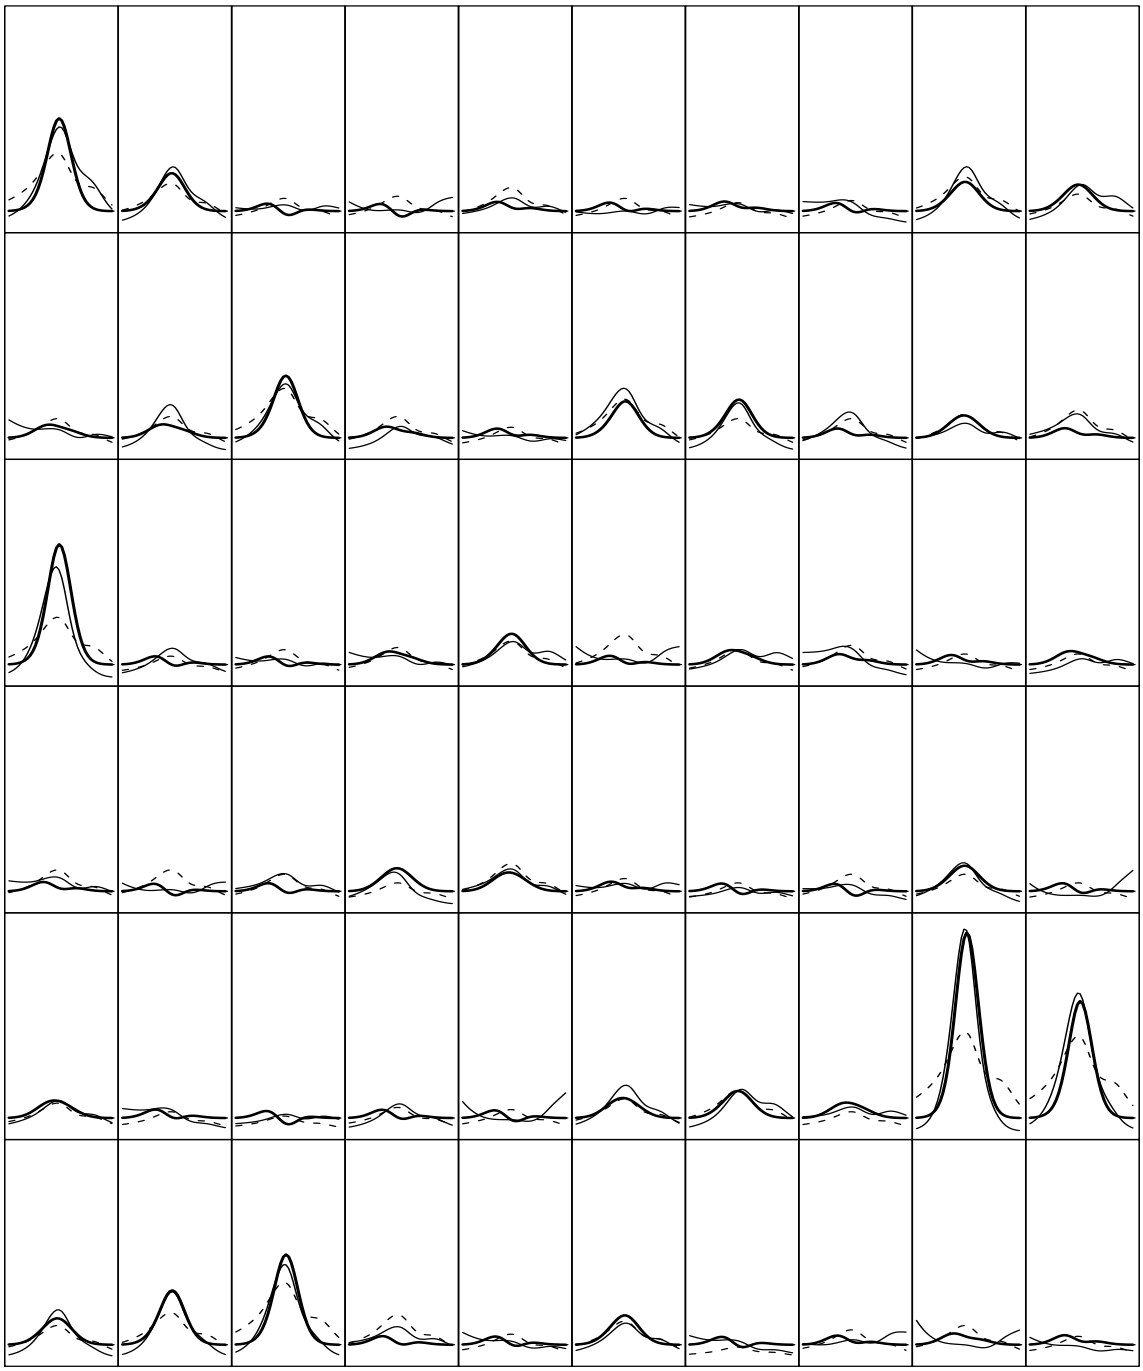

Figure 10: *Neuron C: true (bold curves) and fitted (thin curves) rates for all simulated trials. The dashed curves are the fits from the constant excitability models, which was rejected by the deviance test.*

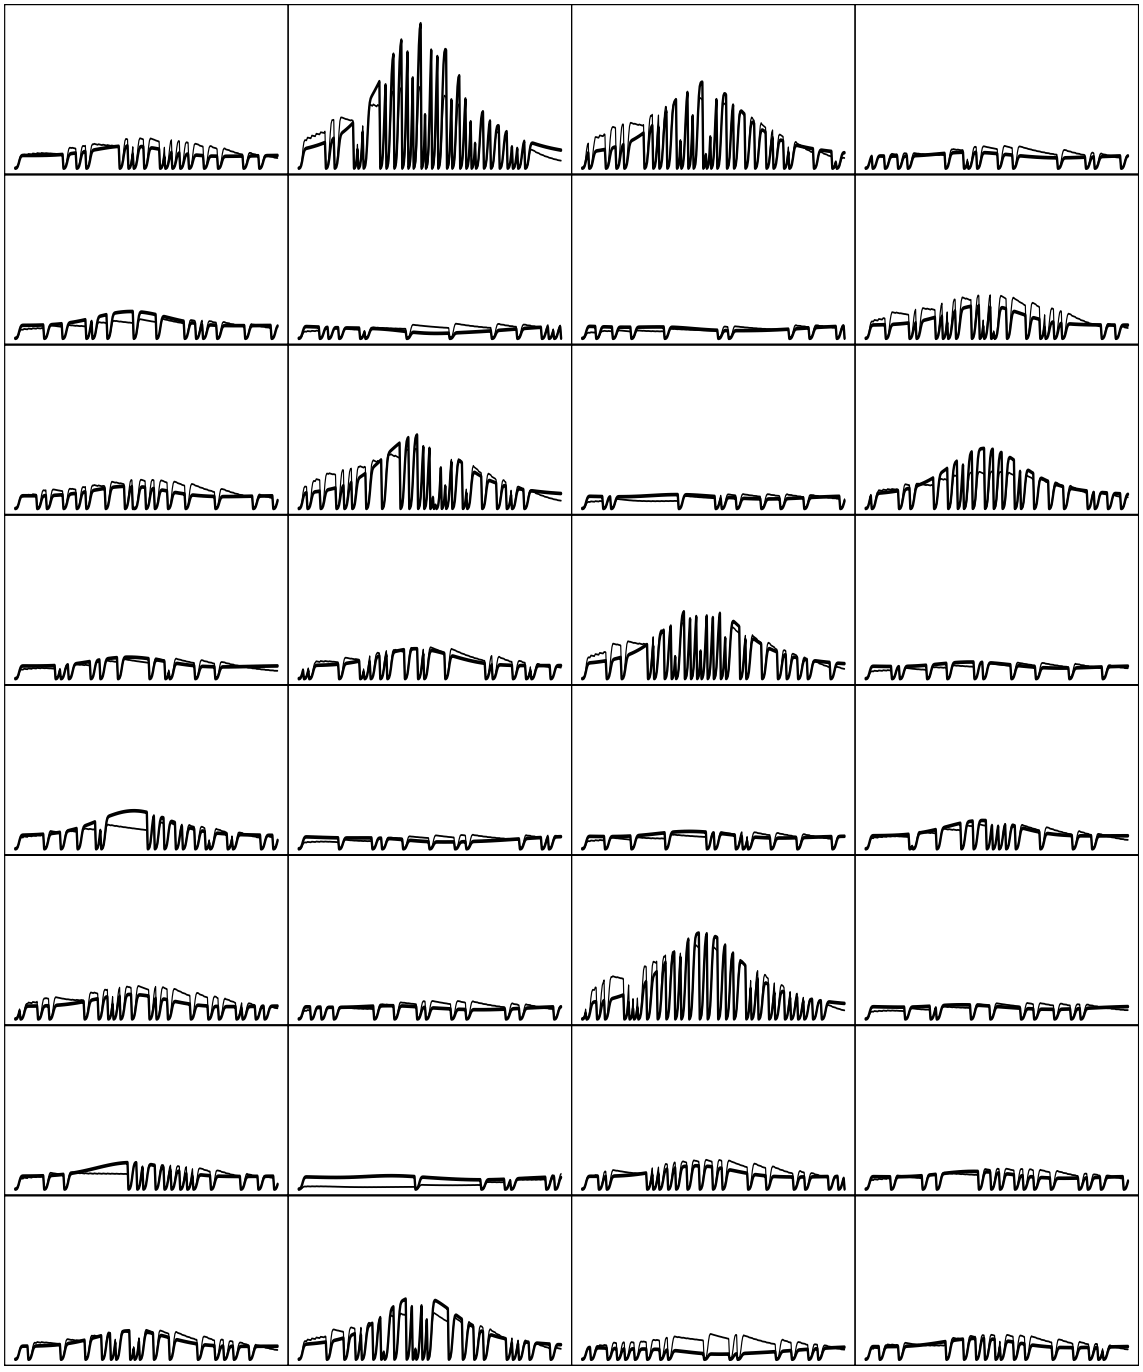

Figure 11: *Neuron D: true (bold curves) and fitted (thin curves) rates for all simulated trials.*

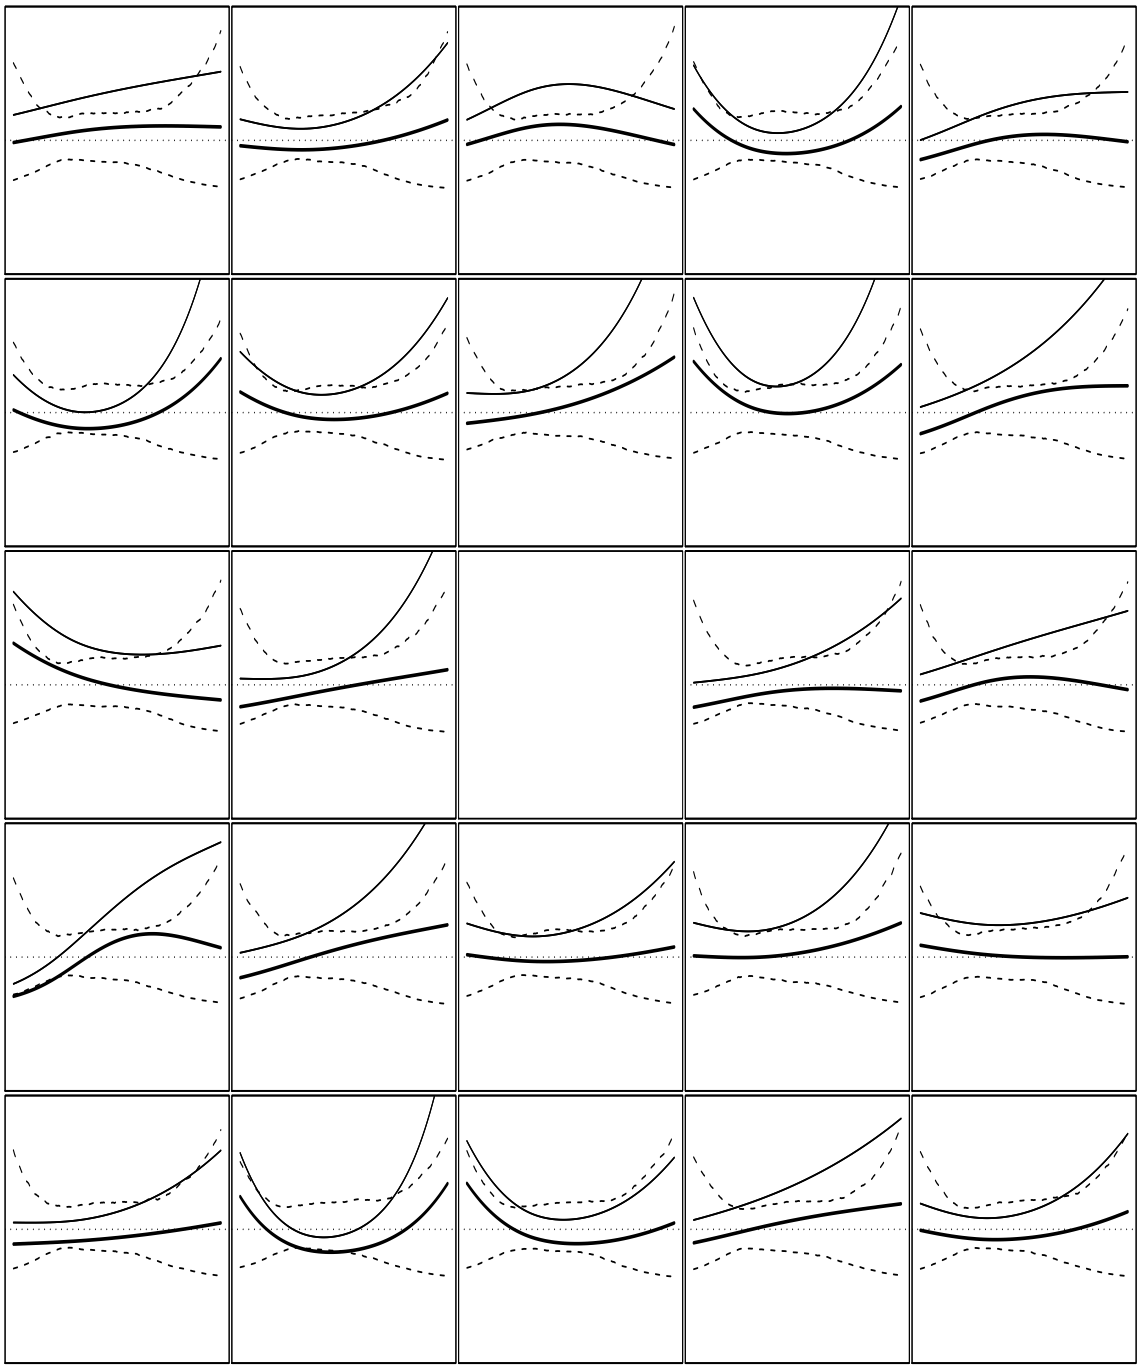

Figure 12: *The estimated $\zeta_\delta(t)$ for neuron pair 380506.s/380506.u before and after adjustment for the excitability covariation. Panels from left to right and top to bottom correspond to lags δ from -33.6 msec to 33.6 msec in steps of 2.8 msec. The synchrony plot for $\delta = 0$ was omitted due to the limitations of the recording of joint spikes.*

Grants

Cai and Kass: MH64537

Ventura: MH64537 and N01-NS-2-2346

References

- **Aersten, A, Gerstein, G, Habib, M, and Palm, G.** Dynamics of neuronal firing correlation - modulation in effective connectivity. *Journal Of Neurophysiology* 61: 900–917, 1989.

- **Aronov, D, Reich, DS, Mechler, F, and Victor, JD.** Neural coding of spatial phase in V1 of the macaque monkey. *J Neurophysiol* 89: 3304–3327, 2003.

- **Azouz, R and Gray, CM.** Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J. Neurosci.* 19: 2209–2223, 1999.

- **Baker, SN and Gerstein, GL.** Determination of response latency and its application to normalization of cross-correlation measures. *Neural Computation* 13: 1351–1377, 2001.

- **Barbieri, R, Quirk, MC, Frank, LM, Wilson, MA, and Brown, EN.** Non-Poisson stimulus-response models of neural spike train activity. *J. Neurosci. Methods* 105: 25–37, 2001.

- **Ben-Shaul, Y, Bergman, H, Ritov, Y, and Abeles, M.** Trial to trial variability in either stimulus or action causes apparent correlation and synchrony in neuronal activity. *J. Neurosci. Methods* 111: 99–110, 2001.

- **Brillinger, D.** Some statistical methods for random process data from seismology and neurophysiology. *The Annals of Statistics* 16: 1–54, 1988.

- **Brody, CD.** Correlations without synchrony. *Neural Computation* 11: 1537–1551, 1999a.

- **Brody, CD.** Disambiguating different covariation. *Neural Computation* 11: 1527–1535, 1999b.

- **Brown, EN, Barbieri, R, Ventura, V, Kass, RE, and Frank, LM.** (2002). A note on the time-rescaling theorem and its implications for neural data analysis. *Neural Computation* 14(2): 325–46, 2001.

- **Daley, DJ and Vere-Jones, D.** *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods, Second Edition*, New York: Springer, 2003.

- **Cai,C, Kass, RE, and Ventura, V.** *Statistical Assessment of Time-Varying Dependence Between Two Neurons.* Technical Report ..., Statistics Department, Carnegie Mellon University, 2004a.

- **Daley, DJ and Vere-JOnes, D.** *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods, Second Edition*, New York: Springer, 2003.

- **DiMatteo, I, Genovese, CR, and Kass, RE.** Bayesian curve fitting with free–knot splines. *Biometrika*, 88: 1055-1073, 2001.

- **Grün, S, Riehle, A, and Diesmann, M.** Effect of across trial non-stationarity on joint-spike events. *Biol. Cybernetics*, 2003 (in press).

- **Hanes, DP and Schall, JD.** Neural control of voluntary movement initiation. *Science* 274: 427–430, 1996.

- **Hastie, T and Tibshirani, R.** *Generalized Additive Models.* London: Chapman and Hall, 1990.

- **Johnson, D.** Point process models of single-neuron discharges. *J. Comput. Neurosci.* 3: 275–299, 1996.

- **Kass, RE and Ventura, V.** A Spike-train probability model. *Neural Comput.* 13: 1713–1720, 2001.

- **Kass, RE, Ventura, V, and Cai, C.** Statistical smoothing of neuronal data. *Network: Comput. Neural Syst.* 14: 5-15, 2003.

- **McCullagh, P and Nelder, J.A.** *Generalized Linear Models*, Second Edition. New York: Chapman and Hall, 1990.

- **Nawrot M, Aertsen A, and Rotter, S.** Single-trial estimation of neuronal firing rates. *J Neurosci. Methods*, 94: 81-91, 1999.

- **Reich, DS, Mechler, F, and Victor, JD.** Independent and redundant information in nearby cortical neurons. *Science* 294: 2566–2568, 2001.

- **Shadlen, MN and Newsome, WT.** The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *J. Neurosci.* 18(10): 3870–3896, 1998.

- **Ventura, V.** Testing for, and estimating latency effects for Poisson and non-Poisson spike trains. *Neural Computation*, 2004 (To appear).

List of Figures

		Deviance *(p-value)*			
		Neuron			
Model	*df*	*A*	*B*	*C*	*D*
No excitability	*0*	*1525 **	*3090*	*2205*	*2983*
Constant excitability	R	*1482* *(0.94)*	*1572 ** *($< 10^{-5}$) **	*2018* *($< 10^{-5}$)*	*2920* *(0.001)*
1 PCA excitability	$2R$	*1420* *(0.5)*	*1515* *(0.59)*	*1862 ** *($< 10^{-5}$) **	*2872 ** *(0.034) **
2 PCAs excitability	$3R$	× ×	× ×	*1801* *(0.44)*	*2872* *(1)*

Table 2: Deviances and significance tests p-values for Poisson neurons A, B, and C, and non-Poisson neuron D: smoothed-PSTH model (no excitability) $\lambda(t)$ vs. constant excitability model $e^{w_{0r}}\lambda(t)$ vs. time-varying single-principal component model $e^{w_{0r}+w_{1r}\phi_1(t)}\lambda(t)$ vs. time-varying two-principal component model $e^{w_{0r}+w_{1r}\phi_1(t)+w_{2r}\phi_2(t)}\lambda(t)$. R is the number of trials; we used $R = 60$ for neurons A, B, and C, and $R = 32$ for neuron D. The "" indicates the model of choice.*

Neuron	Excitability effect $g_r(t)$	Global rate $\lambda(t)$
E1 / F1	$g_r^1(t) = g_r^2(t) = 1 + c_r f(t, 390, 35)$	$\lambda^1(t) = 0.04 + 24 f(t, 390, 40)$
E2 / F2	$c_r = b_r - \bar{b}$, $b_r \sim Gamma(1, 0.5)$	$\lambda^2(t) = 0.04 + 24 f(t, 390, 60)$
	Common latencies : $\tau_r \sim f(t, 0, 40)$	
	$\zeta_0(t) = 1 + 45 f(t, 380, 30)$	

Table 3: *Firing rates $\lambda_r(t)$ given by (1) for neurons E1/E2 and F1/F2. $f(t, a, b)$ denotes the normal density function with mean a and standard deviation b. " $\sim$ " means "randomly generated from".*

			Deviance *(p-value)*	
Model		*df*	*Neuron pair* *E1/E2*	*F1/F2*
1 PCA models	*Same excitability*	$2R$	$\times$	*4447* * *($< 10^{-5}$)* *
	Proportional excitability	$2R+1$	$\times$	*4447* *(0.96)*
	Different excitability	$4R$	$\times$	*4344* *(0.40)*
2 PCAs models	*Same excitability*	$3R$	*4270* * *($< 10^{-5}$)* *	$\times$
	Proportional excitability	$3R+1$	*4270* *(0.9 6)*	$\times$
	Different excitability	$6R$	*4232* *(1)*	$\times$

Table 4: Deviances and p-values for neuron pairs E1/E2 and F1/F2. This table records the deviance and p-values of the tests that determines if the two neurons in the pair have the same, proportional, or different excitability effects. The two neurons of the pair are fitted together with models: same excitability model $e^{w_{0r}+w_{1r}\phi_1(t)}\lambda_1(t)$ *and* $e^{w_{0r}+w_{1r}\phi_1(t)}\lambda_2(t)$ *vs. proportional excitability model* $e^{w_{0r}+w_{1r}\phi_1(t)}\lambda_1(t)$ *and* $e^{w_{0r}+w_{1r}\phi_1(t)}\lambda_2(t)\alpha$ *vs. different excitability model* $e^{w^1_{0r}+w^1_{1r}\phi^1_1(t)}\lambda_1(t)$ *and* $e^{w^2_{0r}+w^2_{1r}\phi^2_1(t)}\lambda_2(t)$ *(note that for this last model,* $e^{w^1_{0r}}$ *and* $e^{w^2_{0r}}$*,* $e^{w^1_{1r}}$ *and* $e^{w^2_{1r}}$*, and* $\phi^1_1(t)$ *and* $\phi^2_1(t)$ *are not constrained to be equal). The same sequences of models are fitted for the 2 PCAs models. Here,* $R = 60$*, the number of trials for one neuron. The "*" indicates the model of choice.*

Model corrected	*P-values for* $\zeta_0(t)$	
for modulations in	*Neurons E1/E2*	*Neurons F1/F2*
firing rate	$< 10^{-5}$	$< 10^{-5}$
firing rate, latencies	*0.001*	$< 10^{-5}$
firing rate, latencies, excitability	*0.63*	$< 10^{-5}$

Table 5: P-values of the excursion of $\hat{\zeta}_0(t)$ *outside the Bootstrap bands.*