



Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons

Valérie Ventura, Can Cai and Robert É. Kass

JN 94:2928-2939, 2005. doi:10.1152/jn.00644.2004

You might find this additional information useful...

This article cites 24 articles, 16 of which you can access free at:

<http://jn.physiology.org/cgi/content/full/94/4/2928#BIBL>

This article has been cited by 1 other HighWire hosted article:

Statistical Assessment of Time-Varying Dependency Between Two Neurons

V. Ventura, C. Cai and R. E. Kass

J Neurophysiol, October 1, 2005; 94 (4): 2940-2947.

[\[Abstract\]](#) [\[Full Text\]](#) [\[PDF\]](#)

Updated information and services including high-resolution figures, can be found at:

<http://jn.physiology.org/cgi/content/full/94/4/2928>

Additional material and information about *Journal of Neurophysiology* can be found at:

<http://www.the-aps.org/publications/jn>

This information is current as of October 26, 2005 .



Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons

Valérie Ventura, Can Cai, and Robert E. Kass

Department of Statistics and Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, Pennsylvania

Submitted 25 June 2004; accepted in final form 23 March 2005

Ventura, Valérie, Can Cai, and Robert E. Kass. Trial-to-trial variability and its effect on time-varying dependency between two neurons. *J Neurophysiol* 94: 2928–2939, 2005; doi:10.1152/jn.00644.2004. 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 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 in earlier studies by Brody and Ben-Shaul et al., 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 decomposes the spike-train variability into a stimulus-related component and a trial-specific component, allowing many degrees of freedom to characterize the former while assuming a small number suffices to characterize the latter. The Bootstrap significance test of the companion paper 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. In data from two VI neurons we find that highly statistically significant evidence of dependency disappears after adjustment for time-varying trial-to-trial variation.

INTRODUCTION

Spike trains recorded from behaving animals display variation in spike timing both within and across repeated trials. In some cases, the irregularity 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 internal state of 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 dependency between spike trains from two neurons by extending the method of Ventura et al. (2005a).

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 realignment of trials becomes desirable (Baker and Gerstein 2001; Ventura 2004; Woody 1967). 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) or by smoothing the interspike intervals of each trial and inverting to get a firing rate (Pauluis and Baker 2000). However, smoothing each trial ignores a special feature of this situation: although there is substantial information on the general shape of the firing-rate function obtained from the aggregated trials (i.e., the smoothed peristimulus time histogram [PSTH]), there is relatively little information per trial from which to estimate a time-varying function, at least for smaller firing rates (e.g., <40 Hz). This creates a compelling need for statistically efficient procedures; otherwise, standard errors and confidence intervals become very wide and statistical tests have little power (Kass et al. 2005). 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), who 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. Although likely to capture some dominant features of trial-to-trial variation, this model may be too restrictive 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

Address for reprint requests and other correspondence: V. Ventura, Department of Statistics and Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213 (E-mail: vventura@stat.cmu.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

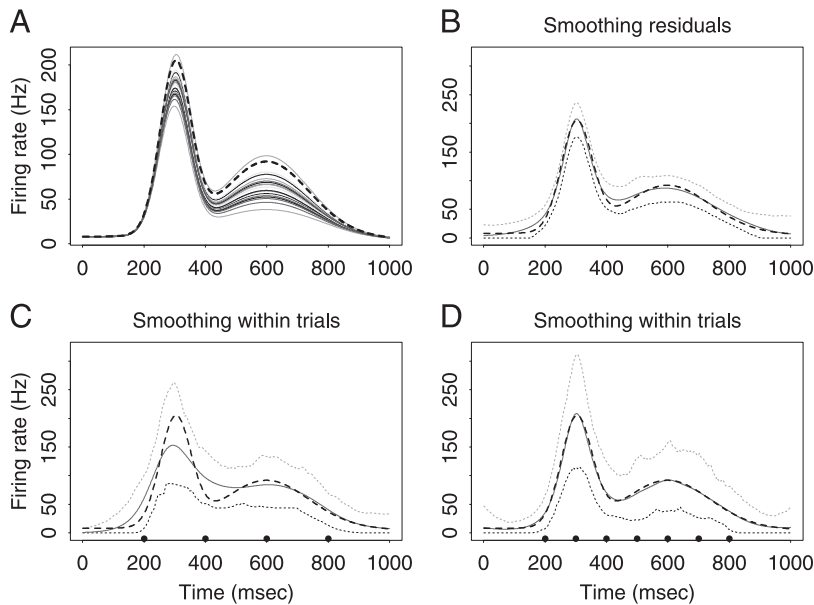


FIG. 1. Smoothing excitability effects. *A*: firing rate functions $P_r(t)$, $r = 1, \dots, 20$, of a simulated neuron having an excitability effect described by model Eq. 1, with latency effects $\tau_r = 0$ for all r . Magnitude and shape of the firing rate function vary from trial to trial. *B*, *C*, *D*: spike trains were generated from the rate functions in *A*, and the firing rate function $P_r(t)$ was estimated, for $r = 8$, using 3 methods described in the text. Spike trains were generated in sets of 20 trials, and a total of 200 such sets were used to obtain, for each method, the mean estimated $\hat{P}_r(t)$ (solid), and 95% pointwise probability bands for $P_r(t)$ (dotted). True $P_r(t)$ for $r = 8$ is also shown (dashed bold). Bold dots on the x -axes in *C* and *D* indicate the locations of the spline knots used for smoothing. Corresponding plots of the method in *B* for all 20 trials can be found in APPENDIX, Fig. A1.

version of it to estimate the stimulus-induced firing rate, and the assumption of Poisson spiking. The method described here avoids the requirement of a background period because it fits the excitability effects as continuous functions of time, 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.

METHODS

The phenomenon we seek to describe is illustrated in Fig. 1A, which displays the firing rates of a hypothetical neuron for 20 trials of the same experiment. The firing-rate intensity functions are distinct for the different trials. Let $P_r(t)$ denote the probability that the neuron spikes at time t on trial r , and $P(t)$ the probability that the neuron spikes at time t averaged across trials. Note that the PSTH provides a representation of $P(t)$ based on a limited set of trials. To estimate the within-trial spiking probability $P_r(t)$ we make use of two observations. First, $P_r(t)$ contains both the trial-averaged component $P(t)$ and a trial-specific component, which we might write as $g_r(t)$, with g_r standing for “gain factor.” Second, there is substantial information, across many trials, with which to estimate $P(t)$, whereas there is often very limited information to estimate the trial-specific gain factors. Figure 1B displays an average of estimates $\hat{P}_r(t)$ of $P_r(t)$ for $r = 8$, with 95% probability bands, obtained by 1) smoothing (filtering) the data for trial 8 to produce an initial estimate $\hat{P}_r(t)$, 2) smoothing the PSTH to obtain an estimate $\hat{P}(t)$, 3) smoothing further the residual $\hat{P}_r(t) - \hat{P}(t)$ to get the final estimate $\hat{P}_r(t)$ as the sum of $\hat{P}(t)$ from step 2 and the smoothed residual from step 3. This procedure is explained in greater detail in the next subsection. We used splines to do the smoothing, allowing a small number of degrees of freedom (large bandwidth) for step 3 but a larger number of degrees of freedom (smaller bandwidth) for steps 1 and 2.

If steps 2 and 3 are eliminated, the results are much worse. Figure 1, *C* and *D* displays averages of estimates after step 1 alone is applied, first with a small number of degrees of freedom, and then with a larger number of degrees of freedom. Note that using a small number of degrees of freedom will introduce excess bias,

whereas using a large number of degrees of freedom will introduce excess variance. The 95% probability bands in both *C* and *D* of Fig. 1 are much wider than those in Fig. 1B. The mean integrated squared errors (MISE¹) for the estimates in *B*, *C*, and *D* are 1,463, 7,125, and 5,519, respectively. This means, for example, that the estimates in *B* are almost $5,519 \div 1,463 \cong$ four times more efficient than the estimates in *D*, or equivalently, that the method in *D* would require, if it were possible, that we use four repeats of the same trial $r = 8$ to achieve the same accuracy as the estimates in *B*.

In Fig. 1A, the 20 spiking probability functions $P_r(t)$ are related to $P(t)$ according to

$$P_r(t) = g_r(t - \tau_r)P(t - \tau_r) \quad (1)$$

where τ_r is a latency that varies across trials. Thus the variability of the spike trains is attributed partly to the stochastic variation of point processes and partly to the extra variation caused by the gain factor g_r and shift (latency) τ_r , which might represent trial-dependent changes in the internal state of the subject. We have introduced the shift parameters τ_r for completeness. We do not discuss here the estimates of latency effects because of the availability of methods for estimating these parameters; they may be obtained, for example, by the procedure of Ventura (2004) as a preliminary step.

In the next subsection we address the statistical problem of recovering the factors $g_r(t)$. Because the spike train for a single trial is relatively sparse, it is important to reduce the dimensionality of the functional representation. To achieve that, we introduce a principal-component decomposition of the functional variation $g_r(t)$, which serves in part to reduce dimensionality and in part to provide additional interpretation of the variability.

We show how significance tests may be used to examine whether a constant-excitability model is appropriate, and whether any additional variability may be summarized by one or two parameters, in terms of the principal components of the variability.

¹ If $\hat{f}(t)$ is an estimate of $f(t)$, the MISE of \hat{f} is $E \{ \int [\hat{f}(t) - f(t)]^2 dt \}$, where the expectation, which is with respect to the data used to produce $\hat{f}(t)$, was approximated by a sample average across many simulated samples.

One purpose of analyzing trial-to-trial variation is to examine its effect on time-varying dependency between two neurons. In the remaining subsections we extend the methods of Ventura et al. (2005a) to incorporate trial-to-trial variation; we show how to adjust the cross-correlogram and JPSTH for trial-to-trial variation; and we show how significance tests may be used to produce evidence that two neurons have shared input in the form of shared trial-to-trial variation.

Equation 1 is the basis for the methodology developed here. Firing rate functions that depend on t are valid for Poisson processes only. For non-Poisson processes, a conditional intensity $P(t - \tau_r | H_{rt})$ must replace $P(t - \tau_r)$ in Eq. 1, as in Ventura et al. (2005a), where H_{rt} is the history of trial r up to time t . An example of a simulated non-Poisson neuron is treated in RESULTS, *Simulated data: non-Poisson spike trains*.

Low-dimensional representation of time-varying excitability effects

If Eq. 1 is believed to apply, so that the firing rate of a neuron is different across trials, an estimate of $P(t)$, generally the PSTH, will not provide a complete description of the firing properties of the neuron. One standard way to proceed is to fit, by Gaussian filtering or by other means, the firing-rate functions on all trials separately. Because one trial has few spikes, at least relative to the number of spikes available to smooth a PSTH, smoothing each trial separately will produce highly variable estimates. We propose a more efficient alternative.

We first obtain a smoothed estimate $\hat{P}(t)$ of $P(t)$ by smoothing the PSTH. We prefer a spline-based method called BARS (Bayesian adaptive regression splines) (DiMatteo et al. 2001; Kass et al. 2003) for this purpose, but many alternative smoothing techniques would work well for typical data sets. The advantage of BARS is that it determines automatically the optimal number and placement of knots. An analogous procedure, were Gaussian filtering used, would be to determine the optimal time-varying bandwidth.

The second step is to apply binary regression to the data from each trial separately to estimate $P_r(t)$. Binary regression, within the framework of generalized linear models (McCullagh and Nelder 1990), is the equivalent of the usual regression framework, but for binary data, here spike trains transformed to a binary sequence with 1 coding for a spike at t , and 0 coding for no spike. Although BARS could also be applied here to estimate $P_r(t)$, we found that our procedure was not sensitive to knot number and placement. We therefore obtain smooth estimates of all $P_r(t)$ by using regression splines with the same knots used to smooth $P(t)$. Letting $\hat{P}_r(t)$ denote the estimates thus obtained, then according to Eq. 1, estimates of $g_r(t)$ are $\hat{g}_r(t) = \hat{P}_r(t)/\hat{P}(t)$.

Because there are often relatively few spikes per trial, there are relatively few data for fitting $\hat{P}_r(t)$, so that the deviations $\hat{g}_r(t)$ are highly variable. To reduce their variability, our next step is to identify a suitable low-dimensional representation of these deviations. This could be done very simply, for example, by smoothing further the deviations $\hat{g}_r(t)$, relying on splines having a small number of knots (e.g., one to three knots), or equivalently on Gaussian filtering with a large bandwidth. According to this scheme the trial-to-trial variability is allowed to be curvilinear, but is constrained to be slowly varying, across peristimulus time. The resulting smoothed deviations $\tilde{g}_r(t)$ are less variable than $\hat{g}_r(t)$ and thus yield less variable estimates of $P_r(t)$, as demonstrated in Fig. 1.

Our preferred method, however, is a low-dimensional principal-components representation, where the trial-to-trial fluctuation in firing-rate intensity will be assumed to have the form

$$\log g_r(t) = w_{0r} + \sum_{j=1}^J w_{jr} \phi_j(t) \tag{2}$$

where the $\phi_j(t)$ functions are suitably chosen curves. A special case of Eq. 2 is the *constant excitability* model, in which $\log g_r(t) = w_{0r}$, discussed, for example, in Brody (1999a,b). Note that we model $\log g_r(t)$ rather than $g_r(t)$ as a linear combination of the principal components to ensure that the resulting fitted $g_r(t)$ is a positive function. By restricting the form and number J 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 available per trial. As discussed in the following text, the $\phi_j(t)$ functions will be taken to be principal components of the trial-to-trial variation. The set of weights $w_r = (w_{0r}, w_{1r}, w_{2r}, \dots, w_{Jr})$ then describes the variability specific to trial r more simply and succinctly than would smooth functions of time.

To fit Eq. 2, we begin with the vectors of values of $\tilde{g}_r(t)$ evaluated along the time axis: if there are time points t_1, \dots, t_{\max} , then for each r we obtain the vector $\tilde{g}_r(t_1), \dots, \tilde{g}_r(t_{\max})$. From these R vectors we compute the principal components of their covariance matrix $\tilde{\Sigma}_g$. Let us write these R principal components as $\phi_1(t), \dots, \phi_R(t)$. We then need to determine how many principal components to use, say J , and obtain the coefficients $\{w_{jr}\}$ in Eq. 2. 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 again performed, except that $g_r(t)$ in Eq. 1 is replaced by $\exp[w_{0r} + \sum_{j=1}^J w_{jr} \phi_j(t)]$. Stepwise significance tests are used to determine how many of the principal components J are actually needed to represent the trial-to-trial variability. Specifically, to determine whether the contribution of some of the principal components is negligible, sequentially nested regression models should be considered: first the regression model Eq. 1 with $g_r(t) = 1$, then the model involving trial-dependent constants with $g_r(t) = \exp(w_{0r})$, then the model involving trial-dependent constants and the first principal component with $g_r(t) = \exp[w_{0r} + w_{1r} \phi_1(t)]$, and so forth, up to the model involving the constants, and all principal components. These models can be fit sequentially and the deviance difference compared with 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.

All smooth fits are 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., Hastie and Tibshirani 1990; McCullagh and Nelder 1990). As we showed in Ventura et al. (2005a), spline-based generalized linear models can also be used while taking account of non-Poisson firing behavior. Such methods are applied to the simulated non-Poisson spike trains in RESULTS.

Bootstrap excursion test of time-varying dependency

Suppose two neurons labeled 1 and 2 are recorded simultaneously across multiple trials, and the task is to assess the time-varying dependency between their spike trains. Ventura et al. (2005a) defined the quantity

$$\zeta_\delta(t) = \frac{P^{12}(t, t + \delta)}{P^1(t)P^2(t + \delta)} \tag{3}$$

to analyze the excess joint spiking probability above that expected assuming independence, where $P^{12}(t, t + \delta)$ is the probability that neuron 1 will spike at time t and neuron 2 will spike at time $t + \delta$, and $P^1(t)$ and $P^2(t)$ are the firing rates for neurons 1 and 2. Ventura et al. (2005a) then demonstrated the use of a Bootstrap procedure to assess any apparent deviation of $\zeta_\delta(t)$ from 1, which would suggest extra

TABLE 1. Firing rates P_r given by Eq. 1 for Poisson neurons A (no trial-to-trial variation), B (constant trial-to-trial variation), and C (time-varying trial-to-trial variation)

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

Non-Poisson neuron D has conditional firing rate intensity $P_r(t, s) = g_r(t)P(t, s) = g_r(t)\lambda_1(t)\lambda_2(s)$, where s is the time elapsed since the last spike before t (see *Simulated data: non-Poisson spike trains*). $f(t, a, b)$ denotes the normal density function with mean a and SD b . “ \sim ” means randomly generated from. The firing rates of all trials are shown in Fig. 2.

correlation between the two neurons above the correlation induced by modulations in the firing rate.

When excitability and/or latency effects are present, we want to assess whether the neurons are correlated above what is expected from modulations in the firing rates, and from the correlation induced by trial-to-trial variability effects. To take account of trial-to-trial variability, the procedure is modified by replacing Eq. 3 with

$$\zeta_{\delta}(t) = \frac{P_r^{12}(t, t + \delta)}{P_r^1(t)P_r^2(t + \delta)} \quad (4)$$

where $P_r^{12}(t, t + \delta)$ is the probability for the r th trial that neuron 1 will spike at time t and neuron 2 will spike at time $t + \delta$, and P_r^1 and P_r^2

are the firing rates for neurons 1 and 2 on trial r . Equation 4 assumes that the excess joint spiking activity does not itself depend on the trial. To obtain an estimate of $\zeta_{\delta}(t)$, we again replace $P_r^1(t)$ and $P_r^2(t + \delta)$ with $\hat{P}_r^1(t)$ and $\hat{P}_r^2(t + \delta)$ as estimated earlier in *Low-dimensional representation of time-varying excitability effects*, and $P_r^{12}(t + \delta)$ with a smooth estimate of the joint spike count on trial r , obtained with BARS or another suitable smoothing method. The estimate of $\zeta_{\delta}(t)$ is then taken to be the average of the estimates of Eq. 4 over trials r .

The Bootstrap significance test now proceeds as in Ventura et al. (2005a). In short, it is based on a computation of null bands for $\zeta_{\delta}(t)$, obtained under the null hypothesis of independence. When $\hat{\zeta}_{\delta}(t)$ is above or below these bands there is *potential* evidence of excess or

Poisson neurons A, B, and C

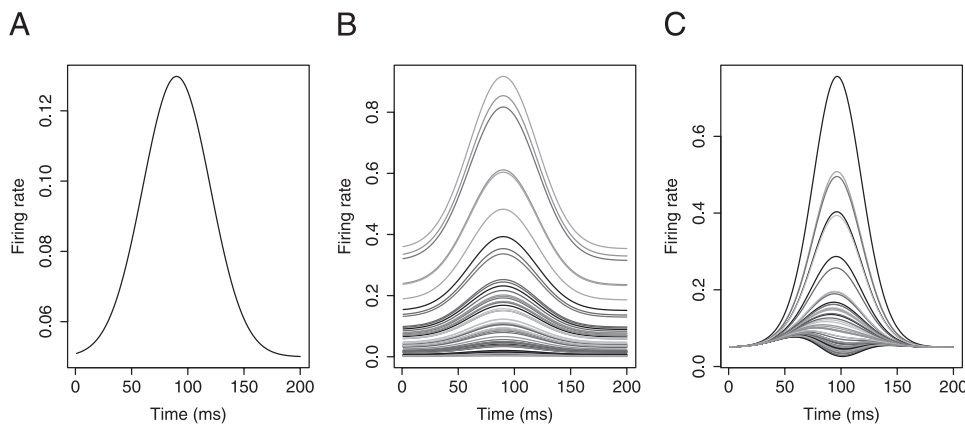


FIG. 2. A, B, and C: true firing rates $P_r(t) = g_r(t)P(t)$ for each trial of simulated neurons A, B, and C described in Table 1. For neuron D, we show $g_r(t)\lambda_1(t)$ for all trials $r = 1 \dots R$ in D, and $\lambda_2(s)$ in E, as described in Table 1 and *Simulated data: non-Poisson spike trains*.

Non-Poisson neuron D

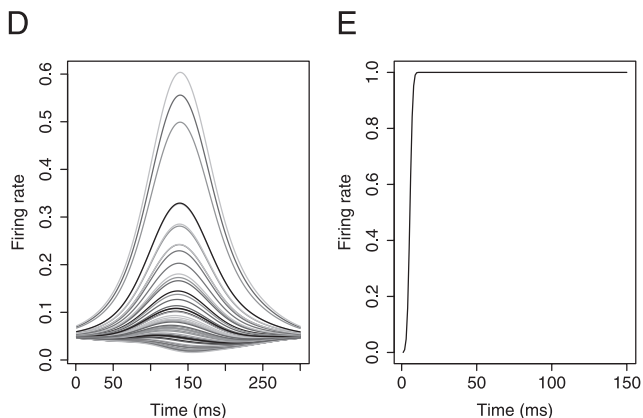


TABLE 2. Deviances and significance tests *P*-values for Poisson neurons A, B, and C, and non-Poisson neuron D to test form of excitability effects

Model	df	Deviance (<i>P</i> -value)			
		A	B	C	D
No excitability	0	1,525*	3,090	2,205	2,983
Constant excitability	<i>R</i>	1,482 (0.94)	1,572* ($<10^{-5}$)*	2,018 ($<10^{-5}$)	2,920 (0.001)
1 PCA excitability	2 <i>R</i>	1,420 (0.5)	1,515 (0.59)	1,862* ($<10^{-5}$)*	2,872* (0.034)*
2 PCAs excitability	3 <i>R</i>	—	—	1,801 (0.44)	2,872 (1)

No excitability model $P(t)$ versus constant excitability model $\exp(w_{or})P(t)$ versus time-varying single principal-component model (1 PCA) $\exp[w_{or} + w_{1r}\phi_1(t)]P(t)$ versus time-varying two principal-components model (2 PCAs) $\exp[w_{or} + w_{1r}\phi_1(t) + w_{2r}\phi_2(t)]P(t)$. $R = 60$ trials for neurons A, B, and C, and $R = 32$ trials for neuron D. “*” indicates the model of choice.

diminished joint spiking activity. If we were assessing a joint spiking activity at a single value of time t , and if $\hat{\zeta}_\delta(t)$ were above or below the null bands then there would be some evidence of excess or diminished joint activity at time t (and lag δ). However, because we are examining a large number of time values, a global assessment is needed.

The only difference between the Bootstrap test of Ventura et al. (2005a) and this bootstrap test is in the sampling of Bootstrap samples. Bootstrap samples should be stochastically similar to the observed spike trains, but they should also conform to the hypothetical reality imposed by the hypothesis under investigation. Here, we are testing whether the neurons are independent above the dependency induced by modulations of firing rates and trial-to-trial variability effects. Therefore Bootstrap samples must contain excitability effects comparable to those in the observed spike trains. The modified Bootstrap procedure is as follows.


- 1) Simulate R trials of spike trains for the two neurons as follows:
 - a) Sample at random and with replacement R numbers from the integers $1, \dots, R$.
 - b) Simulate R pairs of spike trains with firing-rate functions $\hat{P}_r^1(t)$ and $\hat{P}_r^2(t)$, where r runs through the set of trials sampled in step 1a. This is a bootstrap sample.
- 2) Obtain a smooth estimate $\hat{\zeta}_\delta(t)$ as described above but based on this bootstrap sample rather than on the observed data, for each δ of interest.
- 3) Repeat steps 1 and 2 N times to get N estimates $\hat{\zeta}_\delta(t)$.
- 4) For each time t , define $h_L(t)$ and $h_U(t)$ to be the 0.025 and 0.975 quantiles of the N values $\hat{\zeta}_\delta(t)$.

To determine the bands accurately h_L and h_U , we have used bootstrap sample sizes of $N = 1,000$. To reduce the computation effort, a normal approximation can be used with $N \cong 50$, as described in Ventura et al. (2005a).


Note that step 1 is an implementation of a *parametric* bootstrap. Because the trials are not exchangeable as a result of trial-to-trial variability effects, there is no nonparametric bootstrap alternative.

To perform the test of independence, we define G_{obs} to be the largest area of any contiguous portion of $\hat{\zeta}_\delta(t)$ that exceeds the bands defined above. [See Ventura et al. (2005a) for a mathematical definition of G_{obs} .] Our G_{obs} is similar in spirit to k' of Ellaway and Murphy (1985).] To calculate its Bootstrap *P*-value, let $\hat{\zeta}_\delta^{(n)}(t)$, $n = 1, \dots, N$ stand for the estimate of $\zeta_\delta(t)$ obtained from the n th bootstrap sample,

Neuron B



Neuron C



Neuron D




FIG. 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. Fits for all trials can be found in the APPENDIX figures.

and $G_{boot}^{(n)}$ denote the largest area of any contiguous portion of $\hat{\zeta}_\delta^{(n)}(t)$ that exceeds the bands defined above. We then calculate the *P*-value

$$P = \frac{\text{Number of Bootstrap samples for which } G_{boot}^{(n)} > G_{obs}}{N + 1}$$

as is standard for Bootstrap tests, and reject the hypothesis if independence between the two neurons if P is small, say $P < 0.05$.

Adjustment of JPSTH and cross-correlogram for trial-to-trial variability

The joint peristimulus time histogram (JPSTH) and cross-correlogram provide useful visual representations of the correlated activity

TABLE 3. Firing rates $P_r(t)$ given by Eq. 1 for neuron pairs E1/E2 and F1/F2

Neuron Pair	Excitability Effect $g_r(t)$	Global Rate $P(t)$
E1/F1	$g_r^1(t) = g_r^2(t) = 1 + c_r f(t, 390, 35)$	$P^1(t) = 0.04 + 24f(t, 390, 40)$
E2/F2	$c_r = b_r - \bar{b}$, $b_r \sim \text{Gamma}(1, 0.5)$	$P^2(t) = 0.04 + 24f(t, 390, 60)$
Common latencies: $\tau_r \sim f(t, 0, 04)$		
$\zeta_0(t) = 1 + 15f(t, 380, 30)$		

$f(t, a, b)$ denotes the normal density function with mean a and SD b . “~” means randomly generated from.

for a pair of neurons. To take account of firing-rate variation in the neurons, Aertsen et al. (1989) proposed a normalized version of the JPSTH. Both the normalized JPSTH and the cross-correlogram assume stationarity across trials. In the presence of excess trial-to-trial variability, this may be violated. Once we have estimates of the trial-specific firing rates $P_r^i(t)$, however, it is straightforward to adjust the JPSTH and cross-correlogram to remove the general excitability effects. The corrected JPSTH of Aertsen et al. (1989) has $(t, t + \delta)$ pixel equal to

$$J_\delta(t) = P^{12}(t, t + \delta) - P^1(t)P^2(t + \delta) \quad (5)$$

with estimate $\hat{J}_\delta(t)$ taken to be the value of $J_\delta(t)$ when the spiking probabilities $P^1(t)$, $P^2(t)$, and joint spiking probability $P^{12}(t, t + \delta)$ are replaced with their observed-data counterparts, i.e., the spike counts divided by the number of trials (and the pixel width). Large values of $|\hat{J}_\delta(t)|$ are evidence that the two neurons are correlated at lag δ .

When excitability effects are present, values of $|\hat{J}_\delta(t)|$ are inflated, which may lead to the false claim that the two neurons are dependent. For values of $|\hat{J}_\delta(t)|$ close to zero to still suggest independence, we define the $(t, t + \delta)$ pixel of the corrected JPSTH to be

$$J\delta(t) = R^{-1} \sum_r J_{r,\delta}(t) \quad J_{r,\delta}(t) = P_r^{12}(t, t + \delta) - P_r^1(t)P_r^2(t + \delta)$$

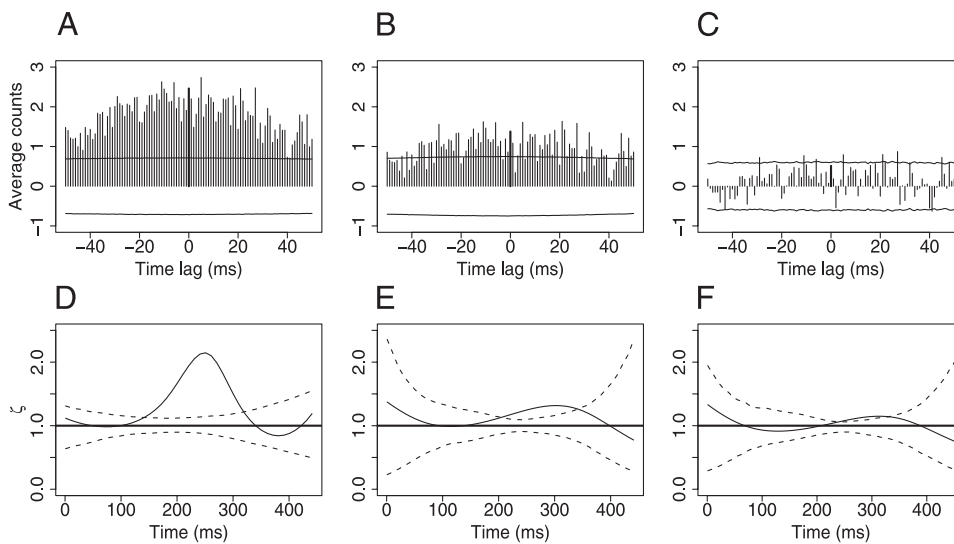
where $J_{r,\delta}(t)$ is the JPSTH based only on trial r , with estimate $\hat{P}_r^{12}(t, t + \delta) - \hat{P}_r^1(t)\hat{P}_r^2(t + \delta)$.

A related adjustment has been used by Baker et al. (2001). The corrected cross-correlogram at lag δ may be then defined as

$$C(\delta) = \sum_r J_\delta(t) \quad (6)$$

Note that Eq. 6 is not normalized, as in Aertsen et al. (1989), although it could be without substantially changing the methodology. Figures 4

Neurons E1/E2 : no synchrony



Neurons F1/F2 : synchrony

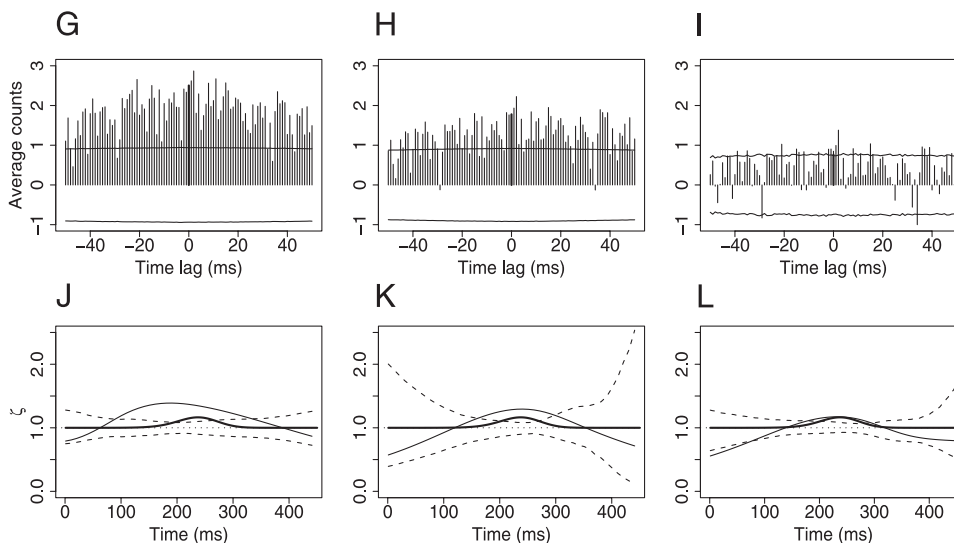


FIG. 4. Neuron pairs with latency and excitability effects, and no synchrony (E1/E2) or synchrony (F1/F2). Shuffle corrected cross-correlogram before correction for latency and excitability (A, G), after correction for latency (B, H), and after correction for latency and excitability (C, I), with 95% Bootstrap bands. (D, J), (E, K), and (F, L): corresponding panels for $\zeta_0(t)$ of the joint spiking model [bold line is the true $\zeta_0(t)$; solid line is the estimate], with 95% bootstrap bands; the P -values for the excursion of $\hat{\zeta}_0(t)$ outside the bands are in Table 5.

TABLE 4. Deviances and P-values, for neuron pairs E1/E2 and F1/F2, of the tests that determine whether the two neurons in the pair have the same, proportional, or different excitability effects

Model	df	Deviance (P-value)	
		Neuron Pair	
		E1/E2	F1/F2
1 PCA models			
Same excitability	2R	—	4,447* ($<10^{-5}$)*
Proportional excitability	2R + 1	—	4,447 (0.96)
Different excitability	4R	—	4,344 (0.10)
2 PCAs models			
Same excitability	3R	4,270* ($<10^{-5}$)*	—
Proportional excitability	3R + 1	4,270 (0.96)	—
Different excitability	6R	4,232 (1)	—

The two neurons of the pair are fitted together with models: same excitability model $\exp[w_{0r} + w_{1r}\phi_1(t)]P_1(t)$ and $\exp[w_{0r} + w_{1r}\phi_1(t)]P_2(t)$ versus proportional excitability model $\exp[w_{0r} + w_{1r}\phi_1(t)]P_1(t)$ and $\exp[w_{0r} + w_{1r}\phi_1(t)]P_2(t)\alpha$ versus different excitability model $\exp[w_{0r}^1 + w_{1r}^1\phi_1^1(t)]P_1(t)$ and $\exp[w_{0r}^2 + w_{1r}^2\phi_1^2(t)]P_2(t)$ [for this last model, $\exp(w_{0r}^1)$ and $\exp(w_{0r}^2)$, $\exp(w_{1r}^1)$ and $\exp(w_{1r}^2)$, and $\phi_1^1(t)$ and $\phi_1^2(t)$ are not constrained to be equal]. The same sequences of models are fitted for the two PCAs model. Here $R = 60$, the number of trials for one neuron. “*” indicates the model of choice.

and 6 show instances of cross-correlograms adjusted for excitability effects.

Trial-to-trial variation shared across neurons

In previous subsections the fitted firing rate functions $\hat{P}_r^1(t)$ and $\hat{P}_r^2(t)$ were not restricted to have any relationship to each other, although it is plausible that the two neurons might have either the same trial-to-trial effects or proportional trial-to-trial effects arising from common inputs. These possibilities may also be examined within the framework we have developed here, using standard generalized linear model methodology.

We will write the equality and proportionality cases as $g_r^1(t) = g_r^2(t)$ and $g_r^1(t) = \alpha \cdot g_r^2(t)$ for all t , where α is a scalar constant. A further special case is the constant excitability model $g_r^1(t) = \exp(w_{0r}^1)$, where w_{0r}^1 is a constant, for which equality and proportionality become $w_{0r}^1 = w_{0r}^2$ and $w_{0r}^1 = \log \alpha + w_{0r}^2$. As explained in many texts on linear regression (the generalized linear model context being analogous), fitting these models and testing them against one another is straightforward. Such tests are performed in the next section, with results recorded in Table 4.

RESULTS

In the next two subsections we illustrate our methods for fitting excitability effects based on simulated Poisson and non-Poisson spike trains. We then illustrate how to correct synchrony detection plots and measures, and how to carry out bootstrap inferences.

Simulated data: Poisson spike trains

We simulate Poisson spike train data with firing rate $P_r(t)$ given by Eq. 1 under three scenarios: no trial-to-trial variation

(neuron A), constant trial-to-trial variation (neuron B), and time-varying trial-to-trial variation (neuron C). Specifics of the firing rates are summarized in Table 1 and shown in Fig. 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 example, for neuron A, we fitted first model Eq. 1 with $g_r(t) = 1$, then with the trial-dependent constants $g_r(t) = \exp(w_{0r})$. The deviance difference between the two model fits (1,525 – 1,482 = 43) is compared with 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 0.05 we would conclude, correctly, that no trial-to-trial variation is present for neuron A. For neuron B the results indicate strong evidence for trial-to-trial variability ($P < 10^{-5}$), but no evidence for time-dependent trial-to-trial variability ($P = 0.59 > 0.05$), 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, Fig. A2. 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 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 ($P > 0.05$). 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, Fig. A3.

Simulated data: non-Poisson spike trains

Data were also simulated from a non-Poisson neuron (neuron D) that followed an IMI model (Kass and Ventura 2001), for which the conditional firing rate intensity depends not only on time t , as for Poisson processes, but also on $s_{*}(t)$, the time of the last spike previous to time t , according to $P[t, s_{*}(t)] = \lambda_1(t)\lambda_2[s_{*}(t)]$. Neuron D had $R = 32$ trials and 300 time bins. Figure 2 shows the functions $g_r(t)\lambda_1(t)$ for all 32 trials, and $\lambda_2(s)$; specifics are in Table 1.

The estimation procedure discussed earlier in *Low-dimensional representation of time-varying excitability effects* was followed (as in the Poisson case), with deviances and P-values recorded in Table 2, and the resulting fitted firing rate functions

TABLE 5. P-values of the excursion of $\hat{\zeta}_0(t)$ outside the Bootstrap bands for the two pairs of neurons described in Table 3 and shown in Fig. 4

Model Corrected for Modulations in	P-Values for $\zeta_0(t)$	
	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	0.031




FIG. 5. Raster plot and peristimulus time histograms (PSTHs) for (A–B) neuron 380506.s and (C–D) neuron 380506.u.

for six representative trials are displayed in Fig. 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, Fig. A4.

Simulated data: adjustment of dependency 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 (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 Fig. 4. The firing rates of these four neurons are not shown because they are similar to neuron C in Fig. 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 values $P < 10^{-5}$. We estimated the latency effects as in Ventura (2004), and shifted the trials according to the latency estimates.

The methods of subsection *Low-dimensional representation of time-varying excitability effects* were then 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 *Trial-to-trial variation shared across neurons* 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-correlograms and the fitted $\hat{\zeta}_0(t)$ for 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 Fig. 4B, when the excitability effect still exists, most of the cross-correlogram values are still outside the 95% bootstrap band. After both the latency and excitability effects are corrected in Fig. 4C, the cross-correlogram indicates that

the two neurons are independent. Figure 4E shows the estimated $\zeta_0(t)$ of the joint spiking model in Eq. 4. Before being corrected for the latency effect, the estimated $\zeta_0(t)$ has a peak outside the Bootstrap band. The area outside the band has P -value of <0.0001 . After the latency effect only is screened out, the magnitude of $\hat{\zeta}_0(t)$, shown in Fig. 4E, is reduced but 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 $\hat{\zeta}_0(t)$, shown in Fig. 4F, is further reduced and the largest area outside the bands has a P -value of 0.63. This leads correctly to the conclusion that the two neurons are independent. The same procedure also yields the correct conclusion for the dependent pair F1/F2.

Application to a pair of VI 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.s and 380506.u, 90 deg spatial phase). The data consist of 64 trials shown in Fig. 5. The stimulus, identical for all trials, consisted of a standing sinusoidal grating that appeared at time 0, and disappeared at 237 ms.

Figure 6A, which displays the rate-adjusted cross-correlogram with bins 2.8 ms wide,² 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 $2 \times 2.8 = 5.6$ ms, $\hat{\zeta}_{5.6}(t)$ is displayed in Fig. 6C, along with 95% Bootstrap bands; $\hat{\zeta}_{5.6}(t)$ exceeds 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

² The choice of 2.8 ms is based on 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 a time lag <2.8 ms, we analyze time lags >2.8 ms or < -2.8 ms.




FIG. 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. Dashed curves are pointwise 95% probability bands computed under the null hypothesis of independence. C: time-varying synchrony curve $\hat{\zeta}_{5.6}(t)$ before (thin solid curve) and after (bold curve) the adjustment for time-varying excitability effects. Dashed curves are pointwise 95% Bootstrap bands computed under the null hypothesis $\hat{\zeta}_{5.6}(t)=1$ for all t .

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 a result of 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 of 0.096 and 0.174 for the two neurons, respectively. However, the tests 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 nonconstant excitability model ($P < 0.00001$) against the constant excitability model ($P = 0.752$) for different vs. same excitability effects). Applying the sequential fitting procedure we then determined that two principal components should be used to capture the nonconstant excitability effects, with each principal component explaining 83 and 14%, respectively, of the variability in the observed pairs of spike trains. Therefore the firing rate of neuron $i = 1, 2$ on trial r , $P_r^i(t)$, can be summarized as

$$P_r^i(t) = P^i(t) \exp[w_{r0} + w_{r1}\phi_1(t) + w_{r2}\phi_2(t)] \quad (7)$$

where the two principal components $\phi_1(t)$ and $\phi_2(t)$ are displayed in Fig. 7, A and B.

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 $-5w_0$, so that the firing rate model in Eq. 7 is close to

$$P_r^i(t) = P^i(t)e^{w_0} \exp[\phi(t)] \quad (8)$$

where $\phi(t) = 5 - \phi_1(t) - \phi_2(t)$ is also plotted in Fig. 7C. That is, the firing rate for a particular trial is modulated by the function in Fig. 7C, times a gain factor.

Figure 6B shows the corrected cross-correlogram with 95% Bootstrap bands. We conclude that the correlation apparent in the uncorrected cross-correlogram (Fig. 6A) is attributed entirely to excitability effects. Figure 6C displays the estimates of $\hat{\zeta}_{5.6}(t)$ (corresponding to the second bin to the right of 0 in the cross-correlograms) before and after the adjustment for excitability. Before excitability adjustment, the P -value for synchrony at lag 5.6 ms is $< 10^{-5}$, whereas after the adjustment $\hat{\zeta}(t)$ lies within the Bootstrap bands. There is thus no evidence of synchrony at that lag, beyond the correlations induced by rate and excitability modulations. We reached similar conclusion at all lags; Fig. A5 in the APPENDIX displays the same as Fig. 6C for all lags.




FIG. 7. Shared excitability effects for 2 neurons from Aronov et al. (2003). First (A) and second (B) principal components, $\phi_1(t)$ and $\phi_2(t)$. C: function $\phi(t) = 5 - \phi_1(t) - \phi_2(t)$ in Eq. 8. D: plot of the coefficients w_1 against w_2 . Each point in the plot represents a coefficient pair on a particular trial. E: plot of $(w_1 + w_2)/2$ against w_0 . Solid lines in D and E have intercept 0 and slopes 1 and 5, respectively.

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.

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 1) are present, 2) are constant or time-varying, 3) are shared by two neurons, and 4) are adequately described by a small number of parameters. In addition, the methodology allows the trial-to-trial variability to be removed, so that assessments of correlation and synchrony, or lagged time-locked firing, may be suitably tested. 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 Eq. 1 statistical models and their non-Poisson generalization allow excitability effects to be removed so that correlation and synchrony assessments may be adjusted, and the restriction in Eq. 2 reduces the noise in the fitting, making the inference procedures more powerful.

One simplification of the method used here is that we have ignored variability introduced by first smoothing $P(t)$ and $P_r(t)$ when we subsequently make statistical inferences based on the coefficients of the principal components. 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).

Kass and Ventura (unpublished observations) show that excitability effects described by Eq. 1 produce spike count correlations that grow as the spike count interval increases. For realistic firing rates spike count correlations that are very small on timescales of tens of milliseconds become nonnegligible on scales of hundreds of milliseconds. This may help explain correlations that have been reported and discussed (Shadlen and Newsome 1998).

Excess trial-to-trial variability 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.

The methods presented here have 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.

GRANTS

This work was supported by National Institutes of Health Grants MH-64537 to C. Cai and R. E. Kass and MH-64537 and N01-NS-2-2346 to V. Ventura.

APPENDIX: SUPPLEMENTARY FIGURES

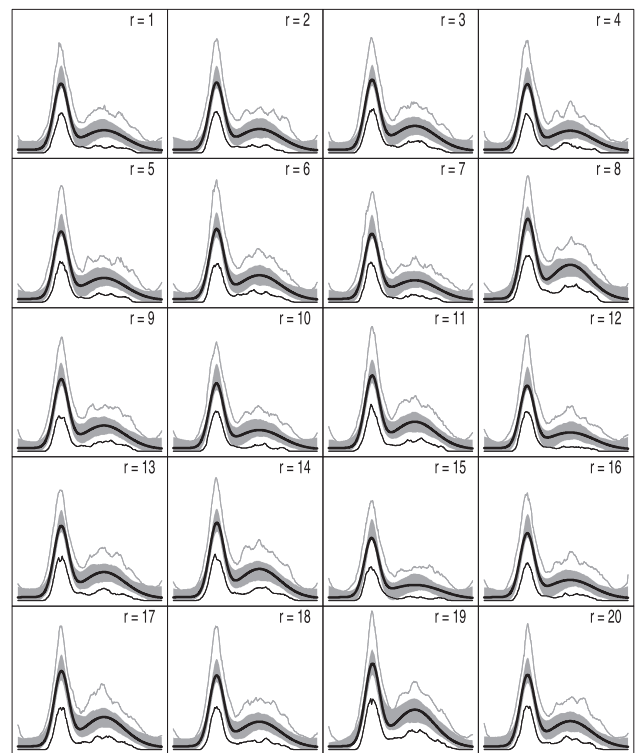


FIG. A1. 95% simulation bands for the estimates of $P_r(t)$, $r = 1, \dots, 20$ for the illustration of METHODS. Solid bands are for $P_r(t)$ smoothed separately for all r with spline knots as in Fig. 1D, shaded areas are for our smoothing procedure described in *Low-dimensional representation of time-varying excitability effects*. Bold lines are the true $P_r(t)$. All panels use the same axes ranges.




FIG. A2. Neuron B: true (bold curves) and fitted (thin curves) firing rates for all simulated trials.




FIG. A3. Neuron C: true (bold curves) and fitted (thin curves) firing rates for all simulated trials. Dashed curves are the fits from the constant excitability model that was rejected by the deviance test.




FIG. A4. Neuron D: true (bold curves) and fitted (thin curves) firing rates for all simulated trials.




FIG. A5. 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 to 33.6 ms in 2.8 -ms steps. Synchrony plot for $\delta = 0$ was omitted because of the limitations of the recording of joint spikes.

REFERENCES

Aersten A, Gerstein G, Habib M, and Palm G. Dynamics of neuronal firing correlation—modulation in effective connectivity. *J Neurophysiol* 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 Comput* 13: 1351–1377, 2001.

Baker SN, Spinks R, Jackson A, and Lemon RN. Synchronization in monkey motor cortex during a precision grip task: I. Task-dependent modulation in single-unit synchrony. *J Neurophysiol* 85: 869–885, 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.

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

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

Brown EN, Barbieri R, Ventura V, Kass RE, and Frank LM. A note on the time-rescaling theorem and its implications for neural data analysis. *Neural Comput* 14: 325–346, 2002.

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

Ellaway PH and Murthy KSK. The source and distribution of short-term synchrony between gamma-motoneurons in the cat. *Q J Exp Physiol Cogn Med Sci* 70: 233–247, 1985.



- Grün S, Riehle A, and Diesmann M.** Effect of across trial non-stationarity on joint-spike events. *Biol Cybern* 00: 000–000, 2003.
- 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 & 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 Brown EN.** Statistical issues in the analysis of neuronal data. *J Neurophysiol* 94: 8–25, 2005.
- Kass RE, Ventura V, and Cai C.** Statistical smoothing of neuronal data. *Network: Comput Neural Syst* 14: 5–15, 2005.
- McCullagh P and Nelder JA.** *Generalized Linear Models* (2nd ed.). New York: Chapman & Hall, 1990.
- Nawrot M, Aertsen A, and Rotter S.** Single-trial estimation of neuronal firing rates. *J Neurosci Methods* 94: 81–91, 1999.
- Pauluis Q and Baker SN.** An accurate measure of the instantaneous discharge probability, with application to unitary joint-event analysis. *Neural Computation* 12: 647–669, 2000.
- Shadlen MN and Newsome WT.** The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18: 3870–3896, 1998.
- Ventura V.** Testing for, and estimating latency effects for Poisson and non-Poisson spike trains. *Neural Comput* 16: 11: 2323–2350, 2004.
- Ventura V, Cai C, and Kass RE.** Statistical assessment of time-varying dependency between two neurons. *J Neurophysiol* 94: 2940–2947, 2005.
- Woody CD.** Characterization of an adaptive filter for the analysis of variable latency neuroelectric signals. *Med Biol Eng* 5: 539–553, 1967.