

Published in final edited form as:

Neural Comput. 2012 August ; 24(8): 2007–2032. doi:10.1162/NECO_a_00307.

A Framework for Evaluating Pairwise and Multi-way Synchrony Among Stimulus-Driven Neurons

Ryan C. Kelly^{1,2} and Robert E. Kass^{1,2}

¹Department of Statistics, Carnegie Mellon University

²The Center for the Neural Basis of Cognition, Carnegie Mellon University

Abstract

Several authors have discussed previously the use of loglinear models, often called maximum entropy models, for analyzing spike train data to detect synchrony. The usual loglinear modeling techniques, however, do not allow for time-varying firing rates that typically appear in stimulus-driven (or action-driven) neurons, nor do they incorporate non-Poisson history effects or covariate effects. We generalize the usual approach, combining point process regression models of individual-neuron activity with loglinear models of multiway synchronous interaction. The methods are illustrated with results found in spike trains recorded simultaneously from primary visual cortex. We then go on to assess the amount of data needed to reliably detect multiway spiking.

Keywords

functional connectivity; loglinear models; multiple spike train analysis

1 Introduction

Synchrony is widely believed to play a fundamental role in neural computation (e.g. Uhlhaas et al., 2009), but its statistical assessment is subtle (for reviews see Grün (2009) and Harrison et al. (2011)). To analyze multi-way synchrony among simultaneously-recorded multiple spike trains that are represented as binary time series, across many trials, it is natural to consider well-established loglinear modeling technology (Gütig and Aertsen, 2003; Martignon et al., 2000; Nakahara and Amari, 2002; Schneidman et al., 2006). The standard approach, however, has two shortcomings. First, it assumes stationarity of firing rates across suitable time intervals. Second, it does not incorporate spiking history, or other covariates, and therefore effectively assumes Poisson spiking. While previous authors have been aware of these issues (e.g., Martignon et al. (2000)) they have not offered specific methods for dealing with them. In addition, the standard approach ignores the inherent relative sparsity of two-way and higher-order synchronous spiking. Here we provide a modification of the usual loglinear modeling methodology to deal with the case of inhomogeneous non-Poisson firing rates that result from stimulus-driven (or action-driven) recordings, and we give a straightforward procedure for parameter estimation, which is a variant of maximum likelihood via iterative proportional fitting. Tests for pairwise and multiway synchrony may then be based on the bootstrap.

To explain more fully the problem we are trying to solve, and the idea behind our proposed solution, let us suppose we have simultaneously-recorded spike trains from some particular

3 neurons, which we label (1, 2, 3), for which spiking patterns take the form (a, b, c) , where $a, b, c \in \{0, 1\}$ so that, for instance, (1,1,0) would signify that the first two neurons fired but the third did not. The probability of obtaining the pattern (a, b, c) at some time t on trial r may depend on the history $H_{t,r}$ of spiking patterns prior to time t and it may also depend on some other covariates (such as a measurement of network activity, as in Kass et al. (2011), or variables representing trial-to-trial variation, as in Ventura et al. (2005a)), which we here take to be a vector $u_{t,r}$. We then write the probability of pattern (a, b, c) at time t on trial r as $p_{abc,r}^{123}(t|H_{t,r}, u_{t,r})$. If we were to treat the spike trains as jointly stationary, and ignore history, covariates, and additional sources of trial-to-trial variation, we could omit t and r and write the usual model for 2-way, but not 3-way interaction as

$$\log p_{abc} = \alpha + a\beta_a^1 + b\beta_b^2 + c\beta_c^3 + ab\beta_{ab}^{12} + ac\beta_{ac}^{13} + bc\beta_{bc}^{23} \quad (1)$$

where the parameters satisfy certain constraints (e.g., $\sum p_{abc} = 1$, $\sum \beta_a^1 = 0$, etc.). In computer science and physics models such as (1) are often called *maximum entropy* models (Schneidman et al., 2006). In the statistics literature the parameters are usually standardized by subtracting means (Agresti, 2002). Using (1), for a given set of data the loglikelihood function may be maximized iteratively to produce fitted parameters and probabilities. The problem we solve here is to introduce a variant of (1) that allows both non-stationarity and the use of history and other covariates.

1.1 Overview of approach

Suppose we have spike trains from n neurons recorded simultaneously over a time interval of length T , across R trials. We consider spiking patterns at a relatively fine time resolution, denoted by δ . In Section 3, we report an analysis of simultaneous spiking data recorded from primary visual cortex where we took δ to be 5 milliseconds. The spike train data may be represented as binary arrays with dimensionality $(T/\delta) \times R \times n$. For the data analysis reported in Section 3 we used arrays of dimension $200 \times 120 \times 3$ to estimate $200 \times 120 \times 8$ probabilities $p_{abc,r}^{123}(t|H_{t,r}, u_{t,r})$. In full generality there are more parameters than data values. Furthermore, non-stationarity and history dependence could, in their most general conception, be very complicated. The approach we take here simplifies the situation greatly by assuming

1. for every neuron i , the firing probability varies smoothly across time and depends on spiking history $H_{t,r}$ only through the history $H_{t,r}^i$ for neuron i , together with the covariate vector $u_{t,r}$; and
2. excess joint pairwise spiking, above that expected under independence, does not depend on either spiking history or the covariate vector.

Importantly, assumption 2 defines excess synchrony relative to “independence,” which here means independence *conditionally* on all history and covariate effects mentioned in assumption 1. In specifying these assumptions we aim to emphasize the way synchrony is judged against a backdrop of explanatory covariates. For example, it is widely appreciated that a pair of neurons may exhibit excess pairwise spiking relative to what might be expected from their time-averaged firing rates because they respond to given stimuli with roughly similar temporal profiles—this would be synchrony due to their individual firing-rate functions, as seen through overlapping PSTHs. Various methods may be used to adjust or “normalize” pairwise spiking to account for the individual time-varying firing rate functions (e.g., Aertsen et al. (1989)). Many other possible sources of pairwise spiking may be present in particular cases, including global network activity. One of our purposes here is to introduce a general framework for quantifying the contributions of alternative sources of

pairwise spiking, while assessing statistical evidence and uncertainty. A second purpose is to examine excess multi-way spiking relative to that expected from pairwise spiking. The approach we develop melds loglinear modeling, as in (1), together with point process regression modeling (which usually comes under the rubric of generalized linear models) as in numerous articles (e.g., Kass and Ventura (2001), Kelly et al. (2010a), Okatan et al. (2005), Pillow et al. (2008), Stevenson et al. (2009), Truccolo et al. (2005), Zhao et al. (2011)). We use point process regression to model the behavior of each individual neuron; we then overlay the structure of loglinear models to account for synchronous connections.

For notational simplicity we concatenate the history and covariate vectors as a single vector

$$x_{t,r} = (H_{t,r}, u_{t,r}).$$

When we consider only the spiking history $H_{t,r}^i$ of neuron i at time t on trial r we write

$$x_{t,r}^i = (H_{t,r}^i, u_{t,r}).$$

Let $p_{1,r}^i(t|x_{t,r})$ be the probability of neuron i spiking at time t on trial r . The central object in the point process framework is the *conditional intensity function*, $\lambda^i(t|x_{t,r}^i)$, which is the firing rate function in continuous time, and the relationship between the continuous and discrete time representations is summarized by

$$p^i(t|x_{t,r}^i) \approx \lambda^i(t|x_{t,r}^i) \delta, \quad (2)$$

the approximation being justified by passage to the limit as $\delta \rightarrow 0$. The corresponding key statistical result is that the likelihood function based on the continuous-time representation, in terms of the conditional intensity function, is approximately equal to the likelihood function based on the binary data and the probability of spiking $p^i(t|x_{t,r}^i)$. We use x in the notation $x_{t,r}^i$ to connote regression-style modeling of probability in terms of explanatory covariates, including history.

Now let $p_{11,r}^{ij}(t|x_{t,r})$ be the probability that neurons i and j will both spike at time t on trial r . Under assumption 1, above, we use the spike trains from each neuron i to fit that neuron's firing probabilities $p_{1,r}^i(t|x_{t,r}) = p_{1,r}^i(t|x_{t,r}^i)$ across time and across trials. Let us write such fits as $\tilde{p}_{1,r}^i(t|x_{t,r}^i)$. Under assumptions 1 and 2 we may define

$$p_{11,r}^{ij}(t|x_{t,r}) = \tilde{p}_{1,r}^i(t|x_{t,r}^i) \tilde{p}_{1,r}^j(t|x_{t,r}^j) \zeta_x^{ij}(t) \quad (3)$$

so that $\zeta_x^{ij}(t)$ represents the excess pairwise spiking above that predicted by independence (as in Ventura et al., 2005b). We have written $\zeta_x^{ij}(t)$ with the argument t and subscript x to indicate that $\zeta_x^{ij}(t)$ is a function only of time but, because it is defined through (3), it depends indirectly on the covariates used in the individual neuron firing probabilities $\tilde{p}_{1,r}^i(t|x_{t,r}^i)$. For example, we may evaluate excess synchrony with and without a covariate that measures network activity. Kass et al. (2011) gave an example in which the data indicated that $\zeta_x^{ij}(t) = 1$ when a covariate for network activity was included in $x_{t,r}$ while $\zeta_x^{ij}(t) > 1$ when it was omitted. This suggested that excess synchrony above that expected from time-varying

firing rates, while present, was due to global network activity rather than a local circuit that affected the particular pair of neurons.

We may now summarize the steps of the strategy we have implemented for estimating multi-way spiking probabilities based on all combinations of excess pairwise spiking, thereby generalizing (1) to account for non-stationarity, history, and covariates. We consider first the case of $n = 3$ neurons labelled i, j, k :

1. For all i and all t , and for every trial r , fit $p_{1,r}^i(t|x_{t,r}^i)$ to get $\widehat{p}_{1,r}^i(t|x_{t,r}^i)$;
2. For all i, j estimate $\zeta_x^{ij}(t)$ to obtain $\widehat{\zeta}_x^{ij}(t)$;
3. Based on the $(T\delta) \times R$ sets of estimated values $\widehat{p}_{1,r}^i(t|x_{t,r}^i)$, $\widehat{p}_{1,r}^j(t|x_{t,r}^j)$, $\widehat{p}_{1,r}^k(t|x_{t,r}^k)$, $\widehat{\zeta}_x^{ij}(t)$, $\widehat{\zeta}_x^{ik}(t)$, $\widehat{\zeta}_x^{jk}(t)$ from steps 1 and 2, use an iterative algorithm to obtain the complete set of $(T\delta) \times R \times 8$ estimates $\widehat{p}_{abc,r}^{ijk}(t|x_{t,r})$ of $p_{abc,r}^{ijk}(t|x_{t,r})$.

Once we have all the fitted probabilities $\widehat{p}_{abc,r}^{ijk}(t|x_{t,r})$ we are able to use them to estimate quantities defined in terms the probabilities $p_{abc,r}^{ijk}(t|x_{t,r})$. We can also use the set of fitted probabilities $\widehat{p}_{abc,r}^{ijk}(t|x_{t,r})$ to simulate artificial spike trains that reflect all of the estimated multi-neuron dependence, and we can thereby use bootstrap simulation to obtain confidence intervals for estimated quantities; these confidence intervals incorporate the statistical uncertainty from all three steps of the fitting procedure. Similarly, we can test the null hypothesis of 2-way interaction, but *no* 3-way interaction, against the alternative that also includes 3-way interaction. The bootstrap hypothesis tests also incorporate uncertainty from all the steps of the fitting procedure.

The same steps may be followed for more than 3 neurons. For example, for 4 neurons we would again use the 2-way probabilities in (5) to fit probabilities of the form $\widehat{p}_{abcd,r}^{ijkl}(t|x_{t,r})$ and thereby test the null hypothesis of pairwise interaction against both 3-way and 4-way interaction. Thus, according to the approach described here, large numbers of spiking probabilities having potentially complicated forms are estimated by first fitting to each neuron a smooth conditional intensity function and any relevant covariates, then fitting 2-way excess synchronous spiking terms, then finding multiway probabilities by iterative optimization.

To help fix ideas, Figure 1 displays a simple situation involving time-varying firing rates from 3 artificial neurons (dotted lines in panels A–C) where 2-way synchronous spiking occurs more often than predicted by independence (dotted lines in panels D–F), and 3-way interaction occurs more often than that predicted by a 2-way model. We simulated a large number of trials from this model. When a 2-way interaction model was fitted to the simulated data, the 1-way and 2-way firing rates were estimated accurately (solid lines in panels A–F), but the amount of 3-way spiking was underestimated (solid line in panel G).

1.2 Overview of article

In Section 2 we provide methodological details. We begin in Section 2.1 by briefly summarizing the fitting methods for individual-neuron probabilities, then go on to pairwise probabilities in Section 2.2, giving procedures for bootstrap confidence intervals in Section 2.3. In Section 2.4 we present the algorithm for fitting multi-way probabilities under the assumption of pairwise interaction but not higher-order interaction, in Section 2.5 we discuss bootstrap confidence intervals for functions of these probabilities, and in Section 2.6 we describe bootstrap hypothesis tests. In Section 2.7 we review the essential motivation for

the three-step approach outlined in Section 1.1 in terms of what we call *hierarchical sparsity*. In Section 3 we provide a real-data illustration by analyzing some simultaneous spiking data recorded from primary visual cortex. We then, in Section 4, use the framework described here to address a fundamental question: how much data would be needed to distinguish 3-way interaction from the 3-way spiking that occurs by chance from pairwise interaction models? In Section 5 we add remarks about the utility of this method in practice.

2 Methodology

Here, and for the remainder of the article, we omit explicit reference to the trial r , leaving it implicit. We also omit the subscript x on $\zeta_x^{ij}(t)$.

2.1 Fitting 1-way probabilities

Individual-neuron (1-way) firing probabilities may be fitted by invoking smooth point process models, where each neuron i has firing rate governed by a conditional intensity function $\lambda^i(t|x_t^i)$ Kass and Ventura (2001) and Kass et al. (2003) discussed spline-based fitting of conditional intensity functions and Pillow et al. (2008) used an alternative set of smooth basis functions that incorporate both rate variation and history dependence. Sometimes history and covariate effects may be ignored, and then the PSTH may be smoothed by Gaussian filters (kernel smoothers), fixed-knot splines, or more sophisticated methods such as BARS (DiMatteo et al., 2001). The first step of the approach suggested here is to apply one of these individual-neuron models in order to obtain estimates $\hat{p}_1^i(t|x_t^i)$ for all t .

2.2 Fitting pairwise probabilities

We now move on to pairwise probabilities, considering neurons i and j as in (3). While we could incorporate time-varying excess spiking effects as in (3), here, instead, we further specialize by assuming the excess pairwise spiking probability effects are constant across time. We therefore take $\zeta^{ij}(t) = \zeta^{ij}$ and define

$$p_{11}^{ij}(t|x_t) = p_1^i(t|x_t^i) p_1^j(t|x_t^j) \zeta^{ij}. \quad (4)$$

This is helpful when there are not large numbers of joint spikes, and it also allows us to use a simple and intuitive estimator $\widehat{\zeta}^{ij}$ of ζ^{ij} given in equation (9). By inserting the estimates $\hat{p}_1^i(t|x_t^i)$, $\hat{p}_1^j(t|x_t^j)$, and $\widehat{\zeta}^{ij}$ into the right-hand side of (4) we obtain estimates of the 2-way probabilities $p_{ab}^{ij}(t|x_t)$ as follows:

$$\begin{aligned} \hat{p}_{11}^{ij}(t|x_t) &= \hat{p}_1^i(t|x_t^i) \hat{p}_1^j(t|x_t^j) \widehat{\zeta}^{ij} \\ \hat{p}_{10}^{ij}(t|x_t) &= \hat{p}_1^i(t|x_t^i) - \hat{p}_{11}^{ij}(t|x_t) \\ \hat{p}_{01}^{ij}(t|x_t) &= \hat{p}_1^j(t|x_t^j) - \hat{p}_{11}^{ij}(t|x_t) \\ \hat{p}_{00}^{ij}(t|x_t) &= 1 - \hat{p}_{11}^{ij}(t|x_t) - \hat{p}_{10}^{ij}(t|x_t) - \hat{p}_{01}^{ij}(t|x_t). \end{aligned} \quad (5)$$

These estimates, for $i, j = 1, 2, 3$, give us all of the 2-way probabilities needed to fit a 2-way interaction model analogous to (1) that, instead, incorporates time-varying, history-dependent, and covariate-dependent firing probabilities.

To estimate ζ^{ij} we fix the values of $p_1^i(t|x_t^i)$ to be $p_1^i(t|x_t^i) = \widehat{p}_1^i(t|x_t^i)$. Using the point process representation of joint spiking for neurons i and j derived by Kass et al. (2011) we replace ζ^{ij} with ξ^{ij} and write the synchronous-spike likelihood function as

$$\log L(\xi^{ij}) = - \int \lambda^{ij}(t|x_t) dt + \sum_{t_{ij}} \log \lambda^{ij}(t_{ij}|x_{t_{ij}}) \quad (6)$$

where t_{ij} are the times of the joint (synchronous) spikes and where

$$\lambda^{ij}(t|x_t) = \lambda^i(t|x_t^i) \lambda^j(t|x_t^j) \xi^{ij}. \quad (7)$$

Equation (7) is the conditional-intensity analogue of (4). Setting the derivative of $\log L(\xi^{ij})$ to zero and solving gives the MLE

$$\widehat{\xi}^{ij} = \frac{N^{ij}}{\int \lambda^i(t|x_t^i) \lambda^j(t|x_t^j) dt} \quad (8)$$

where N^{ij} is the total number of joint spikes for neurons i and j and the denominator is the expected number of joint spikes under independence, after taking account of the covariate x . This corresponds to equation (22) of Kass et al. (2011). The formula we use in practice replaces the integral in (8) with a sum:

$$\widehat{\zeta}^{ij} = \frac{N^{ij}}{\sum_t \widehat{p}_1^i(t|x_t^i) \widehat{p}_1^j(t|x_t^j)}. \quad (9)$$

To interpret (9) let us re-emphasize the covariate x by restoring its use as a subscript: on the left-hand side of (9) we replace $\widehat{\zeta}^{ij}$ by $\widehat{\zeta}_x^{ij}$. If $\widehat{\zeta}_x^{ij} \geq 1$, the value of $\widehat{\zeta}_x^{ij}$ represents the proportionate excess synchronous spiking beyond that explained by the covariate x , while if $\widehat{\zeta}_x^{ij} < 1$ the value of $\widehat{\zeta}_x^{ij}$ is the proportionate diminution of synchrony, below that explained by x . When $\widehat{\zeta}_x^{ij} \geq 1$ it is helpful to consider the reciprocal

$$E_x = \frac{1}{\widehat{\zeta}_x^{ij}} \quad (10)$$

which ranges from 0 to 1 and represents the proportion of synchronous spikes explained by the covariate vector x (including its time-varying trial-averaged firing rate). For example, if $\widehat{\zeta}_x^{ij} = 2$ when x contains only the time-varying trial-averaged firing rate, then $E_x = .5$ and we would say the time-varying trial-averaged firing rate explains half the synchronous spikes. This gives us a way of interpreting the relative effects of various alternative covariate vectors x , as we illustrate in Section 3. If $\widehat{\zeta}_x^{ij} < 1$ we can define the inhibitory measure $I_x = \widehat{\zeta}_x^{ij}$ which then represents proportion of observed spikes relative to the number predicted by x .

2.3 Confidence intervals for pairwise effects

Standard errors and confidence intervals associated with the pairwise estimates in (9) may be obtained by a parametric bootstrap using (5). That is, equations (5) specify a set of multinomial spiking probabilities at each time t (and for each trial, separately, due to separate spiking history or covariate effects) for every spiking pattern $(1,1), (1,0), (0,1), (0,0)$. It is straightforward to generate G complete sets of pseudo-data pairwise spike trains based on these probabilities, each set of pseudo-data replicating the layout of the original data (a

typical value of G being 500). Let us use $g = 1, \dots, G$ to index the sets of pseudo-data. If we wish to obtain a 95% bootstrap confidence interval for a parameter ϕ (such as $\phi = \zeta^{ij}$), we use the pseudo-data set g to obtain $\widehat{\phi}^{(g)}$ for $g = 1, \dots, G$; that is, for every g we obtain $\widehat{\phi}^{(g)}$ by applying (9) to the pseudo-data, where all the terms $\widehat{p}_1^i(t|x_i)$ are computed from the pseudo-data. We then order the values of $\widehat{\phi}^{(g)}$, find the .025 and .975 percentiles, and use them as endpoints of the 95% confidence interval. A standard error is similarly obtained as the standard deviation of the values of $\widehat{\phi}^{(g)}$. When summarizing results with estimates and standard errors we prefer to work with $\phi = \log \zeta^{ij}$, rather than $\phi = \zeta^{ij}$, because the log transform tends to symmetrize the distribution (and thus make it closer to normal, as it should be when one interprets standard errors). Because we include the re-fitting of $\widehat{p}_1^i(t|x_i)$ to the pseudo-data, the bootstrap confidence intervals account for the uncertainty in that first step of the fitting procedure.

2.4 Fitting higher-order probabilities

For $n = 2$ neurons, equations (5) completely determine the estimated probabilities we need for statistical analysis. For $n = 3$ neurons, pairwise interaction models do not provide analogous closed-form expressions for all the necessary probabilities $p_{abc}^{123}(t|x_i)$, and the same is true for $n > 3$. Even in the simpler stationary Poisson setting, where the data become counts aggregated across time and model (1) may be applied, it is well known (e.g. Agresti, 2002) that maximum likelihood estimation of the two-way interaction model requires iterative methods. The standard algorithm in that context is iterative proportional fitting (IPF) (Agresti, 2002; Schneidman et al., 2006).

Before reviewing IPF and the modification of it we employ here, we wish to make sure it is clear what model we are referring to as the “2-way interaction model” that we are fitting to the spike train data. It is helpful to return to the simpler model (1) and explain how our approach would apply in that setting. The complete model for 3-way probabilities p_{abc} , including 3-way interactions (in statistics it is often called “saturated”), has $8 - 1 = 7$ free parameters: the 8 values p_{abc} must satisfy the constraint $\sum p_{abc} = 1$. The 2-way interaction model (1) omits the 3-way interaction and has 6 free parameters. (In (1) there are¹ 7 parameters with non-zero multipliers $\alpha, \beta_1^1, \beta_1^2, \beta_1^3, \beta_{11}^{12}, \beta_{11}^{13}, \beta_{11}^{23}$ and, again, there is the constraint $\sum p_{abc} = 1$.) In our approach we do not need to use the parameters that appear in (1). Instead, we effectively use a different set of 6 free parameters. We define

$$\begin{aligned}
 p_1^1 &= \sum_{b,c} p_{1bc} \\
 p_1^2 &= \sum_{a,c} p_{a1c} \\
 p_1^3 &= \sum_{a,b} p_{ab1} \\
 \zeta^{12} &= \frac{\sum_{c} p_{11c}}{p_1^1 p_1^2} \\
 \zeta^{13} &= \frac{\sum_{b} p_{1b1}}{p_1^1 p_1^3} \\
 \zeta^{23} &= \frac{\sum_{a} p_{a11}}{p_1^2 p_1^3}
 \end{aligned} \tag{11}$$

and then parameterize (1) using $(p_1^1, p_1^2, p_1^3, \zeta^{12}, \zeta^{13}, \zeta^{23})$. We are not able to write a closed-form expression for p_{abc} in terms of the parameter vector $(p_1^1, p_1^2, p_1^3, \zeta^{12}, \zeta^{13}, \zeta^{23})$ for the

¹In the usual way (1) is presented in the statistics literature there are 19 parameters and 13 constraints.

same reason that we are unable to write closed-form maximum likelihood estimates for p_{abc} in (1). However, any particular value of the vector $(p_1^1, p_1^2, p_1^3, \zeta^{12}, \zeta^{13}, \zeta^{23})$ does define a particular set of values of p_{abc} according to (1) based on the non-linear equations in (11). In the general case we are concerned with here we similarly use $(p_1^1(t|x_t^1), p_1^2(t|x_t^2), p_1^3(t|x_t^3), \zeta^{12}, \zeta^{13}, \zeta^{23})$, to define, for every t , a 2-way interaction model of the form (1) that we are fitting to the data.

Now, to explain the modified IPF that achieves the desired fitting, let us begin with IPF in the standard setting involving counts, with n_{abc} being the number of time bins in which the pattern (a, b, c) occurred, and m_{abc} being the expected number according to the two-way interaction model (1) where $m_{abc} = np_{abc}$ and $n = \sum n_{abc}$ is the total number of spikes. IPF produces estimates \widehat{m}_{abc} of m_{abc} . With n_{ab+} being the number of bins in which the first two neurons have the pattern (a, b) , and n_{a+c} and n_{+bc} defined similarly, the first cycle of the IPF algorithm, as in Agresti (2002), involves three steps:

$$\widehat{m}_{abc}^{(1)} = \widehat{m}_{abc}^{(0)} \left(\frac{n_{ab+}}{\widehat{m}_{ab+}^{(0)}} \right)$$

$$\widehat{m}_{abc}^{(2)} = \widehat{m}_{abc}^{(1)} \left(\frac{n_{a+c}}{\widehat{m}_{a+c}^{(1)}} \right)$$

$$\widehat{m}_{abc}^{(3)} = \widehat{m}_{abc}^{(2)} \left(\frac{n_{+bc}}{\widehat{m}_{+bc}^{(2)}} \right)$$

Iterating these steps produces convergence to the maximum likelihood estimates under very general conditions (Haberman (1974)).

To fit our general time-varying and/or history-dependent and/or covariate-dependent 2-way interaction model we replace the counts in the standard IPF above with the probability estimates given by (5), and the expected values by the probabilities $p_{abc}^{ijk}(t|x_t)$. Here, IPF will produce a set of estimates $\widehat{p}_{abc}^{ijk}(t|x_t)$, which we will write as

$$\widehat{p}_{abc}^{ijk,2way}(t|x_t) = \widehat{p}_{abc}^{ijk}(t|x_t).$$

Let

$$\widehat{p}_{ab}^{ij}(t|x_t) = \sum_c \widehat{p}_{abc}^{ijk}(t|x_t)$$

and let $\widehat{p}_{ab}^{ik}(t|x_t)$ and $\widehat{p}_{ab}^{jk}(t|x_t)$ be defined analogously. Given values $\widehat{p}_{abc}^{ijk(g)}(t|x_t)$ at iteration g , we obtain the next 3 updates as

$$\begin{aligned}
\widehat{p}_{abc}^{ijk(g+1)}(t|x_t) &= \widehat{p}_{abc}^{ijk(g)}(t|x_t) \left(\frac{\widehat{p}_{ab}^{ij}(t|x_t)}{\widehat{p}_{ab}^{ij(g)}(t|x_t)} \right) \\
\widehat{p}_{abc}^{ijk(g+2)}(t|x_t) &= \widehat{p}_{abc}^{ijk(g+1)}(t|x_t) \left(\frac{\widehat{p}_{ac}^{jk}(t|x_t)}{\widehat{p}_{ac}^{jk(g+1)}(t|x_t)} \right) \\
\widehat{p}_{abc}^{ijk(g+3)}(t|x_t) &= \widehat{p}_{abc}^{ijk(g+2)}(t|x_t) \left(\frac{\widehat{p}_{bc}^{ik}(t|x_t)}{\widehat{p}_{bc}^{ik(g+2)}(t|x_t)} \right).
\end{aligned} \tag{12}$$

For iteration g the values of the form $\widehat{p}_{ab}^{ij(g)}(t|x_t)$ are obtained from

$$\widehat{p}_{ab}^{ij(g)}(t|x_t) = \sum_c \widehat{p}_{abc}^{ijk(g)}(t|x_t).$$

Because there are $T\delta$ time values and R trials, (12) produces $RT\delta$ sets of 3 equations, e.g., for each of the data analyses in Section 3 there were 24,000 sets of 3 equations, which were iterated to convergence. There are many ways to initialize the algorithm, including taking

$\widehat{p}_{abc}^{ijk(0)}(t|x_t) = \frac{1}{8}$. We have found, in practice, that the algorithm converges in only a few iterated sets of the three steps in (12). We summarize the fitting algorithm in the box below.

Algorithm for Fitting the Generalized 2-Way Interaction Model

1. For all i and all t , and for every trial, fit $\widehat{p}_1^i(t|x_t^i)$ to get $\widehat{p}_1^i(t|x_t^i)$;
2. For all i, j use (9) to obtain $\widehat{\zeta}^{ij}$;
3. Initialize, then iterate (12) to convergence.

The output of this algorithm is the complete set of $T\delta \times R \times 8$ estimated probabilities $\widehat{p}_{abc}^{ijk,2way}(t|x_t)$

When there are $n > 3$ neurons we may apply the same method for fitting the 2-way interaction model, thereby obtaining $T\delta \times R \times 2^n$ estimated probabilities. Clearly as n grows it quickly becomes infeasible to compute a multiple of 2^n quantities. We mention this again in Section 5.

2.5 Estimates and confidence intervals for functions of probabilities

Once we have these fitted probabilities $\widehat{p}_{abc}^{ijk,2way}(t|x_t)$ we may also estimate any quantity ϕ that may be written as a function of them

$$\phi = \phi\left(\left\{\widehat{p}_{abc}^{ijk,2way}(t|x_t), \text{ for all } t, (a, b, c)\right\}\right)$$

simply by plugging in the probability estimates:

$$\widehat{\phi} = \phi\left(\left\{\widehat{p}_{abc}^{ijk,2way}(t|x_t), \text{ for all } t, (a, b, c)\right\}\right).$$

We may then compute standard errors and confidence intervals using a parametric bootstrap as in the $n = 2$ case, outlined following equation (9), using the probabilities $\widehat{p}_{abc}^{ijk,2way}(t|x_t)$ to generate the 3-way pseudo-data spike trains.

One of the benefits of 2-way models such as (1) is that the 2-way interaction coefficients may be used to provide definitions of the specific functional connectivity between two neurons after taking account of a third neuron (e.g., Martignon et al., 2000). For this purpose let us define

$$p_{1+0}^{ijk}(t|x_t) = p_{110}^{ijk}(t|x_t) + p_{100}^{ijk}(t|x_t) \quad (13)$$

and denote other marginal probabilities by analogous notations, then consider $\phi = \zeta_{|k=\emptyset}^{ij}$ defined by

$$p_{110}^{ijk,2way}(t|x_t) = p_{1+0}^{ijk,2way}(t|x_t) p_{+10}^{ijk,2way}(t|x_t) \zeta_{|k=\emptyset}^{ij} \quad (14)$$

which we estimate with

$$\widehat{\zeta}_{|k=\emptyset}^{ij} = \frac{N_{|k=\emptyset}^{ij}}{\sum_t \widehat{p}_{1+0}^{ijk,2way}(t|x_t) \widehat{p}_{+10}^{ijk,2way}(t|x_t)} \quad (15)$$

where $N_{|k=\emptyset}^{ij}$ is the total number of (1, 1, 0) spike patterns. The estimator $\widehat{\zeta}_{|k=\emptyset}^{ij}$ gives us measure of the excess spiking activity of neurons i and j that is unrelated to the activity of neuron k . We do not mean to suggest that this particular measure has some special stature—others might be envisioned—but it is intuitive and easily computed within the framework we are describing here.

We may also define the excess 3-way spiking, and estimate it. For this purpose let $p_{111}^{ijk}(t|x_t)$ denote the 3-way spiking probability under the general 3-way model, which includes 3-way interaction. We define ζ^{ijk} by

$$p_{111}^{ijk}(t|x_t) = p_{111}^{ijk,2way}(t|x_t) \zeta^{ijk} \quad (16)$$

and estimate it with

$$\widehat{\zeta}^{ijk} = \frac{N^{ijk}}{\sum_t \widehat{p}_{111}^{ijk,2way}(t|x_t)} \quad (17)$$

where N^{ijk} is the total number of triplet joint spikes for neurons i, j, k and $\widehat{p}_{111}^{ijk,2way}(t|x_t)$ is obtained from fitting the 2-way interaction model. Formula (17) is analogous to (9) and may be derived by an analogous argument.² Once again, standard errors and confidence intervals may be obtained from the parametric bootstrap, but now the 3-way interaction term is included in the model. Specifically, the spike trains are generated from the $\widehat{p}_{abc}^{ijk}(t|x_t)$ estimates obtained from using (17):

²Formula (17) is asymptotically equivalent, as $\delta \rightarrow 0$, to equation (23) of Kass et al. (2011), but it may differ in practice and seems preferable because of the direct connection between (16) and (1).

$$\begin{aligned}
\widehat{p}_{111}^{ijk}(t|x_t) &= \widehat{p}_{111}^{ijk,2way}(t|x_t) \widehat{\zeta}^{ijk} \\
\widehat{p}_{110}^{ijk}(t|x_t) &= \widehat{p}_{11}^i(t|x_t) - \widehat{p}_{111}^{ijk}(t|x_t) \\
\widehat{p}_{101}^{ijk}(t|x_t) &= \widehat{p}_{11}^j(t|x_t) - \widehat{p}_{111}^{ijk}(t|x_t) \\
\widehat{p}_{011}^{ijk}(t|x_t) &= \widehat{p}_{11}^k(t|x_t) - \widehat{p}_{111}^{ijk}(t|x_t) \\
\widehat{p}_{100}^{ijk}(t|x_t) &= \widehat{p}_1^i(t|x_t) - \widehat{p}_{111}^{ijk}(t|x_t) - \widehat{p}_{110}^{ijk}(t|x_t) - \widehat{p}_{101}^{ijk}(t|x_t) \\
\widehat{p}_{010}^{ijk}(t|x_t) &= \widehat{p}_1^j(t|x_t) - \widehat{p}_{111}^{ijk}(t|x_t) - \widehat{p}_{110}^{ijk}(t|x_t) - \widehat{p}_{011}^{ijk}(t|x_t) \\
\widehat{p}_{001}^{ijk}(t|x_t) &= \widehat{p}_1^k(t|x_t) - \widehat{p}_{111}^{ijk}(t|x_t) - \widehat{p}_{101}^{ijk}(t|x_t) - \widehat{p}_{011}^{ijk}(t|x_t) \\
\widehat{p}_{000}^{ijk}(t|x_t) &= 1 - \sum_{(a,b,c) \neq (0,0,0)} \widehat{p}_{abc}^{ijk}(t|x_t).
\end{aligned} \tag{18}$$

We illustrate by applying (15) and (17), together with bootstrap confidence intervals, in Section 3.

2.6 Hypothesis tests

When $n = 2$ the null hypothesis concerning ζ^{ij} in (4) is that of independence, $H_0 : \zeta^{ij} = 1$. To test H_0 we may use a parametric bootstrap, with spike trains generated by the fitted independence model

$$\widehat{p}_{ab}^{ij|\zeta=1}(t|x_t) = \widehat{p}_a^i(t|x_t^i) \widehat{p}_b^j(t|x_t^j) \tag{19}$$

which simply requires that we generate the neuron i and neuron j spike trains independently using $\widehat{p}_1^i(t|x_t^i)$ and $\widehat{p}_1^j(t|x_t^j)$. There are several choices for the test statistic. Let us suppose we use N^{ij} and let N_{obs}^{ij} denote the value computed from the data. We generate bootstrap pseudo-data, with the same number of trials as the real data, based on $\widehat{p}_1^i(t|x_t^i)$ and $\widehat{p}_1^j(t|x_t^j)$, and repeat this procedure G times (e.g., we might take $G = 10,000$). We let $g = 1, \dots, G$ label the pseudo-data and for set g of pseudo-data we compute N^{ij} and label its value $N^{ij(g)}$. The bootstrap p -value is then

$$p = \frac{\text{number of values } g \text{ such that } N^{ij(g)} \geq N_{obs}^{ij}}{G}. \tag{20}$$

This p -value satisfies the usual properties of the parametric bootstrap, i.e., it furnishes approximately the correct probability of rejecting H_0 under the assumed null model. In practice, if the numerator is 0 we do not report $p = 0$ but rather say $p < 1/G$. Some authors prefer to add 1 to both the numerator and denominator of (20).

When $n = 3$, to test whether the 2-way interaction model is adequate to explain all the 3-way spikes, we take the general 2-way interaction model as H_0 . We may then use a parametric bootstrap by generating pseudo-data spike trains according to the probabilities $\widehat{p}_{abc}^{ijk,2way}(t|x_t)$ of Section 2.4, and can use the total number of 3-way spikes N^{ijk} as the test statistic. If we again label the sets of pseudo-data with $g = 1, \dots, G$ and the data-based value of N^{ijk} with N_{obs}^{ijk} we obtain the p -value by replacing the numerator in (20) with the number of values of g such that $N^{ijk(g)} \geq N_{obs}^{ijk}$.

When $n > 3$ we may apply the bootstrap test not only to the triplet spikes from all combinations of 3 neurons using N^{ijk} , but also the quadruplet spikes from all combinations of 4 neurons using N^{ijkl} , and so forth, in principle up to $N^{ijk\dots n}$.

The p -value in (20) is one-sided in the sense that only excess joint spiking is being assessed. To check for either excess joint spiking or diminished joint spiking we could instead use $\widehat{\zeta}^{ij}$ as the test statistic. Then, with $\widehat{\zeta}^{ij(g)}$ being the value computed from set g of the pseudo-data and $\widehat{\zeta}_{obs}^{ij}$ being the value computed from the real data we would define

$$p = \frac{\text{number of values } g \text{ such that } |\log \widehat{\zeta}^{ij(g)}| \geq |\log \widehat{\zeta}_{obs}^{ij}|}{G} \quad (21)$$

as a p -value for the two-sided test. When $\widehat{\zeta}_{obs}^{ij} > 1$, Equation (21) finds the proportion of sets of pseudo-data for which either $\widehat{\zeta}^{ij(g)} \geq \widehat{\zeta}_{obs}^{ij}$ or $1/\widehat{\zeta}^{ij(g)} \leq 1/\widehat{\zeta}_{obs}^{ij}$. Higher-order tests may be modified similarly.

2.7 Motivation from hierarchical sparsity

Central to the strategy in Sections 2.1–2.5 is the idea that we may perform fitting hierarchically, using the three steps articulated in Section 1.1 and realized in the algorithm of Section 2.5. The main justification for this procedure is based on the rate at which pairwise and multiway spikes occur.

For small binwidth δ (e.g., $\delta = .005$ seconds), each firing probability $p_1^i(t|x_t)$, is also small and, as $\delta \rightarrow 0$ we have

$$p_1^i(t|x_t) = O(\delta). \quad (22)$$

According to (4), we then have

$$p_{11}^{ij}(t|x_t) = O(\delta^2). \quad (23)$$

Equations (22) and (23), together with higher-order counterparts such as

$$p_{111}^{ijk}(t|x_t) = O(\delta^3)$$

provide a formal expression of *hierarchical sparsity*. Kass et al. (2011) applied hierarchical sparsity in deriving point process representations that approximate discrete-time models, analogously to Equation (2). We use the idea here to motivate our hierarchical fitting procedure. Under hierarchical sparsity there is likely to be good information (many spikes) available to estimate 1-way probabilities, but as we move up the hierarchy of interactions the information degrades: there are much fewer 2-way spikes, and then even fewer 3-way spikes, etc. Although we do not offer a more precise theoretical statement, it is apparent that, under hierarchical sparsity, the hierarchical fitting procedure should produce estimates of the multiway probabilities that capture well the available information in the data.

3 Data Analysis

We illustrate the approach by analyzing two sets of 3 neurons recorded from primary visual cortex, as described in (Kelly et al., 2010b) and (Kass et al., 2011). In each case the neural responses were recorded from an anesthetized monkey while sinusoidal gratings were moved across the visual field. The data in Figures 2 and 3 correspond to one second of the recordings, and there are 120 repeated trials. Both figures display raster plots from 3 cells, with dark circles superimposed to indicate triplet firing within 5 millisecond time bins. The

3 cells chosen for Figure 2 are different from the 3 cells chosen for Figure 3, except that cell 3 in Figure 2 is the same as cell 2 in Figure 3. In Figure 2 there are 11 such triplets in total, across the 120 trials, while in Figure 3 there are 12 triplets. To implement step 1 of the fitting procedure, in which we fit time-varying individual-neuron firing rate functions (see Section 2.1), we used a Gaussian filter with bandwidth $\sigma = 75$ milliseconds to smooth each PSTH, as shown at the bottom of panels A, B, and C in Figures 2 and 3. Results were not sensitive to choice of bandwidth. In related analyses we used spline fits, including BARS, but chose Gaussian filters here for speed of implementation which was important for the simulation study reported in Section 4. Panel D in each figure shows the bootstrap distribution of the number of triplets obtained under the null model using $\hat{p}_{abs}^{ijk,2way}(t|x_t)$, as described in Section 2.6. (We used $G = 500$ bootstrap samples.) For the neurons corresponding to Figure 2, $N_{obs}^{123} = 11$ is in the extreme tail of the bootstrap distribution and, applying (20), the test has a significant p -value with $p = .002$. On the other hand, for the neurons corresponding to Figure 3, $N_{obs}^{123} = 12$ is in middle of the bootstrap distribution and the test is not significant ($p = .52$).

For the neurons in Figure 2, applying (17) and then bootstrapping based on (18), we obtained the estimate of ζ^{123} is $\hat{\zeta}^{123} = 2.9$ with approximate 95% confidence interval (1.7, 4.3). We also computed estimates and confidence intervals for each $\zeta_{|k=\emptyset}^{ij}$ using (15) and then again bootstrapping³ based on (18). We obtained

$$\hat{\zeta}_{|3=\emptyset}^{12} = .68 (.31, 1.1)$$

$$\hat{\zeta}_{|1=\emptyset}^{23} = 1.0 (.81, 1.3)$$

$$\hat{\zeta}_{|3=\emptyset}^{12} = 1.0 (.86, 1.2).$$

Thus, interestingly, the estimates were all statistically indistinguishable from the null value of 1. This indicates that once the 3-way interactions are considered there are no longer any significant 2-way effects. We suspect the excess 3-way spiking among these neurons is due to the slow-wave activity discussed by (Kelly et al., 2010b) and (Kass et al., 2011). For the neurons in Figure 3 all the corresponding estimated effects and confidence intervals were consistent with null values of 1.

The results given above did not include individual-neuron history terms. When we included history, based on spike count in the preceding 100 milliseconds, the 3-way hypothesis test results were essentially unchanged: individual-neuron history effects apparently picked up network activity so that pairwise interactions became non-significant, but the 3-way interaction for the neurons in Figure 2 remained highly significant. (This was not sensitive to the width of history window.)

A related example of the approach presented here, focusing on 2-way interaction, was given in Kass et al. (2011). Two alternative pairs of neurons were analyzed, with and without

³We used (18) rather than the 2-way model because the test of $H_0 : \zeta^{123} = 1$ was significant.

covariates x_t^i that had two components: the spike count for neuron i in the previous 100 milliseconds and the total spike count among all neurons $k \neq i, j$ in the previous 100 milliseconds. Thus, x_t^i represents the recent network history, including that of neuron i itself but not including that of neuron j . We may summarize the results in terms of $E_x = 1/\zeta_x^{ij}$, as in (10), the proportion of synchronous spikes explained by covariates x as follows:

Proportion E_x of Synchronous Spikes Explained by x

covariate x	Pair 1	Pair 2
average firing rate	.52	.40
time-varying firing rate	.50	.42
network history	≈ 1	.48

Here we have written ≈ 1 to indicate that the bootstrap test of $H_0 : \zeta^{ij} = 1$ was not rejected for pair 1 when the network history covariates were included. The bootstrap 95% confidence intervals for ζ^{ij} were Pair 1: (.79,1.4) and Pair 2: (1.4,3.6). The main findings are, first, for both pairs there is highly significant synchronous activity beyond that due to firing rate ($p \approx .001$ for pair 1, $p < .0001$ for pair 2) with only about 50% and 40% of the synchronous spikes explained by firing rate⁴ for the two pairs; and, second, for pair 1, network activity appears to explain synchronous spiking, but for pair 2 it does not ($p = .0002$ from the bootstrap test of H_0). Indeed, from the pair 2 values of E_x above, the network activity explains only a very small additional proportion of spikes beyond the time-varying firing rate. For pair 2 there is excess synchronous spiking that is presumably associated with the stimulus.

For these two pairs of neurons we have also examined whether there is evidence of time-varying synchrony. The interesting case would be time-varying synchrony within a particular stimulus epoch. We considered the 300 milliseconds corresponding to each stimulus and decomposed these into the first 100 ms and the latter 200 ms, then used (9) within each of these time intervals, and computed confidence intervals for both the early part of the time interval and the later part of the time interval. In every case these two confidence intervals strongly overlapped, indicating no evidence in favor of time-varying synchrony. This may have been due to the relatively small number of synchronous spikes available within these 300 millisecond stimulus conditions across the 120 trials.

4 The Power of Tests for 3-way Interaction

We next investigate the amount of data needed to reliably detect 3-way interaction. We chose 4 data-generating scenarios. For the first two we used a model representing individual-neuron contributions together with 3-way interaction,

$$\log p_{abc}^{ijk}(t|x_t) = \alpha(t|x_t) + a\beta_a^1(t|x_t) + b\beta_b^2(t|x_t) + c\beta_c^3(t|x_t) + abc\beta_{abc}^{123}(t|x_t) \quad (24)$$

while for the latter two we used a model representing 2-way interaction together with 3-way interaction,

⁴The values of E_x are similar for the overall time-averaged firing rate (a constant firing rate) and for the time-varying firing rate; this depends on the “signal correlation,” i.e., the overlap of the PSTHs, with the latter typically being bigger than the former when there is substantially shared time-varying response to the stimulus and, thus, relatively large signal correlation.

$$\begin{aligned} \log p_{abc}^{ijk}(t|x_t) = & \alpha(t|x_t) + a\beta_a^1(t|x_t) b\beta_b^2(t|x_t) + c\beta_c^3(t|x_t) \\ & + ab\beta_{ab}^{12}(t|x_t) + ac\beta_{ac}^{13}(t|x_t) + bc\beta_{bc}^{23}(t|x_t) \\ & + abc\beta_{abc}^{123}(t|x_t). \end{aligned} \quad (25)$$

For each model we simulated spike trains across 1 second, using bin width $\delta = .005$ (5 milliseconds), with individual-neuron firing rates adjusted to be either 5 Hz or 10 Hz. In (25) we then adjusted the coefficients β^{ij} so that the 2-way excess firing would be $\zeta^{ij} = 2$. The combination of the two models with the two firing rates constituted our four scenarios. For each scenario we chose a grid of values of excess 3-way synchrony ζ^{123} given by (16), and a grid of values R for the number of trials. For each combination of ζ^{123} and R we computed the probability of rejecting H_0 after fixing the rejection cutoff so that the probability of rejection under H_0 (the α -level) was .05, as is customary. The results are shown in Figure 4, with black lines indicating the combinations of ζ^{123} and R that produce the customary minimally-acceptable power values of .8. The figure is based on 1000 replications of each simulation. Specifically, for each ζ^{123} and R , the figure displays smoothed proportions, out of 1000, of the replications that rejected H_0 .

The plots in Figure 4 indicate that in order to have substantial probability of detecting 3-way interaction one needs either a large value of ζ^{123} or a large number of trials R , or both. For example, when there are, on average, twice as many triplets than would occur under H_0 ($\zeta^{123} = 2$), in order to reject H_0 with power .8 in the 5 Hz scenario of panel *A* one would need more than 700 trials (of 1 second duration). For values of ζ^{123} clearly below 2, which could be realistic, it becomes extremely difficult to reliably reject H_0 . More than 1000 trials would be needed, and the extrapolation of the power curves indicates the number grows quickly as ζ^{123} decreases. When the number of three-way spikes is larger, the situation is more favorable. When the firing rates are higher for the 10 Hz scenario of panel *B*, but even there one needs about 150 trials to reliably reject H_0 when $\zeta^{123} = 2$. Panels *C* and *D* show the same two firing rates, but with injected pairwise correlation. The pairwise correlation effectively increases the number of three-way spikes and thus the power is larger when pairwise correlation is present. For the scenario in panel *C*, 200 trials are need to reliably reject H_0 , and in the most favorable situation of panel *D*, 75 trials are needed.

5 Discussion

Our main purpose has been to generalize statistical assessment of synchrony based on loglinear modeling (using maximum entropy models) so that it can accommodate time-varying firing rates and/or non-Poisson history effects and/or time-varying covariates. We have presented methodological details to supplement the theoretical treatment of Kass et al. (2011), which served to provide a point process foundation for discrete-time modeling. The hierarchical approach described here is sufficiently simple that the generalized 2-way model, together with multiway estimates and hypothesis tests, can be implemented relatively easily and effectively. We have produced Matlab code that may be accessed at <http://www.cnbc.cmu.edu/~rkelly/code/synchrony>. The code includes inputs for user-defined history and covariate effects.

The chief theoretical novelty in our approach is to recognize the strong heuristic of hierarchical sparsity, as articulated in Section 2.7: if multi-way spikes are relatively rare compared to individual-neuron spikes, then it should be useful to conceptualize excess spiking as involving multi-way gain factors, and to combine these with individual-neuron firing probabilities described by point process regression models. The strategy adopted here is different than other statistical treatments of multiple binary time series, such as in Liang and Zeger (1989), which are not aimed at point processes and do not consider the special

circumstance, and opportunity, presented by hierarchical sparsity; they would be considerably more cumbersome in this setting. On the other hand, hierarchical sparsity is not universally valid. It may happen that population-level activity produces substantial bursts of multiway spiking, as illustrated in panel C of Figure 5. The data in this figure come from the same experiment as the data used here in Section 3. Our approach assumes the individual-neuron conditional intensity functions account for this kind of shared activity. If conditional intensity functions failed to include appropriate covariates to identify population activity then multiway spiking might no longer be rare compared with individual-neuron spiking and hierarchical sparsity might no longer be applicable. Also, network activity such as that in Figure 5 would be an example source of trial-to-trial variation; in our approach all important sources of trial-to-trial variation must be included by defining suitable covariates. It would be interesting to consider incorporating into the framework described here latent variables to accommodate network bursting (Chen et al., 2009; Tokdar et al., 2010; Wu et al., 2009), but that is a topic for future research. In Kass et al. (2011) a covariate based on population spike counts was used in point process regression models to analyze synchrony, following the approach specified in greater detail here.

The usual parameters in (1) have some statistical virtues, as do the orthogonalized parameters discussed by Amari (2009). We have worked with a different parameterization, given in (11), with the ζ^{ij} chosen due to interpretability as a gain factor for increased synchronous firing rate. This also led to our suggestion of $\zeta_{|k=\emptyset}^{ij}$ in (14) as a measure of functional connectivity of neurons i and j in the presence of neuron k . One could, instead, introduce a generalized version of the usual parameters β_{11}^{ij} in (1) to account for time-varying firing rates, etc. In the context of synchrony investigation we find $\zeta_{|k=\emptyset}^{ij}$ more directly interpretable, but this is a matter of taste and convenience.

We also conducted a power study, the main conclusion of which is that large numbers of trials are likely needed in order to detect realistic multiway spiking above that determined by 2-way interactions. In the most favorable case we examined, where cells fire at 10 Hz, approximately 75 1-second trials are needed to reliably detect excess 3-way spiking that produces double the number of spikes expected under the 2-way model (and this assumed that diminished 3-way spiking is not of interest so that a one-sided test could be used; a two-sided test would require considerably more data). When there are lower firing rates, or less dramatic excess firing, the ability to reliably detect excess 3-way spiking deteriorates and very long recording sessions will be needed. With the code we have made available, additional scenarios may be investigated so that the power of new experiments to find excess 3-way synchrony could be considered carefully.

We have assumed that excess multiway spiking is constant in time. At least in the case of 2-way spiking there will be opportunities to examine stimulus-related time-locked increases in synchrony, with modest amounts of data, which could have important physiological relevance (Riehle et al., 1997). In our framework, this would require estimating $\zeta^{ij}(t)$ in (3), and an illustration of such an estimate was given by Ventura et al. (2005b) based on low-order spline fitting. An alternative would be to apply formula (9) repeatedly across distinct time intervals, and a continuous-time estimate could also be obtained by windowing or smoothing the numerator and denominator of (9), analogously to what was done in Kass et al. (2003). From our power results in Section 4, however, we would expect time-varying multi-way interactions to require either high firing rates or extended recording sessions. Indeed, when we examined data from 2 pairs of neurons in Section 3 and found no evidence in favor of non-constant $\zeta^{ij}(t)$, we observed that we had relatively few synchronous spikes to work with, and thus little statistical power to detect non-constant $\zeta^{ij}(t)$. In the context of time-varying synchrony it is also worth pointing out that, in the presence of time-varying

individual-neuron firing rates, alternative models of excess synchrony are not equivalent (Ito and Tsuji, 2000). In particular, the assumption in (1) that β_{11}^{ij} is constant in time is not equivalent to the assumption that $\zeta^{ij}(t)$ is constant in time.

We have applied the IPF procedure of Section 2.2 to triples of neurons. It could be applied to larger sets of $n > 3$ neurons, but the number of probabilities that must be fitted is 2^n so as n grows the computation will quickly become infeasible. We believe it remains possible to treat n neurons, even when n is large, but this will require additional methods and is a subject for future investigation. An important alternative to the parametric bootstrap described here is based on spike jittering, as in Harrison and Geman (2009). The jitter null hypothesis, however, is different than the bootstrap null of Section 2.3 (see Harrison et al., 2011). In future work we also plan to investigate the use of jitter in conjunction with the model-based approach described here.

Finally, no discussion of synchrony is complete without some reference to the problem of spike sorting, which undoubtedly can have an impact on synchrony detection. Considerable effort went into the characterization of spike waveforms in the data analyzed here (Kelly et al., 2007), but the problem is difficult (Harris et al., 2000; Ventura, 2009). We trust spike misclassification will be mitigated as recording technologies advance.

References

- Aertsen AMHJ, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of effective connectivity. *J. Neurophys.* 1989;900917. 61.
- Agresti, A. *Categorical Data Analysis*. 2nd ed.. Wiley; New York: 2002.
- Amari S-I. Measure of correlation orthogonal to change in firing rate. *Neural Comput.* 2009; 21:960–972. [PubMed: 19018699]
- Chen Z, Vijayan S, Barbieri R, Wilson MA, Brown EN. Discrete- and continuous-time probabilistic models and inference algorithms for neuronal decoding of up and down states. *Neural Comput.* 2009; 21:17971862.
- DiMatteo I, Genovese CR, Kass RE. Bayesian curve-fitting with free-knot splines. *Biometrika.* 2001; 88:1055–1071.
- Grün S. Data-driven significance estimation for precise spike correlation. *J. Neurophysiol.* 2009; 101:1126–1140. [PubMed: 19129298]
- Gütig R, Aertsen A. Analysis of higher-order neuronal interactions based on conditional inference. *Biol. Cybernetics.* 2003; 88:352–359.
- Haberman, SJ. *The Analysis of Frequency Data*. U. Chicago Press; 1974.
- Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsaki G. Accuracy of Tetrode Spike Separation as Determined by Simultaneous Intracellular and Extracellular Measurements. *J. Neurophysiol.* 2000; 84:401–414. [PubMed: 10899214]
- Harrison, MT.; Amarasingham, A.; Kass, RE. Statistical identification of synchronous spiking. In: Di Lorenzo, P.; Victor, J.; Taylor; Francis, editors. *Spike Timing: Mechanisms and Function*. 2011. submitted
- Harrison MT, Geman S. A rate and history-preserving algorithm for neural spike trains. *Neural Comput.* 2009; 21:1244–1258. [PubMed: 19018703]
- Ito H, Tsuji S. Model dependence in quantification of spike interdependency by joint peri-stimulus time histogram. *Neural Comput.* 2000; 12:195–217. [PubMed: 10636939]
- Kass RE, Ventura V. A spike-train probability model. *Neural Comput.* 2001; 13:1713–1720. [PubMed: 11506667]
- Kass RE, Kelly RC, Loh W-L. Assessment of synchrony in multiple neural spike trains using loglinear point process models. *Annals Appl. Statist.* 2011 In Press.
- Kass RE, Ventura V. A spike-train probability model. *Neural Comput.* 2001; 13:1713–1720. [PubMed: 11506667]

- Kass RE, Ventura V, Cai C. Statistical smoothing of neuronal data. *Network: Comput. in Neural Systems*. 2003; 14:5–15.
- Kass RE, Ventura V, Brown EN. Statistical issues in the analysis of neuronal data. *J. Neurophysiol*. 2005; 94:8–25. [PubMed: 15985692]
- Kass RE, Ventura V, Cai C. Statistical smoothing of neuronal data. *Network: Comput. Neural Systems*. 2003; 14:5–15.
- Kelly RC, Smith MA, Samonds JM, Kohn A, Bonds AB, Movshon JA, Lee TS. Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex. *J. Neurosci*. 2007; 27:261–264. [PubMed: 17215384]
- Kelly RC, Smith MA, Kass RE, Lee T-S. Accounting for network effects in neuronal responses using L1 penalized point process models. *NIPS*. 2010a; 23
- Kelly RC, Smith MA, Kass RE, Lee T-S. Local field potentials indicate network state and account for neuronal response variability. *J. Comput. Neurosci*. 2010b; 29:567–579. [PubMed: 20094906]
- Liang KY, Zeger SL. A class of logistic regression models for multivariate binary time series. *J. Amer. Statist. Assoc*. 1989; 84:447–451.
- Martignon L, Deco G, Laskey K, Diamond M, Freiwald W, Vaadia E. Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput*. 2000; 12:2621–2653. [PubMed: 11110130]
- Nakahara H, Amari S. Information geometric measure for neural spikes. *Neural Comput*. 2002; 14:2269–2316. [PubMed: 12396564]
- Okatan M, Wilson MA, Brown EN. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput*. 2005; 17:1927–61. [PubMed: 15992486]
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EJ. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008; 454:995–999. [PubMed: 18650810]
- Riehle A, Grn S, Diesmann M, Aertsen A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*. 1997; 278:19501953.
- Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*. 2006; 440:1007–1012. [PubMed: 16625187]
- Stevenson IH, Rebesco JM, Hatsopoulos NG, Haga Z, Miller LE, Kording KP. Bayesian inference of functional connectivity and network structure from spikes. *IEEE Trans. Neural Systems and Rehabilitation*. 2009; 17:203–213.
- Tokdar S, Xi P, Kelly RC, Kass RE. Detection of bursts in extracellular spike trains using hidden semi-Markov point process models. *J. Comput. Neurosci*. 2010; 29:203–212. [PubMed: 19697116]
- Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophys*. 2005; 93:1074–1089.
- Uhlhaas PJ, Pipa G, Lima B, Melloni L, Neuenschwander S, Nikoli D, Singer W. Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience*. 2009; 3
- Ventura V, Cai C, Kass RE. Trial-to-trial variability and its effect on time-varying dependence between two neurons. *Journal Neurophysiol*. 2005a; 94:2928–2939.
- Ventura V, Cai C, Kass RE. Statistical assessment of time-varying dependence between two neurons. *Journal Neurophysiol*. 2005b; 94:2940–2947.
- Ventura V. Traditional waveform based spike sorting yields biased rate code estimates. *PNAS*. 2009; 106:6921–6926. [PubMed: 19372379]
- Wu W, Kulkarni J, Hatsopoulos N, Paninski L. Neural decoding of goal-directed movements using a linear state-space model with hidden states. *IEEE Trans. Neural Syst. Rehab. Engineering*. 2009; 17:370–378.
- Zhao M, Batista AP, Cunningham JP, Chestek CA, Rivera-Alvidrez Z, Kalmar R, Ryu SI, Shenoy KV, Iyengar S. An L1-regularized logistic model for detecting short-term neuronal interactions. *J. Comput. Neurosci*. 2011

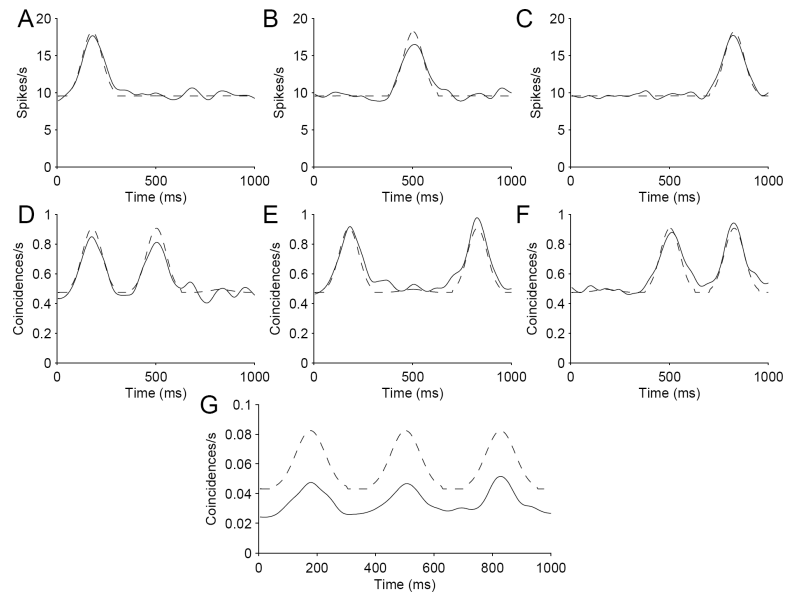


Figure 1.

Illustration of time-varying 2-way interaction model in the presence of 3-way spiking. In all panels, the dotted line is the simulated rate, and the solid line is the estimated rate based on the fitting algorithms discussed in Section 2. *A–C*: Three simulated cells, each with a different firing rate profile. Each cell has an elevation of the firing rate at a different point in the trials. *D–F*: Coincidence rates for pairwise synchronous events, which are uniformly elevated above those predicted by independence. *D* corresponds to coincidences between neurons shown in *A,B*; *E* to coincidences between *A,C*; and *F* to coincidences between *B,C*. *G*: Three-way coincidence rates. The pairwise model underestimates the number of three-way synchronous events.

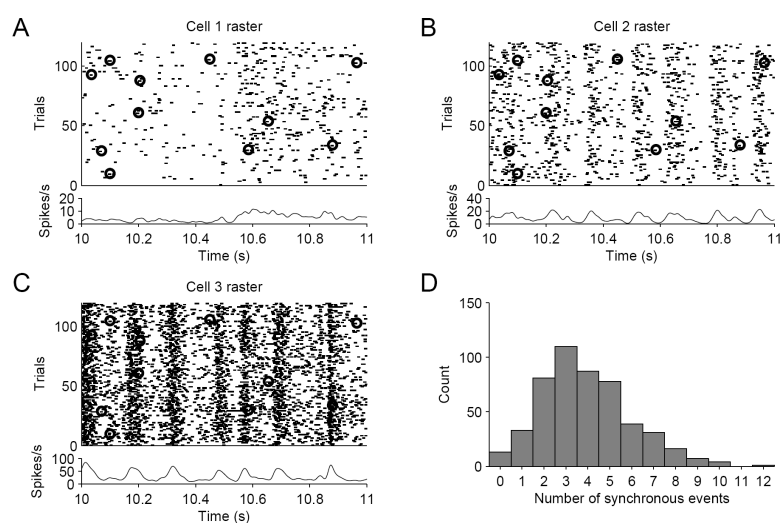


Figure 2.

A, B, C: Raster plots, with PSTHs shown below, for 3 neurons recorded simultaneously from primary visual cortex. Dark circles indicate 3-way synchronous spikes. *D:* Histogram of bootstrap values $N_{obs}^{ijk(g)}$. The value $N_{obs}^{ijk}=11$ is in the tail of this distribution.

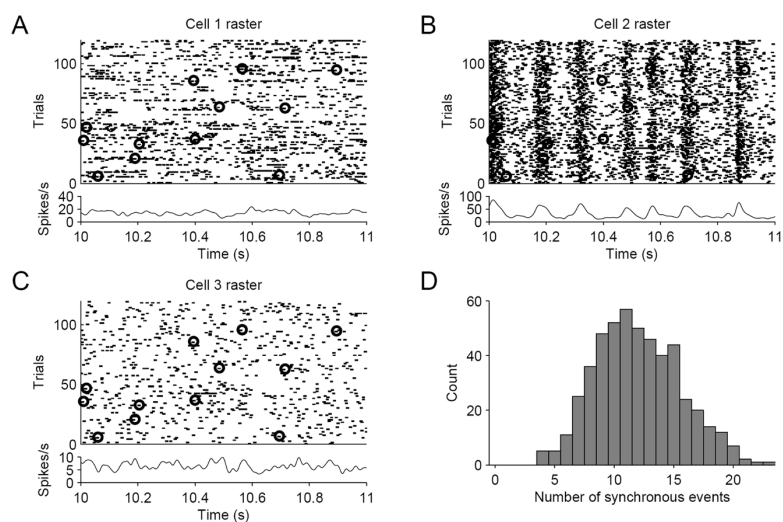


Figure 3.

A,B,C,D: Same as Figure 2, except the data are from 3 different neurons. In this case the value $N_{obs}^{ijk}=12$ is in the middle of the bootstrap histogram in *D*.




Figure 4.

Power analysis for detecting 3-way interaction. Results for four scenarios are given in the four plots. Model (24) was used for panels *A* and *B* while model (25) was used for panels *C* and *D*. For *A* and *C* the individual-neuron firing rates were set to 5 Hz and for *B* and *D* they were set to 10 Hz. The pairwise synchrony coefficient for *C* and *D* is 2. In each panel the *x*-axis is the value of the excess 3-way firing rate ζ^{123} , e.g., $\zeta^{123} = 2$ indicates twice as many triplet spikes as would be observed under the null model. The *y*-axis is the number of trials *R*. The bold lines indicate the values of ζ^{123} and *R* for which the power was .8.

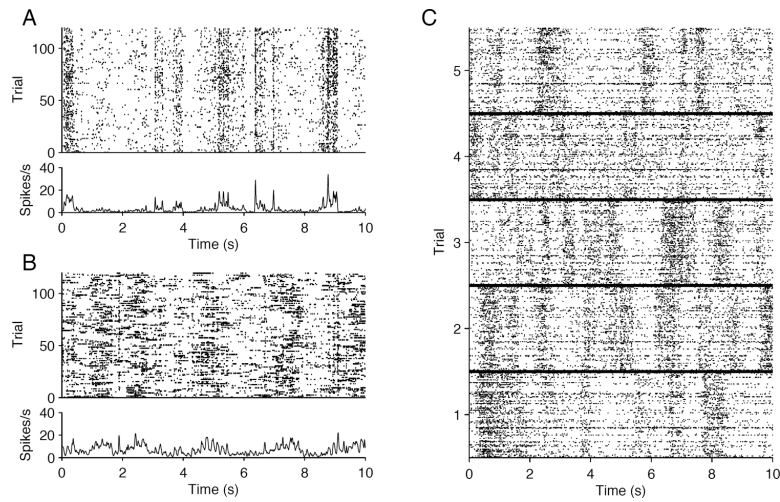


Figure 5.

Neural spike train raster plots for repeated presentations of a drifting sine wave grating stimulus (from Kass et al. (2011)). Recordings were made in V1. (A): Single cell responses to 120 repeats of multiple sinusoidal grating stimuli. At the top is a raster corresponding to the spike times, and below is a peri-stimulus time histogram (PSTH) for the same data. (B): Same as (A), for a different cell. (C): Population responses to the same stimulus, for 5 repeats. Each block, corresponding to a single trial, is the population raster for $n = 125$ units. On each trial there are several dark bands, which constitute bursts of network activity.