Discovering Functional Neuronal Connectivity from Serial Patterns in Spike Train Data

Casey Diekman
diekman@njit.edu
Department of Mathematical Sciences, New Jersey Institute of Technology,
Newark, NJ 07102, U.S.A.

Kohinoor Dasgupta
kohinoor@umich.edu
Vijay Nair
vnn@umich.edu
Department of Statistics, University of Michigan, Ann Arbor, MI 48109, U.S.A.

K. P. Unnikrishnan
kunnikrishnan@northshore.org
Center for Biomedical Research Informatics, NorthShore University HealthSystem,
Evanston, IL 60201, U.S.A.

Repeating patterns of precisely timed activity across a group of neurons (called frequent episodes) are indicative of networks in the underlying neural tissue. This letter develops statistical methods to determine functional connectivity among neurons based on nonoverlapping occurrences of episodes. We study the distribution of episode counts and develop a two-phase strategy for identifying functional connections. For the first phase, we develop statistical procedures that are used to screen all two-node episodes and identify possible functional connections (edges). For the second phase, we develop additional statistical procedures to prune the two-node episodes and remove false edges that can be attributed to chains or fan-out structures. The restriction to nonoverlapping occurrences makes the counting of all two-node episodes in phase 1 computationally efficient. The second (pruning) phase is critical since phase 1 can yield a large number of false connections. The scalability of the two-phase approach is examined through simulation. The method is then used to reconstruct the graph structure of observed neuronal networks, first from simulated data and then from recordings of cultured cortical neurons.

1 Introduction

Multielectrode array recordings have become a standard tool in neuroscience, leading to large collections of neuronal spike train data. However,

there has not been an analogous development in statistical methods to extract the spatial and temporal information in these large data sets quickly and efficiently (Brown, Kass, & Mitra, 2004). The goal of this letter is to develop statistical methods for analyzing patterns in multineuronal spike train data and use them to reconstruct the graph structure of the underlying neuronal network.

The primary contribution of the letter is the development of a two-phase strategy and associated statistical procedures for identifying functional connections based on frequent episodes in spike train data. In phase 1, the methodology counts the occurrences of all two-node episodes and screens them to identify possible functional connections (edges). In phase 2, the set of identified two-node episodes is pruned to remove the potentially large number of false edges that can be attributed to chains or fan-out structures. We focus on episode counts based on the concept of nonoverlapping occurrences. The restriction to nonoverlapping occurrences is important for computational reasons. Computationally efficient algorithms have been developed in the literature (Laxman, Sastry, & Unnikrishnan, 2005, 2007; Patnaik, Sastry, & Unnikrishnan, 2008). However, the statistical theory associated with nonoverlapping counts has not been fully developed (Diekman, Sastry, & Unnikrishnan, 2009; Sastry & Unnikrishnan, 2010). We develop this theory and use it to derive associated statistical procedures for phases 1 and 2. The two-phase approach is then applied to spike train data in order to reconstruct the graph structure of both simulated and experimentally observed networks of neurons.

Counts of serial episodes are statistical summaries that characterize the directional relationships among neurons. They are formally defined in section 2, where we also provide additional background for the data types considered in the letter. Section 3 develops the distribution theory for serial episodes under a stochastic framework. The results are then used in section 4 to develop statistical thresholds for determining the important episodes (those that correspond to strong functional connectivity among the neurons) and estimate the strengths of the connectivity. In section 5, we show how the results can be used to recover the structure of simulated networks and examine the scalability of the method. Section 6 illustrates application of the methods to multi-electrode array data from cortical cultures. The letter concludes with section 7 discussing how the results can be extended.

2 Background

Spike train data are time-ordered sequences of events, as indicated by the illustrative data set involving five neurons, *A* to *E*:

[(A, 1), (D, 3), (C, 4), (A, 5), (B, 7), (C, 8), (B, 11),

(C, 12), (E, 12), (A, 13), (B, 15), (C, 16)]

In this case, neuron A fired at time unit 1, neuron D fired at time unit 3, and so on. When a neuron influences another, the influence is mediated by processes (e.g., synaptic communication) that have characteristic time delays on the order of milliseconds. Thus, the difference between the times of the two spikes should differ by some time unit k. The subset $\langle (A, 1), (C, 4) \rangle$ indicates C occurring $k = 3$ milliseconds after A. Temporal patterns in spike train data (such as the firing of one neuron followed by the firing of another after a fixed time) that occur frequently are indicative of neuronal microcircuits. Automatically detecting such patterns, and using statistical methods to separate patterns due to underlying connections from those that occur randomly, is very useful for understanding the functional connectivity in the neural tissue.

Serial episodes are temporal patterns with a specific ordering of the events, while parallel episodes have no constraints on the relative order of events in the pattern (Mannila, Toivonen, & Verkamo, 1997). The analysis in this letter focuses on serial episodes with interevent time constraints. For an episode involving the two neurons, A and B, we represent the time information through the notation $A[k]B$, where k is the interevent delay. For example, in the above data set, there are three distinct occurrences of the episode $A[3]C$: $\langle (A, 1), (C, 4) \rangle$; $\langle (A, 5), (C, 8) \rangle$, and $\langle (A, 13), (C, 16) \rangle$. A serial episode that occurs a large number of times in the data is called a frequent episode. Clearly the term *large* needs further quantification, and this is in fact one of the goals of the letter. Further, we will restrict attention to two-node interactions. We interpret frequent multinode episodes of the form $A[k_1]B[k_2]C$ as indicative of either chain structure (consecutive occurrences of two-node episodes such as $A[k_1]B$ followed by $B[k_2]C$), fan-out structure (co-occurrence of the two-node episodes $A[k_1]B$ and $A[k_1 + k_2]C$), or both (feedforward loop structure) (Milo et al., 2002; Marbach et al., 2010) (see Figure 1). Chains and fan-out structures are known to lead to systematic prediction errors (in particular, false positives) in algorithms for inferring gene regulatory networks (Marbach et al., 2010). The procedures developed in this letter are designed to infer these structures accurately. There is some discussion in the neuroscience literature about "true" third- and higher-order interactions (Martignon et al., 2000; Amari, Nakahara, Wu, & Sakai, 2003; Gerhard, Pipa, Lima, Neuenschwander, & Gerstner, 2011). However, Schneidman, Berry, Segev, and Bialek (2006) suggests that a significant amount of the connectivity in neuronal networks can be explained by just two-node interactions.

The time delay for a connection between two neurons in general will not be known, so we will count the occurrences of episodes for a range of possible values of k. This in turn implies that we need computationally efficient algorithms for counting. Frequent episode discovery, a general method in temporal data mining (Mannila et al., 1997), has been proposed as a method for efficiently counting the numbers of serial episodes with interevent time constraints (Patnaik et al., 2008). It is also possible that the
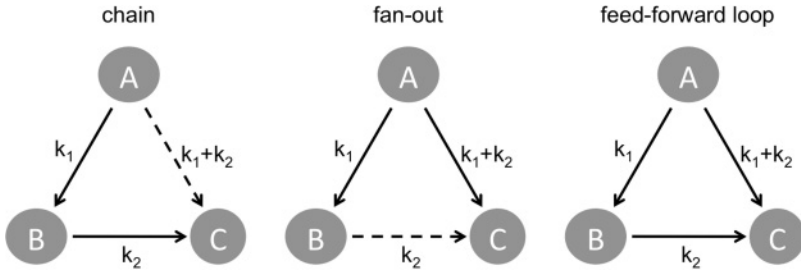
Figure 1: The three network structures that are likely to lead to frequent occurrences of the two-node episodes $A[k_1]B$, $B[k_2]C$, and $A[k_1 + k_2]C$. For chain and fan-out structures, one of these frequent episodes represents a false edge (dashed lines).

time delay is random. While the methods in this letter can be adapted to handle this situation, we will not pursue it here.

The analysis in this letter focuses on episode counts based on nonoverlapping occurrences. To explain this concept, consider the sequence for two neurons A and B:

$$[(A, 1), (B, 2), (A, 3), (A, 5), (B, 6), (B, 8), (A, 9),$$
$$(B, 10), (A, 12), (B, 14)]$$

There are several occurrences of the episode $A[5]B$ in this sequence: $\langle (A, 1), (B, 6) \rangle$, $\langle (A, 3), (B, 8) \rangle$, and $\langle (A, 5), (B, 10) \rangle$. The last two occurrences overlap with the first in the sense that events (A, 3) and (A, 5) occur before (B, 6). In this case, we cannot determine whether the firing of (B, 6) was influenced by (A, 1), (A, 3), or (A, 5). If we are to interpret a serial episode as part of a microcircuit in the underlying neural tissue, then it should represent a series of causative or triggering events. A frequent episode based on counts of nonoverlapping occurrences would be such a series. To allow for this, some researchers (Achar, Laxman, & Sastry, 2011) restrict attention to nonoverlapping occurrences, namely, occurrences with nonoverlapping time periods. Figure 2 summarizes the number of overlapping and nonoverlapping occurrences for the episode $A[5]B$ for a longer data sequence. There are $N = 6$ counts of $A[5]B$ and only $M = 3$ of these are nonoverlapping occurrences.

An additional advantage of restricting attention to nonoverlapping occurrences is that it is more efficient to count than total occurrences, an important concern when dealing with large data sets and searching for episodes at all possible values of k, the time delays. Efficient algorithms have been developed for such situations by Laxman et al. (2005, 2007) and
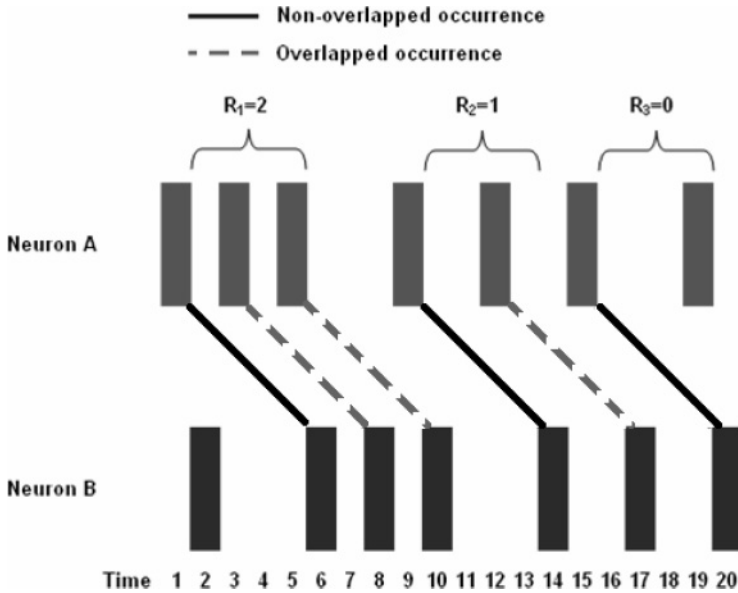
Figure 2: Nonoverlapping occurrences of episode A[5]B. The total number of occurrences in $L = 20$ bins is $N = 6$, consisting of $M = 3$ nonoverlapping occurrences and $R_1 + R_2 + R_3 = 3$ overlapping occurrences.

Patnaik et al. (2008). These algorithms use finite-state automata, and part of their efficiency comes from the fact that counting nonoverlapping occurrences requires only one automaton per episode.

The rest of the letter is organized as follows. Section 3.1 describes the basic model and develops the distributions of the number of occurrences of episodes. The main contribution here is the development of the distribution based on nonoverlapping occurrences. Sastry and Unnikrishnan (2010) use recurrence relations to compute just the mean and variance of this random variable. Our results are more general, and we provide a complete characterization that allows us to compute the moment-generating function and, hence, all moments. The results also allow us to compute the covariances of the random variables for different types of episodes, which are needed for testing various hypotheses of interest. We use the distributional results to develop statistical thresholds for separating the observed episodes into two groups: (1) random (arising from pairs of nodes that are independent) and (2) active (pairs of nodes that are connected). A simple approach is to test for independence of the pairs of nodes. However, in practice, researchers are interested in detecting only sufficiently strong connections (as indicated by some measure of strength). Section 4 considers an appropriate notion of strength and develops methods for determining connectivity above (or

below) the specified measure of strength. Two issues arise in implementing the methods: screening out false positives that occur due to chains and fan-outs and controlling the overall error rates due to multiple testing. Both topics are discussed in section 4. The results are then applied to simulated networks in section 5 and data from cortical cultures in section 6. Section 7 briefly discusses extensions to nonstationary models and other practical issues.

3 Distributions of Episode Counts

3.1 The Model. Multivariate point processes (and their discrete versions) are often used to model multineuronal spike trains (see Brillinger, 1988; Chornoboy, Schramm, & Karr, 1988; Truccolo, Eden, Fellows, Donoghue, & Brown, 2005; Rigat, Gunst, & Pelt, 2006). We consider a class of dependent Bernoulli processes to characterize the behavior of a set of neurons. These can be viewed as discretized versions of the Poisson processes that have been extensively discussed in the literature.

Let J be the number of neurons being observed. Let Δ be a small enough interval so that there is at most one spike per neuron in the interval. We observe the process over a period of L consecutive intervals. Let $X_j(t) = 1$ if neuron j fires at time t and zero otherwise, for $t = 1, \ldots, L$ and $j = 1, \ldots, J$. Let P_j be the probability of neuron j firing in an interval, and denote the number of firings during L intervals as N_j for $j = 1, \ldots, J$. In the artificial case where the neurons fire independently, the $X_j(t)$s follow independent Bernoulli processes, and N_js will be independent binomial(L, P_j) random variables.

We are interested in the case where there is dependence among the neurons, and we formulate it as follows. Suppose we have two neurons, A and B, with unidirectional dependence: A influences B, but B does not influence A, and there is a delay of k time units between the time A fires and when it affects the probability of B firing (written as $A \xrightarrow{k} B$). Since the firing of A is unaffected by B, $X_A(t)$ is still a Bernoulli process with success probability P_A. The firing of neuron B is due to a mixture of two processes: it can fire on its own or due to excitation by A. Since each of these processes is independent over time, the firing of B across different intervals remains independent. The mixture probability of B firing in any interval t is

$$P(X_B(t) = 1) = P(X_B(t) = 1|X_A(t-k) = 1)P(X_A(t-k) = 1)$$
$$+ P(X_B(t) = 1|X_A(t-k) = 0)P(X_A(t-k) = 0)$$
$$= P_A P_{B|A} + (1 - P_A)P_{B|\bar{A}},$$
$$P(X_B(t) = 0) = 1 - P(X_B(t) = 1).$$

Further, $X_B(t)$ is marginally a Bernoulli process with $P_B^c = P_A P_{B|A} + (1 - P_A) P_{B|\bar{A}}$, where the superscript c stands for combined. However, $(X_A(t), X_B(t))$ are now dependent.

While the above formulation is relatively simple, it has been found to be reasonable for modeling the firing behavior of cortical neurons (Shadlen & Newsome, 1998). Nevertheless, neurons do exhibit time-varying behavior, so the reasonableness of the Bernoulli model for a particular data set of interest should be tested empirically (Kass, Ventura, & Brown, 2005). If the temporal behavior is piece-wise stationary or slowly varying, the results can be applied to segments of time intervals within which the behavior is approximately stationary. Another issue is that immediately following a spike, a neuron may exhibit a short (typically less than 1 ms) absolute refractory period during which it cannot spike again (Kass & Ventura, 2001; Olson, Gettner, Ventura, Carta, & Kass, 2000). Such refractory effects become detectable when neurons have high enough firing rates.

3.2 Distribution in the Unrestricted Case. Consider first the case in which all occurrences of an episode are counted; attention is not restricted to nonoverlapping occurrences. Suppose the episode of interest is $E = A[k]B$: a firing of neuron A followed k seconds later by a firing of neuron B. For $t = 1, \ldots, L - k$, consider the binary process $I_E(t)$ that equals 1 if $X_A(t) = 1$ AND $X_B(t + k) = 1$. The $I_E(t)$'s are independent and identically distributed (i.i.d.) Bernoulli with success probability $P_{A[k]B} = P_A P_{B|A}$ for $t = 1, \ldots, L - k$. Let N denote the number of times that $I_E(t)$ equals 1, that is, the number of occurrences of the serial episode E. The possible values of N range from 0 to $(L - k)$; if A fires after time $t = L - k$, an occurrence of the episode would not be complete by time $t = L$. It is clear that N has a binomial distribution with parameters $(L - k)$ and $P_{A[k]B}$. Therefore, one can estimate $P_{A[k]B}$ using the simple estimator $\hat{P}_{A[k]B} = N/(L - k)$. $\text{Var}(\hat{P}_{A[k]B}) = P_{A[k]B}(1 - P_{A[k]B})/(L - k)$. This result holds provided there are no cycles among the neurons; the network structure connecting the neurons is a directed acyclic graph (DAG). If neuron A influences B and B in turn influences A, the model becomes much more complicated. We do not consider such situations here.

3.3 Distribution in the Nonoverlapping Case. The technically challenging case is the distribution of serial episode counts based on nonoverlapping occurrences. Consider again two neurons A and B, and let M be the number of nonoverlapping occurrences of the episode $E = A[k]B$. Iterative methods for calculating the mean and variance of M were developed in Sastry and Unnikrishnan (2010). We provide a complete characterization of the distribution and also obtain an expression for its moment-generating function (and hence moments of all orders).

We can partition N, the total number of occurrences, as follows. Consider the jth nonoverlapping occurrence of E, and let R_j be the number of

(unrestricted) occurrences that occur between the $(j-1)$th and jth nonoverlapping occurrences (see Figure 2 for an illustration). Then it is easily seen that $N = M + \sum_{j=1}^{M} R_j$. Further, the R_js and M are independent of each other. So we have the following result:

Proposition. *N has the same distribution as*

$$N = M + \sum_{j=1}^{M} R_j. \tag{3.1}$$

Further, the R_j's are i.i.d. binomial(k, $P_{A[k]B}$) random variables.

This provides a complete characterization of the distribution of M. Although we cannot use the result to compute the distribution of M analytically, the moment-generating function (and hence all the moments of M) can be computed easily. The mean and variance can be obtained directly. Taking expectations of both sides of equation 3.1 and using the independence of the R_js and the distribution of N, we have

$$E[M] = \frac{E[N]}{1 + E[R_1]} = \frac{L-k}{1/P_{A[k]B} + k}. \tag{3.2}$$

Similarly, the variance of M is obtained as

$$\text{Var}[M] = \frac{(L-k)P_{A[k]B}(1 - P_{A[k]B})}{(1 + kP_{A[k]B})^3} \tag{3.3}$$

$$= \frac{\text{Var}[N]}{(1 + kP_{A[k]B})^3}. \tag{3.4}$$

To get the moment-generating function (MGF) of M, we write

$$E[\exp(tN)] = E\left[\exp\left(t\left(M + \sum_{j=1}^{M} R_j \right) \right) \right]$$

$$= E_M\left[E_X\left[\exp\left(t\left(M + \sum_{j=1}^{M} R_j \right) \right) | M \right] \right]$$

$$= E[\exp(tM)(Q_{A[k]B} + P_{A[k]B}\exp(t))^{Mk}]$$

$$= E[\exp(M(t + k\log(Q_{A[k]B} + P_{A[k]B}\exp(t))))], \tag{3.5}$$

where $Q_{A[k]B} = 1 - P_{A[k]B}$. Let $s(t) = t + k \log(Q_{A[k]B} + P_{A[k]B} \exp(t))$. This is a one-to-one and strictly increasing function of t, so a unique inverse exists. Denote this inverse function by $g(s)$. (Note that $g(s)$ has to be obtained numerically.) Then the MGF of M can be expressed as

$$\mathrm{E}\,[\exp(sM)] = \mathrm{E}\,[\exp(g(s)N)] = (Q_{A[k]B} + P_{A[k]B} \exp(g(s)))^{L-k}.$$

We can get the nth moment of the distribution of M by differentiating the MGF n times with regard to s and evaluating it at $s = 0$.

The distribution of M (CDF or pmf) can be approximated up to any degree of precision by Monte Carlo sampling and the characterization in equation 3.1. We also examined more direct approximations and found a fourth-order polynomial function of a normal random variable to be reasonable. Specifically, one can approximate the distribution of M by $\sum_j a_j Z^j$, where Z is a standard normal random variable and the a_js are chosen to match the first four moments of M (Fleishman, 1978; Zhao, Ono, & Ishii, 2002).

Equation 3.2 yields an estimator of $P_{A[k]B}$ based on M as

$$\hat{P}^M_{A[k]B} = \left(\frac{L-k}{M} - k \right)^{-1}, \qquad (3.6)$$

a nonlinear function of M. Its variance can be approximated using Taylor series as

$$\mathrm{Var}[\hat{P}^M_{A[k]B}] \approx \frac{(1 + kP_{A[k]B})P_{A[k]B}(1 - P_{A[k]B})}{(L-k)}. \qquad (3.7)$$

As a side issue, one might wonder about the loss in efficiency in basing the analysis on nonoverlapping occurrences M rather than N. A comparison of the (asymptotic) variances of $\hat{P}^N_{A[k]B}$ (estimator based on N) and $\hat{P}^M_{A[k]B}$ (based on M) provides insight into this question. The ratio of the asymptotic variances of $\hat{P}^N_{A[k]B}$ and $\hat{P}^M_{A[k]B}$ is

$$\text{Relative efficiency} = \frac{1}{1 + kP_{A[k]B}}. \qquad (3.8)$$

It is independent of L and decreases as k or $P_{A[k]B}$ get large, suggesting a greater loss in efficiency in using M in such situations. Of course, these are also precisely the cases where the number of occurrences will be large, and therefore, are the situations where the gains in computational efficiency due to using M over N will be greatest.

4 Detecting Active Connections

Consider again two neurons, A and B. If an excitatory connection is present, $P_{A[k]B} > P_A \times P_B$; for an inhibitory connection, $P_{A[k]B} < P_A \times P_B$. We will restrict attention to the excitatory case as the arguments for the inhibitory case are analogous. The simplest question to ask is if the number of observed occurrences of an episode (based on N or M) is larger than what one would expect under randomness. This can be formulated as testing the null hypothesis $P_{A[k]B} = P_A \times P_B$ against the alternative $P_{A[k]B} > P_A \times P_B$. If we let $S = P_{A[k]B}/P_A P_B$, the question becomes testing $H_0 : S = 1$ against $H_1 : S > 1$. The value of S can be viewed as a measure of the connection strength between A and B.

Researchers are typically interested in detecting only connections that are sufficiently strong or active; weak connections are often not of interest. Let S_0 be a user-specified threshold for strength. Let $\tau(S) = P_{A[k]B} - S P_A P_B$. Then detecting $S > S_0$ is equivalent to testing the hypothesis:

$$H_0 : \tau(S_0) \leq 0 \quad \text{against} \quad H_A : \tau(S_0) > 0. \tag{4.1}$$

In other words, we decide in favor of $S > S_0$ only if the null hypothesis is rejected, a decision that requires sufficiently strong evidence.

4.1 Test Procedures. Consider first the case where one observes N, the episode count based on all occurrences of $A[k]B$. Then we can estimate the parameters in $\tau(S)$ based on the underlying binomial data. Denote the corresponding estimator as $\hat{\tau}_N(S)$. The variance can be expressed as

$$\text{Var}\left(\hat{\tau}_N(S)\right) = \text{Var}\left(\hat{P}_{A[k]B}\right) + S^2 \text{Var}\left(\hat{P}_A \hat{P}_B\right) - 2S \text{Cov}\left(\hat{P}_{A[k]B}, \hat{P}_A \hat{P}_B\right). \tag{4.2}$$

Expressions for these components are obtained in the appendix, sections A.1 and A.2. The problem is nontrivial since $\hat{P}_A$ and $\hat{P}_B$ are not independent when $S > 1$. The estimated variance of $\hat{\tau}_N(S)$ can be obtained by substituting the estimates for the unknown parameters in the variance expressions in section A.1. One can then test the hypothesis of an active connection given in equation 4.1 using a normal approximation for the standardized statistic $Z_\tau = \hat{\tau}_N(S_0)/\sqrt{\widehat{\text{Var}}(\hat{\tau}_N(S_0))}$. In practice, one can implement the test procedure by computing the approximate lower confidence bound corresponding to the desired α-level for $\tau(S_0)$ and checking if the confidence bound does or does not include the value zero.

The more technically challenging problem is to develop analogous procedures based on M or, equivalently, $\hat{\tau}_M(S_0)$. While formulation of the problem and the structure of the estimator $\hat{\tau}_M(S_0)$ are the same as before, the

computation of $\text{Var}(\hat{\tau}_M(S_0))$ is now much more involved. Knowledge of the third- and fourth-order moments of M, discussed in section 3.3, is needed for this purpose. The variance expressions are developed in sections A.3 and A.4. An illustration of these procedures is discussed in section 5. We have examined the adequacy of the normal approximations to the test statistics through simulation studies and found them to be quite adequate.

4.2 Multiple Testing. The proposed procedure involves testing all pairs of nodes at all possible values of the delay k to identify the active two-node episodes. When either the number of neurons or the number of possible values of k is large, one would have to be concerned about the problem of multiple testing. Performing hundreds or thousands of tests can lead to the spurious detection of false positives. There are several methods in the literature for handling multiple comparisons, including Bonferroni correction to control the overall error and the more recent technique of false discovery rates (FDRs; see, e.g., Shaffer, 1995; Benjamini & Hochberg, 1995). We used simulation studies to determine the extent of false positives from multiple testing and which methods for multiple comparisons perform well in our context. We found the Bonferroni correction to be too conservative, as is well known to be the case in the literature. The FDR method is not ideally suited for our situation and did not perform particularly well either. The approach we found to be most effective in our simulation studies was to use S_0, the user-specified threshold for strength, as a tuning parameter. As we increased the value of S_0 from 1 (null hypothesis of independence) to about 3 or 4, most of the false positives due to multiple testing disappeared (see Table 1). The statistical methods, as expected, perform well in detecting connections with relatively strong strengths. Therefore, we propose the use of the strength threshold as a tuning parameter: try different values of S_0 and select the value where the conclusions (number of edges detected) stabilize.

4.3 Eliminating False Edges. In addition to false positives due to multiple testing or type 1 error, certain types of network structure can lead to systematic false positives (Marbach et al., 2010). We refer to false positives due to network structure as false edges. Figure 1 shows the two kinds of false edges that are likely to occur (the dashed lines). The left panel shows the situation where we have a pair of active connections $A[k_1]B$ and $B[k_2]C$. This may lead to the detection of $A[k_1 + k_2]C$ as statistically significant even if it is not an active connection. The center panel shows a different scenario where false edges can occur. When there are active connections $A[k_1]B$ and $A[k_1 + k_2]C$, then $B[k_2]C$ may be detected as active even if it is not. So we need to do a second pass through the connections that are detected as active to determine if there are false edges due to this phenomenon.

We propose procedures for testing and screening out false edges in the different scenarios described above. Suppose we have detected all three of

Table 1: Average Number of Significant Two-Node Connections for Different Connection Strengths S and Strength Thresholds S_0.

Connection Strength	Strength Threshold								
	$S_0 = 1$	2	3	4	5	10	15	20	25
$S = 1$	198.51	0.97	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	208.16	1.60	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3	209.75	3.27	0.25	0.00	0.00	0.00	0.00	0.00	0.00
4	217.25	5.07	1.42	0.33	0.00	0.00	0.00	0.00	0.00
5	230.42	7.34	3.44	1.11	0.44	0.00	0.00	0.00	0.00
10	239.33	8.25	7.00	6.88	6.63	0.38	0.00	0.00	0.00
15	241.39	7.67	7.00	7.00	7.00	6.25	0.28	0.00	0.00
20	245.27	7.61	7.00	7.00	7.00	7.00	5.38	0.35	0.00
30	249.71	9.76	8.86	7.14	7.14	7.11	7.02	7.00	4.91
40	265.47	9.14	9.00	8.57	8.57	7.45	7.13	7.00	7.00

Note: The network is as shown in Figure 3, 100 replicates, $\alpha = 0.05$.

the episodes $A[k_1]B$, $A[k_1 + k_2]C$, and $B[k_2]C$ as significant. Consider first the case in the left panel of Figure 1 and the events $AC = A[k_1 + k_2]C = \{X_A(t) = 1, X_C(t + k_1 + k_2) = 1\}$, $ABC = \{X_A(t) = 1, X_B(t + k_1) = 1, X_C(t + k_1 + k_2) = 1\}$, and $A\bar{B}C = \{X_A(t) = 1, X_B(t + k_1) = 0, X_C(t + k_1 + k_2) = 1\}$ for $t \in [0, L - k_1 - k_2]$. Denote the number of nonoverlapping occurrences of the respective events as M_{AC}, M_{ABC}, and $M_{A\bar{B}C}$, respectively. (The arguments for overlapping occurrences are similar.) If the connection AC is active, then $M_{A\bar{B}C}$ will be much more than what one would expect under independence. So we can test the hypothesis

$$H_0: P_{A\bar{B}C} \leq P_A P_{\bar{B}} P_C,$$
$$H_1: P_{A\bar{B}C} > P_A P_{\bar{B}} P_C, \tag{4.3}$$

to determine if the AC connection is active. Let $\xi = P_{A\bar{B}C} - P_A P_{\bar{B}} P_C$; then the above test is equivalent to testing whether $\xi \leq 0$ or $\xi > 0$.

Consider now the second type of false edge shown in the center panel of Figure 1. Again suppose that we have detected all three of the episodes $A[k_1]B$, $A[k_1 + k_2]C$, and $B[k_2]C$ as significant. Our goal is to test if the third episode is a false edge due to co-occurrences of the first two episodes (i.e., a single firing of A followed by both a firing of B and a firing of C with the specified delays).[1] We now use the difference between $P_{\bar{A}BC}$, the probability of the episode $B[k_2]C$ in the absence of A, and the probability of the same episode under independence, namely, $P_{\bar{A}} P_B P_C$. The hypothesis

---

[1] The same arguments apply to detecting the first episode as a false edge; however in such a scenario, there is no reason to expect the last two episodes to co-occur frequently.
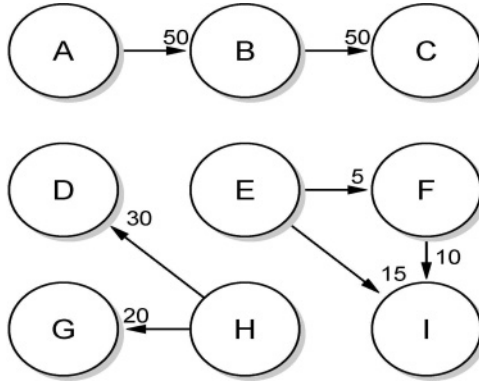
Figure 3: A network with nine neurons and embedded connections. Each neuron had an intrinsic firing rate of 5 Hz. The connections had the time delays indicated next to each arrow (in ms) and strengths ranging from $S = 1$ to 40.

test to eliminate this type of false edge is as follows:

$$H_0: \ \eta = P_{\bar{A}BC} - P_{\bar{A}}P_B P_C \leq 0,$$
$$H_1: \ \eta > 0. \tag{4.4}$$

Denote the test statistics for these procedures as $Z_\xi = \hat{\xi}/\sqrt{\widehat{\mathrm{Var}}(\hat{\xi})}$ and $Z_\eta = \hat{\eta}/\sqrt{\widehat{\mathrm{Var}}(\hat{\eta})}$. The variances of these estimators are developed in sections A.2 and A.4. As before, we use a normal approximation of the standardized statistic to implement the procedure. One can develop other procedures for screening out false edges, but these tests are simple to implement when dealing with statistics based on nonoverlapping occurrences, and we found them to perform well in simulation studies.

5 Simulation Studies

To illustrate the usefulness of our results, we simulated data from a simple network of nine neurons with a baseline firing rate of 5 Hz and the seven connections shown in Figure 3. We varied the strength of the connections by setting S at values ranging from 1 to 40. Note that $S = 1$ corresponds to no connectivity as it implies that the neurons always spike according to the baseline firing rate.

For each value of S, we simulated 100 data sets from the network. To analyze the data, we first count the total number of spikes from the nine neurons; these can be viewed as one-node episodes. We then obtain the count of all possible two-node serial episodes with delays ranging from

1 to 200 ms. There are 81 pairwise episodes, so this is a total of $81 \times 200 = 16,200$ cases. We can then test for active two-node connections. Since we are performing $n = 81 \times 200 = 16,200$ tests, we must consider the issue of multiple testing as described in section 4.2.

We implemented our two-node detection tests as described in section 4. More specifically, we used the procedure in equation 4.1 with different values of S_0 and examined the number of significant two-node connections. The mean number of such connections for the 100 replicates is shown in Table 1. Consider first the cell corresponding to $S = S_0 = 1$. This is the null case where there are no connections and we are testing the hypothesis of independence. At significance level $\alpha = 0.05$, we would expect about $16,200 \times 0.05$, or approximately 800 of these tests to be declared as significant. However, only 200 connections were identified as significant. One of the main reasons for this is that the 16,200 tests are highly positively dependent. In such cases, one would expect fewer tests to be identified as significant. For the row with $S_0 = 2$, we see that on average, one edge is falsely detected as significant; when $S_0 = 3$, we detect zero connections, is the correct value. Looking at moderate-sized values of S (ranging from 2 to 4), we can see that the tests with $S_0 = 2$ miss some of the seven active edges. When the connection strength is stronger ($S > 5$), the tests with $S_0 = 2$ detect at least seven connections. A detailed analysis showed that all the active edges were indeed identified. For high values of S, we also detect some false edges. An inspection of Figure 3 shows that we may detect two false edges ($A[100]C$ and $G[10]D$) for the reasons discussed in section 4.3. An occurrence of $A[100]C$ will happen whenever an occurrence of $A[50]B$ is followed by an occurrence of $B[50]C$. Similarly, co-occurrence of $H[20]G$ and $H[30]D$ will lead to an occurrence of $G[10]D$.

To demonstrate this, we discuss the case with $S = 30$ for one particular simulated data set. We applied all the test procedures, including the false-edge detection tests in section 4.3. Table 2 shows the nine connections that were significant based on their Z_τ values ($S_0 = 2$). Consider first the group of neurons A, B, and C. We see that the episode $A[100]B$ has a negative Z_ξ value, indicating that it is a false edge, while $B[50]C$ has a large Z_η value, indicating that it is not a false edge. Thus, we conclude that these three neurons have a chain structure. Consider next the group of neurons D, H, and G. Here, the large value of Z_ξ and negative value of Z_η for $H[30]D$ and $G[10]D$, respectively, indicate that $G[10]D$ is a false edge while $H[30]D$ is not. Thus, we conclude that these three neurons have a fan-out structure. Finally, consider the group of neurons E, F, and I. Both the Z_ξ and Z_η values are large, indicating that neither $E[15]I$ nor $F[10]I$ is a false edge. Thus, we conclude that these three neurons have a feedforward loop structure.

5.1 Scalability. To assess the scalability of our approach, we analyzed simulated networks of various sizes and connection densities. We measured the accuracy and computational cost of our method, considering networks

Table 2: Test Statistics for the Initial Detection Test (Z_τ) and the False Edge Elimination Tests (Z_ξ, Z_η) for One Replicate with $S = 30$.

Neuron1[k]Neuron2	Z_τ	Z_ξ	Z_η
A[50]B	15.76	–	–
A[100]C	7.85	−3.70	–
B[50]C	18.04	–	15.72
E[5]F	16.15	–	–
E[15]I	17.85	13.09	–
F[10]I	20.37	–	16.34
G[10]D	8.70	–	−1.38
H[30]D	16.36	13.13	–
H[20]G	17.32	–	–

Table 3: Networks of n Neurons with $e = 0.01n^2$ Edges ($S = 30$, Delays 5 or 10 ms) Simulated and Analyzed with $S_0 = 2$ and $k = 1$ to 20 ms.

Neurons	$n = 10$	20	30	40	50	60	70	80	90	100
Edges	$e = 1$	4	9	16	25	36	49	64	81	100
FP (unpruned)	0.2	1.3	3.5	8.2	18.4	25.8	42.6	64.4	86.2	119.0
FP (pruned)	0.2	0.7	1.0	1.9	4.4	4.0	6.1	6.7	8.0	12.1

Notes: For each configuration, 10 random networks were generated. The mean number of false positives (FP) is much lower when the false edge elimination tests are applied (pruned) than if they are not applied (unpruned).

where the number of neurons ranged from $n = 10$ to 100, with $e = pn^2$ edges for $p = 0.01$ to 0.05. The connectivity was random, with connection strength $S = 30$ and delays of either 5 or 10 ms.

5.1.1 Accuracy. Table 3 shows the results of our methodology on data from networks of different sizes for a fixed density of connections ($p = 0.01$). With $S_0 = 2$ and considering delays from $k = 1$ to 20 ms, there were no false negatives for any n. However, as n increases, the number of false positives increases. The growth in false positives is substantially reduced by our false edge elimination tests. For example, for the networks with $n = e = 100$, these additional hypothesis tests eliminated approximately 90% of the false positives.

We then explored how the positive predictive value (defined as the proportion of positive test results that are true positives) of our two-phase procedure scales with network size and connection density. If only the screening phase is used, then increasing network size and connection density p leads to a decrease in positive predictive value (PPV). However, when the false edge elimination phase is employed, PPV remains high for all the
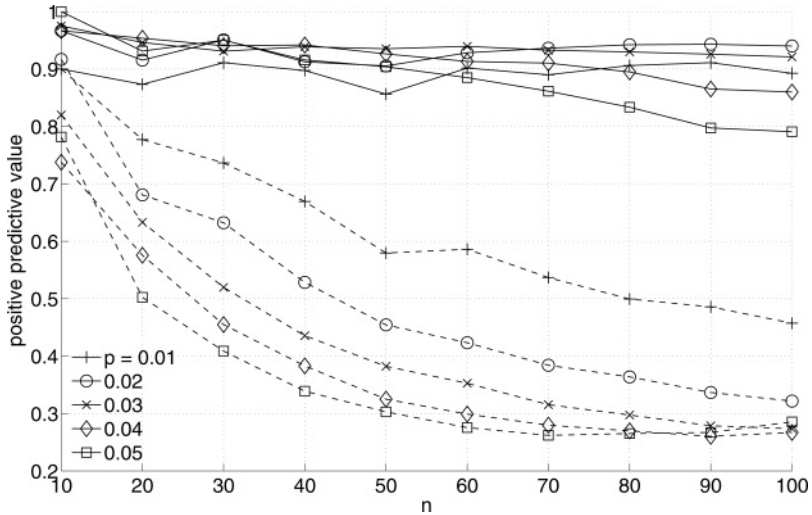
Figure 4: Scaling of PPV, defined as true positives divided by the sum of true positives and false positives, as a function of network size n and connection density p. Mean PPV for 10 random networks, with and without applying the false edge elimination tests (indicated by solid and dashed lines, respectively).

networks considered (see Figure 4). These results indicate that the performance of our two-phase methodology is relatively insensitive to increases in network size and connection density.

   5.1.2 Computational Cost. The bulk of the computational time involves counting all the two-node episodes in phase 1 of the strategy. We cannot think of any statistical technique for this problem where these counts will not be needed. Furthermore, as we noted, computationally efficient algorithms have been developed for counting the nonoverlapping occurrences of episodes (Laxman et al., 2005, 2007; Patnaik et al., 2008). Implementing the pruning step in phase 2 takes relatively less time. For example, a complete analysis of the largest and densest network in Figure 4 ($n = 100$, $e = 500$) took 68.8 minutes on a standard laptop computer (MacBook Pro, 2.4 GHz Intel Core i7 processor, 8 GB 1600 MHz DDR3 memory). For networks with $p = 0.01$, counting the three-node episodes required for implementation of phase 2 (false edge elimination tests) accounted for less than 1% of this computational time. These figures can be reduced on faster computational platforms and using parallel computing.
   As the number of neurons n increases, the computational time scales linearly with the number of two-node episodes ($\sim n^2$) to be counted (see Figure 5). For networks with $n = 100$, as p is increased, the amount of time
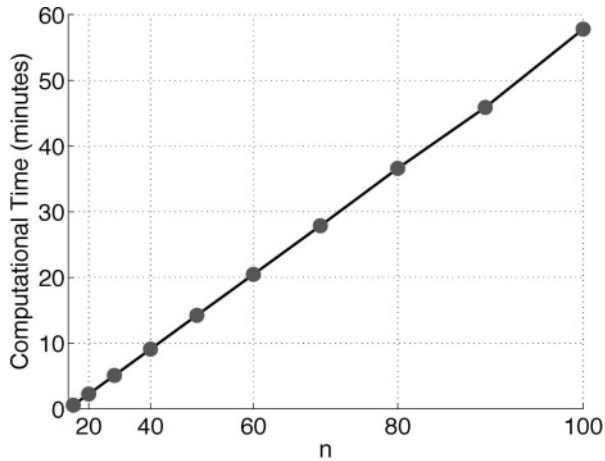
Figure 5: Computational time for analysis of the $p = 0.01$ networks shown in Figure 4. The time required scales linearly with the number of two-node episodes to be counted, which is given by n^2 times the number of delays k being considered (here, 20 different delays were used).

spent counting three-node episodes increases but remains less than 10% of total computational time at $p = 0.05$.

5.1.3 Networks with Non-DAG Structure. We have assumed throughout that the underlying networks have a DAG structure. It is much more difficult to develop the distribution theory for the statistical procedures for non-DAG structures. We can do this under a specified non-DAG network structure, but it is not possible to do this under a general unspecified structure. In addition, and more importantly, there can be identifiability as well as interpretability problems with non-DAG structures. For example, suppose $A \rightarrow B$, $B \rightarrow C$, and $C \rightarrow A$. Further, suppose that the first two connections are excitatory and the last one is inhibitory. Then it becomes very difficult to untangle the opposing relationships and reconstruct the connections purely from the spike train data. Nevertheless, in reality, one has to deal with non-DAG networks, so we have examined the potential limitation of this assumption through a simulation study. The simulation studies in Table 3 and Figure 4 were done under the assumption of a DAG network; the random connections were constrained to avoid cycles. We repeated our analysis on non-DAG networks generated without this constraint. Figure 6 shows that the results for DAG and non-DAG networks are very similar in terms of PPV. While it is not possible to generalize from a single study, the results for these configurations are encouraging. However, as we have noted, this problem requires a more thorough investigation in terms of both
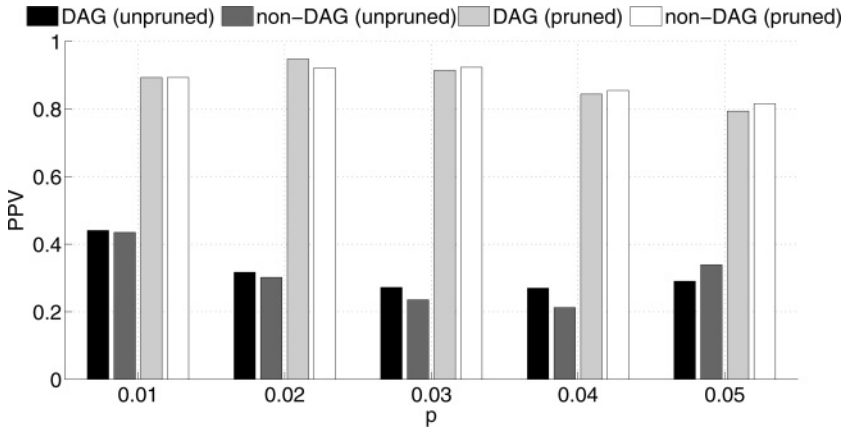
Figure 6: PPV for networks with $n = 100$ and various connection densities p is similar whether the random connections are unconstrained (non-DAG) or constrained to avoid cycles (DAG).

the performance on non-DAG structures and the implications for identifiability and interpretability. We note that this is an issue not just for the methods in this letter but for many others in the literature as well.

5.2 Effect of Subsampling. Spike train data sets recorded from mammalian tissues are typically from a small subset of the neuronal network active in the tissue. This subsampling is due to the large number of neurons present in mammalian brain structures and technical limitations on the number of neurons that can be recorded from simultaneously. So it is possible (and even likely) that the connections discovered in the observed network are not true connections and can be explained by chains and fan-out structures in the unobserved network. This problem is, of course, present in any statistical method that relies on data from the observed network. Thus, connections discovered in any observed network must be considered as functional connections and cannot be assumed to be a direct connection in the large network.

Nevertheless, we used a simulation study to assess the effect of unobserved neurons on the ability to discover functional connectivity. We simulated a network of 500 neurons and then randomly subsampled 20 of these neurons on which to perform our analysis. The original network had 2500 random connections of strength $S = 30$ and random delays of either 5 or 10 ms. For networks with 500 neurons and 2500 connections, we would expect 20-neuron subsamples to yield on average $(20^2/500^2)(2500) = 4$ direct connections between neurons in the subsample. For the network and subsample analyzed here, 3 of the 2500 connections were between neurons

Table 4: Subsamples of Size $n_s = 20$ to 100 Neurons from a 500-Neuron Network Simulation (2500 edges, $S = 30$, delays 5 or 10 ms) Analyzed with $S_0 = 2$ to 4.

Subsamples	$n_s = 20$	30	40	50	60	70	80	90	100
TP	3	12	17	29	38	45	60	81	92
FN	0	0	0	0	0	0	0	0	0
FP ($S_0 = 2$)	5	13	27	45	71	95	116	144	163
FP ($S_0 = 3$)	0	1	4	5	11	18	24	30	34
FP ($S_0 = 4$)	0	1	2	2	4	4	4	6	7

Notes: Increasing S_0 did not affect the number of true positives (TP) or false negatives (FN) but did reduce the number of false positives. Most of these false positives are due to partial observation of chains and fan-out structures present in the 500-neuron network—cases where the middle neuron in a chain or the first neuron in a fan-out is not in the observed sample.

in our observed subsample. When we subject these 20 neurons to our analysis procedure, there are 9 significant two-node episodes with $S_0 = 2$ (before false edge elimination). Three of these are true direct connections between neurons in the subsample: 5[10]11, 5[10]15, and 11[5]16. The 5[10]11 and 11[5]6 connections form a chain that lead to 5[15]6 being detected as significant; this is a false edge, which is eliminated by hypothesis test 4.3. The other five episodes are due to chains or fan-outs involving unobserved neurons. For example, the significant episode 1[10]19 results from true edges 1[5]479 and 479[5]19, whereas the significant episode 13[5]4 comes from edges 186[5]13 and 186[10]14.

As the number of neurons sampled is increased from 20 to 100, the number of true positives (observed direct connections) increases, but so does the number of false positives. The majority of these false positives can be explained by the existence of chains where neuron B is unobserved, or fan-outs where neuron A is unobserved (see Figure 1 and Table 4). The estimated strengths of the false connections induced by the partially observed chains and fan-outs are much less than the estimated strengths of the true connections. Thus, moderately increasing S_0 is an effective means of reducing false positives while still avoiding false negatives.

6 Application to Cultured Cortical Neurons

To understand the interactions of neurons in a network, neurobiologists conduct experiments in which they record the activity of several neurons from a region of the brain simultaneously. For in vitro experiments, the neurons often are cultured on top of a grid of microelectrodes known as a multielectrode array (MEA). A typical MEA setup consists of 8×8 grid of 64 electrodes. Each electrode is capable of recording the activity of one or more neurons around it. Some of the brain areas of interest for such

experiments are the primary motor cortex, hippocampus, and prefrontal cortex, as well as retinal ganglion cells.

After suitable preprocessing of the data, an MEA recording provides spike trains on each electrode, which can be expressed as an event stream. We used our methods to analyze spike train data available in the literature on cultures of cortical neurons (Rolston, Wagenaar, & Potter, 2007; Wagenaar, Pine, & Potter, 2006). In Wagenaar et al. (2006), 'population bursts' in the spiking activity of the cultures were characterized. Such bursts, defined as brief periods of time during which the firing rate of several cells or electrodes greatly exceeds the baseline rate, are a common feature observed in cultures of many different types of neurons. Our focus is not to characterize the bursts but rather to detect precisely timed spiking patterns involving multiple neurons in order to estimate the strength of functional connectivity between different neurons in the culture. Since our methods assume that the firing rates of individual neurons are relatively stationary in the analysis window, we selected a segment of the data that did not contain bursts.

We analyzed 120 seconds of data from culture 2-1-34. To visualize the functional connectivity present in the culture, we show the network graphs based on our pruned list of significant two-node episodes. If there is a significant episode with any delay for a pair of neurons (i, j), we draw a directed edge $i \rightarrow j$ connecting them. Figures 7 and 8 show the inferred network structure for two different S_0 thresholds. As S_0 is increased from $S_0 = 2$ to $S_0 = 10$ (see Figures 7 and 8), the weaker connections are no longer significant and the inferred network becomes sparser. In this way, our methods can be used to identify the strongest microcircuits present in the culture. Although we considered episodes with delays of up to 200 ms, all the significant episodes had delays of less than 10 ms. Fast delays such as this are consistent with the timescale of the action for AMPA, a common excitatory neurotransmitter in the cortex (Destexhe & Sejnowski, 2001; Platt, 2007).

7 Discussion

We have presented a two-phase approach for discovering the functional connectivity of neuronal networks from spike trains based on repeating occurrences of temporally precise patterns of spikes. Most methods for detecting repeated occurrences of precise spike patterns and assessing their significance involve correlations between time-shifted spike trains (Abeles & Gerstein, 1988; Tetko & Villa, 2001; Grün, Diesmann, & Aertsen, 2001; Pipa, Wheeler, Singer, & Nikolic, 2008; Diekman et al., 2009). Approaches using dynamic Bayesian networks (Patnaik, Laxman, & Ramakrishnan, 2011), Granger causality (Krumin & Shoham, 2010; Kim, Putrino, Ghosh, & Brown, 2011; Kispersky, Gutierrez, & Marder, 2011), and information theory (Quinn, Coleman, Kiyavash, & Hatsopoulos, 2011; Ito et al., 2011) have also been
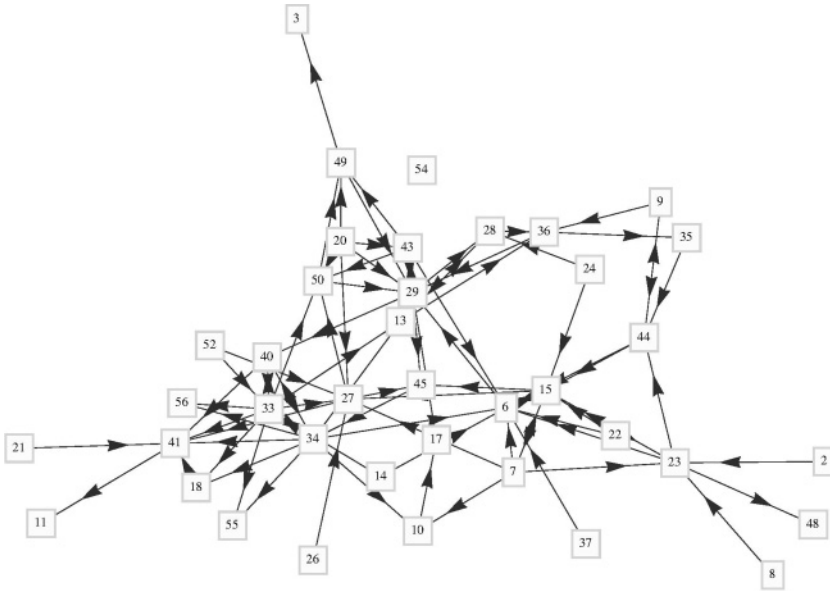
Figure 7: Inferred functional connectivity in cortical culture 2-1-34 with $S_0 = 2$.

employed. In contrast, the framework developed here is based on a computationally efficient temporal data mining method. Our hypothesis tests allow one to infer episodes that represent strong interactions among neurons and also avoid certain systematic errors commonly made by network inference procedures (Marbach et al., 2010). We have demonstrated the effectiveness of our methods on both simulated and in vitro neuronal spike data. We do note that the methodology developed here can be characterized in terms of (simple) Granger causality methods.

The strength of connection between a pair of neurons is characterized using the conditional probability that one neuron fires after a specified delay given that the other neuron has fired. This delay is usually caused by the axonal delay of spike propagation and the delay at synapses due to chemical diffusion. In our analysis, we have assumed that the interevent time constraints that represent the time delays in neuronal connections are constants. In general, there would be variation in these delays. Also, since we do not have complete control over which specific neurons are recorded from, we may be seeing some connections mediated by more than one synapse. Hence, one very useful extension for the method proposed here is to incorporate variable delays as interevent time constraints.

Currently, our statistical framework handles only serial episodes. Another interesting problem would be to extend the statistical theory for
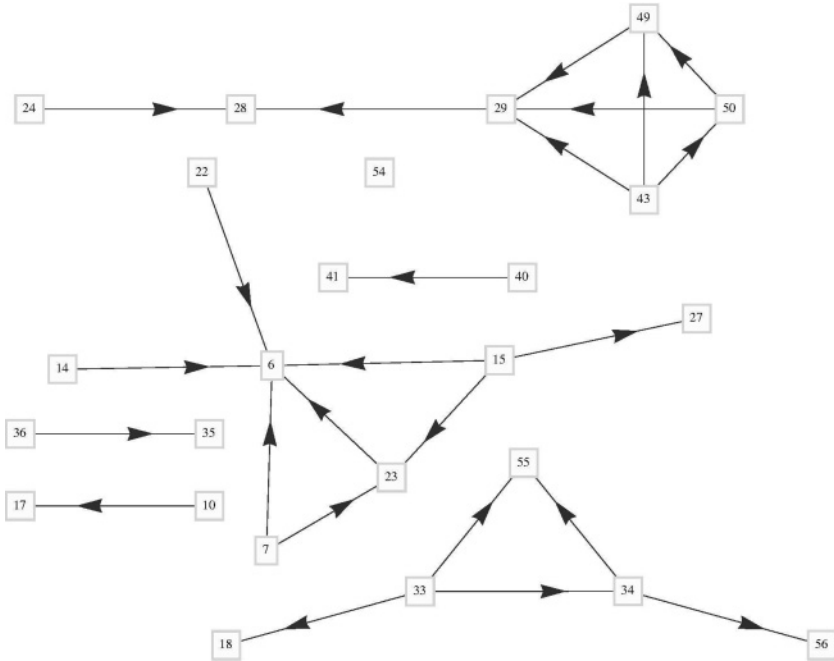
Figure 8: Inferred functional connectivity in cortical culture 2-1-34 with $S_0 = 10$.

the case of more general episodes that may be able to represent additional types of connectivity graphs. For example, Viswanathan, Sastry, and Unnikrishnan (2010) have developed a framework for determining the statistical significance of parallel episodes. Parallel episodes can capture the co-occurrence of spikes from many different neurons within a time window, a type of synchronous activity of interest to many researchers (Grün et al., 2001; Pipa et al., 2008).

Appendix: Derivation of Test Statistics

A.1 Variance of $\tau_N(S_0)$ Based on N. We derive the variance of the test statistic $\tau_N(S_0)$ based on the total number of occurrences. For simplicity of notation, we denote $A[k]B$ as AB. Recall that N_{AB} is binomial($L - k$, P_{AB}), where P_{AB} is the probability of an occurrence of episode AB as defined in section 3.1. Equation 4.2 for the variance is repeated below, and it is a sum of three expressions:

$$\text{Var}\left(\hat{\tau}_N(S_0)\right) = \text{Var}\left(\hat{P}_{AB}\right) + S_0^2\,\text{Var}\left(\hat{P}_A\hat{P}_B\right) - 2S_0\,\text{Cov}\left(\hat{P}_{AB},\hat{P}_A\hat{P}_B\right). \quad \text{(A.1)}$$

From the binomial property of N_{AB}, the first term is

$$\text{Var}\left(\hat{P}_{AB}\right) = \frac{P_{AB}(1 - P_{AB})}{(L - k)}. \tag{A.2}$$

We now show that the second term, up to the constant S_0^2, is

$$\begin{aligned}
\text{Var}\left(\hat{P}_A \hat{P}_B\right) = {} & \frac{P_{AB}}{(L - k)^3}\left(1 - P_{AB} - P_A - P_B + 5P_A P_B\right) \\
& + \frac{P_{AB}}{(L - k)^2}\left(P_{AB} + P_A + P_B - 8P_A P_B\right) \\
& - \frac{P_A P_B}{(L - k)^3}\left(1 - 2P_A - 2P_B + 6P_A P_B\right) \\
& + \frac{P_A P_B}{(L - k)^2}\left(1 - 3P_A - 3P_B + 10P_A P_B\right) \\
& + \frac{P_A P_B}{(L - k)}\left(P_A + P_B + 2P_{AB} - 4P_A P_B\right),
\end{aligned} \tag{A.3}$$

and the third term, up to the constant $-2S_0$, is

$$\begin{aligned}
\text{Cov}\left(\hat{P}_{AB}, \hat{P}_A \hat{P}_B\right) = {} & \frac{1}{(L - k)^2}\left(1 - P_{AB} - P_{AB}P_A - P_{AB}P_B + 2P_{AB}P_A P_B\right) \\
& - \frac{1}{(L - k)}\left(P_{AB}^2 - P_{AB}P_A - P_{AB}P_B + 2P_{AB}P_A P_B\right) \\
& + P_{AB}^2. \tag{A.4}
\end{aligned}$$

We start with the second term, or equation A.3. Note that we are computing the variance of $\text{Var}(\hat{P}_A \hat{P}_B)$ under possible dependence, so these two terms are not independent. We can write

$$\begin{aligned}
N_A N_B &= \sum_{t=1}^{L-k} I_A(t) \sum_{r=1}^{L-k} I_B(r + k) \\
&= \sum_{t=1}^{L-k} I_A(t) I_B(t + k) + \sum_{s=1}^{L-k} \sum_{r \neq s} I_A(s) I_B(r + k) \\
&= N_{AB} + N'_{AB}.
\end{aligned}$$

Thus,

$$\text{Var}(N_A N_B) = \text{Var}(N_{AB}) + \text{Var}\left(N'_{AB}\right) + 2\,\text{Cov}\left(N_{AB}, N'_{AB}\right).$$

Recall that $\mathrm{Var}(N_{AB}) = (L-k)P_{AB}(1-P_{AB})$. Further, it can be shown, through some tedious calculations, that $\mathrm{Cov}(N_{AB}, N'_{AB}) = (L-k)(L-k-1)$ $P_{AB}(P_A + P_B - 2P_A P_B)$ and that

$$
\begin{aligned}
\mathrm{Var}\left(N'_{AB}\right) = {} & (L-k)(L-k-1)P_A P_B(1-P_A P_B) \\
& + (L-k)(L-k-1)(L-k-2)P_A P_B(P_A + P_B) \\
& + (L-k)^2 P_{AB}^2 + (L-k)(2(L-k)^2 - 6(L-k) + 3)P_{AB}P_A P_B \\
& + (L-k)(L-k-1)(L-k-2)(L-k-3)\left(P_A P_B\right)^2 \\
& - (L-k)(L-k-1)((L-k)^2 - (L-k) - 1)\left(P_A P_B\right)^2.
\end{aligned}
$$

We can now get equation A.3 by combining these expression and using the fact that $\mathrm{Var}(\hat{P}_A \hat{P}_B) = \frac{1}{(L-k)^4} \mathrm{Var}(N_A N_B)$.

Finally, consider the third term, or equation A.4. Write

$$
\mathrm{Cov}(N_{AB}, N_A N_B) = \mathrm{E}\left(N_{AB}N_A N_B\right) - \mathrm{E}\left(N_{AB}\right)\mathrm{E}\left(N_A N_B\right). \tag{A.5}
$$

We can express the first term on the right as

$$
\begin{aligned}
\mathrm{E}\left(N_{AB}N_A N_B\right) = \mathrm{E}\Bigg[& \sum_{t=1}^{L-k} I_A(t)I_B(t+k) \sum_{s=1}^{L-k} I_A(s)I_B(s+k) \\
& + \sum_{t=1}^{L-k} I_A(t)I_B(t+k) \sum_{s=1}^{L-k}\sum_{r\neq s} I_A(s)I_B(r+k) \Bigg] \\
= {} & \mathrm{E}\left[N_{AB}^2\right] + \mathrm{E}\left[N_{AB}N'_{AB}\right]. \tag{A.6}
\end{aligned}
$$

Further,

$$
\mathrm{E}\left[N_{AB}^2\right] = (L-k)P_{AB}(1-P_{AB}) + (L-k)^2 P_{AB}^2, \tag{A.7}
$$

and

$$
\begin{aligned}
\mathrm{E}\left[N_{AB}N'_{AB}\right] = \mathrm{E}\Bigg[& \sum_{t=1}^{L-k} I_A(t)I_B(t+k) \sum_{s=1}^{L-k}\sum_{r\neq s} I_A(s)I_B(r+k) \Bigg] \tag{A.8} \\
& = (L-k)(L-k-1)P_{AB}(P_A + P_B + (L-k-2)P_A P_B).
\end{aligned}
$$

Moreover, the last term in equation A.5 (again noting that the two terms are not independent) can be calculated as

$$
\begin{aligned}
\mathrm{E}\left[N_A N_B\right] = {} & \mathrm{E}\left[N_{AB}\right] + \mathrm{E}\left[N'_{AB}\right] \\
& = (L-k)P_{AB} + (L-k)(L-k-1)P_A P_B. \tag{A.9}
\end{aligned}
$$

Combining the results in equations A.5 to A.9 and noting that $\mathrm{Cov}(\hat{P}_{AB}, \hat{P}_A\hat{P}_B) = \mathrm{Cov}(N_{AB}, N_A N_B)/(L-k)^3$, we get equation A.4.

A.2 Variances of the Test Statistics for False Edge Elimination Based on N. We first derive the variance of the test statistic $\hat{\xi}_N$ based on the total number of occurrences N. It can be written as the sum of three terms:

$$\mathrm{Var}\left(\hat{\xi}_N\right) = \mathrm{Var}\left(\hat{P}_{A\bar{B}C}\right) + \mathrm{Var}\left(\hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) - 2\,\mathrm{Cov}\left(\hat{P}_{A\bar{B}C}, \hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right). \quad \text{(A.10)}$$

Let k_1 and k_2 be the time delays, that is, the event of interest is $A[k_1]\bar{B}[k_2]C$. From the binomial property of $N_{A\bar{B}C}$, the first term in equation A.10 is

$$\mathrm{Var}\left(\hat{P}_{A\bar{B}C}\right) = \frac{P_{A\bar{B}C}(1 - P_{A\bar{B}C})}{(L - k_1 - k_2)}. \quad \text{(A.11)}$$

We will show that the second term in equation A.10 is given by

$$\begin{aligned}
&\mathrm{Var}\left(\hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) \\
&= \frac{1}{(L - k_1 - k_2)^3}P_A P_{\bar{A}} P_B P_{\bar{B}} P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)^2}P_A P_{\bar{A}} P_B P_{\bar{B}} P_C^2 \\
&\quad + \frac{1}{(L - k_1 - k_2)^2}P_A P_{\bar{A}} P_{\bar{B}}^2 P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)^2}P_A^2 P_B P_{\bar{B}} P_C P_{\bar{C}} \\
&\quad + \frac{1}{(L - k_1 - k_2)}P_A^2 P_{\bar{B}}^2 P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)}P_A^2 P_B P_{\bar{B}} P_C^2 \\
&\quad + \frac{1}{(L - k_1 - k_2)}P_A P_{\bar{A}} P_{\bar{B}}^2 P_C^2,
\end{aligned}$$

and that the third term in equation A.10 is given by

$$\begin{aligned}
\mathrm{Cov}\left(\hat{P}_{A\bar{B}C}, \hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) = {} &\frac{1}{(L - k_1 - k_2)^3}P_{A\bar{B}C}\bigg[\left(1 - P_{A\bar{B}C} - P_A - P_{\bar{B}} - P_C\right. \\
&+ 2P_A P_{\bar{B}} + 2P_A P_C + 2P_{\bar{B}} P_C - 6P_A P_{\bar{B}} P_C\big) \\
&+ \frac{1}{(L - k_1 - k_2)^3}\left(P_{A\bar{B}C} + P_A + P_{\bar{B}} + P_C\right. \\
&- 3P_A P_{\bar{B}} - 3P_A P_C - 3P_{\bar{B}} P_C\big) + \frac{1}{(L - k_1 - k_2)^2} \\
&\times \left(P_A P_{\bar{B}} + P_A P_C + P_{\bar{B}} P_C - 6P_A P_{\bar{B}} P_C\right)\bigg].
\end{aligned}$$

We start with the second term in equation A.10 and note that N_A, $N_{\bar{B}}$, and N_C are independent under H_0. Therefore,

$$
\begin{aligned}
\mathrm{Var}(N_A N_{\bar{B}} N_C) =\ & \mathrm{Var}(N_A)\mathrm{Var}(N_{\bar{B}})\mathrm{Var}(N_C) + \mathrm{Var}(N_A)\mathrm{Var}(N_{\bar{B}})\big[\mathrm{E}\,(N_C)\big]^2 \\
& + \mathrm{Var}(N_A)\,\mathrm{Var}(N_C)\big[\mathrm{E}\,(N_{\bar{B}})\big]^2 \\
& + \mathrm{Var}(N_{\bar{B}})\,\mathrm{Var}(N_C)\big[\mathrm{E}\,(N_A)\big]^2 \\
& + \big[\mathrm{E}\,(N_A)\big]^2\big[\mathrm{E}\,(N_{\bar{B}})\big]^2\,\mathrm{Var}(N_C) \\
& + \big[\mathrm{E}\,(N_A)\big]^2\big[\mathrm{E}\,(N_C)\big]^2\,\mathrm{Var}(N_{\bar{B}}) \\
& + \big[\mathrm{E}\,(N_{\bar{B}})\big]^2\big[\mathrm{E}\,(N_C)\big]^2\,\mathrm{Var}(N_A).
\end{aligned}
$$

We can get the answer for the second term as

$$
\mathrm{Var}\,\big(\hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\big) = \frac{1}{\left(L - k_1 - k_2\right)^6}\,\mathrm{Var}(N_A N_{\bar{B}} N_C). \tag{A.12}
$$

Turning to the third term in equation A.10,

$$
\begin{aligned}
\mathrm{Cov}(N_{A\bar{B}C},\, N_A N_{\bar{B}} N_C) &= \mathrm{E}\,(N_{A\bar{B}C}N_A N_{\bar{B}} N_C) - \mathrm{E}\,(N_{A\bar{B}C})\mathrm{E}\,(N_A N_{\bar{B}} N_C) \\
\mathrm{E}\,(N_{A\bar{B}C}N_A N_{\bar{B}} N_C) &= \mathrm{E}\big[N_{A\bar{B}C}^2\big] + \mathrm{E}\big[N_{A\bar{B}C}N'_{A\bar{B}C}\big]. \tag{A.13}
\end{aligned}
$$

These expectations can be obtained as

$$
\mathrm{E}\big[N_{A\bar{B}C}^2\big] = (L - k_1 - k_2)(1 - P_{A\bar{B}C}) + \left(L - k_1 - k_2\right)^2 P_{A\bar{B}C}^2 \tag{A.14}
$$

and

$$
\begin{aligned}
\mathrm{E}\big[N_{A\bar{B}C}N'_{A\bar{B}C}\big] = \mathrm{E}\Bigg[&\sum_{t=1}^{L-k_1-k_2} I_A(t)I_{\bar{B}}(t + k_1)I_C(t + k_1 + k_2) \\
&\times \sum_{t=1}^{L-k_1-k_2}\sum_{s\neq t}\sum_{r\neq s\neq t} I_A(t)I_{\bar{B}}(s + k_1)I_C(r + k_1 + k_2)\Bigg] \\
= (L - k_1 - k_2)&(L - k_1 - k_2 - 1)P_{A\bar{B}C}[P_A + P_{\bar{B}} + P_C \\
+ (L - k_1 - k_2 - 2)&P_A P_{\bar{B}} \\
+ (L - k_1 - k_2 - 2)&P_A P_C + (L - k_1 - k_2 - 2)P_{\bar{B}}P_C \\
+ (L - k_1 - k_2 - 2)&(L - k_1 - k_2 - 3)P_A P_{\bar{B}}P_C]. \tag{A.15}
\end{aligned}
$$

Combining equations A.13 to A.15 and using the fact that

$$\text{Cov}\left(\hat{P}_{A\bar{B}C}, \hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) = \frac{\text{Cov}\left(N_{A\bar{B}C}, N_A N_{\bar{B}} N_C\right)}{\left(L - k_1 - k_2\right)^4}, \tag{A.16}$$

we get the final answer.

The variance of η_N can be obtained using similar arguments. We provide just the expressions and omit the derivations. Thus,

$$\text{Var}\left(\hat{\eta}_N\right) = \text{Var}\left(\hat{P}_{\bar{A}BC}\right) + \text{Var}\left(\hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right) - 2\,\text{Cov}\left(\hat{P}_{\bar{A}BC}, \hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right). \tag{A.17}$$

The first term in equation A.17 is

$$\text{Var}\left(\hat{P}_{\bar{A}BC}\right) = \frac{P_{\bar{A}BC}\left(1 - P_{\bar{A}BC}\right)}{\left(L - k_1 - k_2\right)}. \tag{A.18}$$

The second term is

$$
\begin{aligned}
&\text{Var}\left(\hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right) \\
&= \frac{1}{(L - k_1 - k_2)^3} P_A P_{\bar{A}} P_B P_{\bar{B}} P_C P_{\bar{C}} + \frac{1}{\left(L - k_1 - k_2\right)^2} P_A P_{\bar{A}} P_B P_{\bar{B}} P_C^2 \\
&\quad + \frac{1}{\left(L - k_1 - k_2\right)^2} P_A P_{\bar{A}} P_B^2 P_C P_{\bar{C}} + \frac{1}{\left(L - k_1 - k_2\right)^2} P_{\bar{A}}^2 P_B P_{\bar{B}} P_C P_{\bar{C}} \\
&\quad + \frac{1}{\left(L - k_1 - k_2\right)} P_{\bar{A}}^2 P_B^2 P_C P_{\bar{C}} + \frac{1}{\left(L - k_1 - k_2\right)} P_{\bar{A}}^2 P_B P_{\bar{B}} P_C^2 \\
&\quad + \frac{1}{\left(L - k_1 - k_2\right)} P_A P_{\bar{A}} P_B^2 P_C^2
\end{aligned}
$$

Further,

$$
\begin{aligned}
&\text{Cov}\left(\hat{P}_{\bar{A}BC}, \hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right) \\
&= \frac{1}{\left(L - k_1 - k_2\right)^3} P_{\bar{A}BC}\Bigg[(1 - P_{\bar{A}BC} - P_{\bar{A}} - P_B - P_C + 2 P_{\bar{A}} P_B + 2 P_{\bar{A}} P_C \\
&\quad + 2 P_B P_C - 6 P_{\bar{A}} P_B P_C) + \frac{1}{\left(L - k_1 - k_2\right)^3} (P_{\bar{A}BC} + P_{\bar{A}} + P_B + P_C - 3 P_{\bar{A}} P_B \\
&\quad - 3 P_{\bar{A}} P_C - 3 P_B P_C) + \frac{1}{\left(L - k_1 - k_2\right)^2} (P_{\bar{A}} P_B + P_{\bar{A}} P_C + P_B P_C - 6 P_A P_B P_C)\Bigg].
\end{aligned}
$$

A.3 Variance of $\tau_M(S_0)$ Based on M. All the estimators in this section are based on the number of nonoverlapping occurrences of an episode denoted M_{AB}. We compute the variance of the test statistic $\hat{\tau}_M(S_0)$, which is the sum of three terms,

$$\text{Var}\left(\hat{\tau}_M\right) = \text{Var}\left(\hat{P}_{AB}\right) + S_0^2 \text{Var}\left(\hat{P}_A\hat{P}_B\right) - 2S_0 \text{Cov}\left(\hat{P}_{AB}, \hat{P}_A\hat{P}_B\right), \quad \text{(A.19)}$$

where the first term is

$$\text{Var}\left(\hat{P}_{AB}\right) \approx \frac{(1 + kP_{AB})P_{AB}(1 - P_{AB})}{(L - k)}, \quad \text{(A.20)}$$

the second term, up to the constant S_0^2, is

$$\text{Var}\left(\hat{P}_A\hat{P}_B\right) = \frac{P_B}{(L-k)^3}(1 - P_{AB} - P_A - P_B + 5P_AP_B)$$

$$+ \frac{P_{AB}}{(L-k)^2}(P_{AB} + P_A + P_B - 8P_AP_B)$$

$$- \frac{P_AP_B}{(L-k)^3}(1 - 2P_A - 2P_B + 6P_AP_B)$$

$$+ \frac{P_AP_B}{(L-k)^2}(1 - 3P_A - 3P_B + 10P_AP_B)$$

$$+ \frac{P_AP_B}{(L-k)}(P_A + P_B + 2P_{AB} - 4P_AP_B), \quad \text{(A.21)}$$

and the third term, up to the constant $-2S_0$, is

$$\text{Cov}\left(\hat{P}_{AB}, \hat{P}_A\hat{P}_B\right) \approx (1 + kP_{AB})\left[\frac{P_{AB}(1 - P_{AB})}{(L-k)^2(1 + kP_{AB})}\right.$$

$$- \frac{1}{(L-k)^2}(P_{AB}P_A + P_{AB}P_B - 2P_AP_B)$$

$$\left. - \frac{1}{(L-k)}\left(P_{AB}^2 + P_{AB}P_A + P_{AB}P_B - 3P_AP_B\right) + P_AP_B\right].$$

$$\text{(A.22)}$$

We now derive each of these expressions. From equation 3.7, the first term is

$$\text{Var}\left(\hat{P}_{AB}\right) \approx \frac{(1 + kP_{AB})P_{AB}(1 - P_{AB})}{(L - k)}. \quad \text{(A.23)}$$

The second term, $\text{Var}(\hat{P}_A\hat{P}_B)$, is the same as in section A.1. To compute the third term, write

$$\text{Cov}(M_{AB}, N_A N_B) = \text{E}(M_{AB}N_A N_B) - \text{E}(M_{AB})\text{E}(N_A N_B)$$

$$\text{E}(M_{AB}N_A N_B) = \text{E}\left[\sum_{t^*=1}^{L-k} I_A(t^*)I_B(t^*+k) \sum_{s=1}^{L-k} I_A(s)I_B(s+k) \right.$$

$$\left. + \sum_{t^*=1}^{L-k} I_A(t^*)I_B(t^*+k) \sum_{s=1}^{L-k}\sum_{r\neq s} I_A(s)I_B(r+k) \right]$$

$$= \text{E}\left[M_{AB}N_{AB}\right] + \text{E}\left[M_{AB}N'_{AB}\right]. \qquad (A.24)$$

Here t^* refers to the nonoverlapped t's, and there are M_{AB} of them. After some algebraic manipulations, we get

$$\text{Cov}(M_{AB}, N_A N_B) = \frac{(L-k)P_{AB}}{(1+kP_{AB})}\left[(L-k-1)(P_A+P_B) \right.$$

$$+ (L-k-2)P_A P_B) + (1+kP_{AB})$$

$$+ (1-P_{AB})/(1+kP_{AB}) - (L-k)P_{AB}$$

$$\left. + (L-k)(L-k-1)P_A P_B \right].$$

Since

$$\text{Cov}\left(\hat{P}_{AB}, \hat{P}_A\hat{P}_B\right) \approx \frac{\text{Cov}(M_{AB}, N_A N_B)}{(L-k)\text{E}\left[M_{AB}\right]^2 \left(\frac{L-k}{\text{E}[M_{AB}]} - k\right)^2}, \qquad (A.25)$$

we get the final answer.

A.4 Variances of the Test Statistics for False Edge Elimination Based on M. We derive the expressions for the test statistics based on the number of nonoverlapping occurrences M. As before, the variance of $\hat{\xi}_M$ can be written as the sum of three terms,

$$\text{Var}\left(\hat{\xi}_M\right) = \text{Var}\left(\hat{P}_{A\bar{B}C}\right) + \text{Var}\left(\hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) - 2\,\text{Cov}\left(\hat{P}_{A\bar{B}C}, \hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right), \quad (A.26)$$

where the first term is

$$\text{Var}\left(\hat{P}_{A\bar{B}C}\right) \approx \frac{(1 + (k_1 + k_2)P_{A\bar{B}C})P_{A\bar{B}C}(1 - P_{A\bar{B}C})}{(L - k_1 - k_2)}, \tag{A.27}$$

the second term is

$$
\begin{aligned}
\text{Var}&\left(\hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) \\
&= \frac{1}{(L - k_1 - k_2)^3} P_A P_{\bar{A}} P_B P_{\bar{B}} P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)^2} P_A P_{\bar{A}} P_B P_{\bar{B}} P_C^2 \\
&\quad + \frac{1}{(L - k_1 - k_2)^2} P_A P_{\bar{A}} P_{\bar{B}}^2 P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)^2} P_A^2 P_B P_{\bar{B}} P_C P_{\bar{C}} \\
&\quad + \frac{1}{(L - k_1 - k_2)} P_A^2 P_{\bar{B}}^2 P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)} P_A^2 P_B P_{\bar{B}} P_C^2 \\
&\quad + \frac{1}{(L - k_1 - k_2)} P_A P_{\bar{A}} P_{\bar{B}}^2 P_C^2, \tag{A.28}
\end{aligned}
$$

and the third term is

$$
\begin{aligned}
\text{Cov}\left(\hat{P}_{A\bar{B}C}, \hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C\right) \approx \frac{P_{A\bar{B}C}(1 + (k_1 + k_2)P_{A\bar{B}C})}{(L - k_1 - k_2)^2} &\left[\frac{(1 - P_{A\bar{B}C})}{1 + k_1 + k_2 P_{A\bar{B}C}} \right. \\
&- (P_A + P_{\bar{B}} + P_C - 2P_A P_{\bar{B}} - 2P_A P_C - 2P_{\bar{B}} P_C \\
&+ 6P_A P_{\bar{B}} P_C) + (L - k_1 - k_2)(P_A + P_{\bar{B}} \\
&+ P_C - 3P_A P_{\bar{B}} - 3P_A P_C - 3P_{\bar{B}} P_C \\
&+ 11 P_A P_{\bar{B}} P_C) + (L - k_1 - k_2)^2 (P_A P_{\bar{B}} + P_A P_C \\
&\left. + P_{\bar{B}} P_C - 6P_A P_{\bar{B}} P_C) \right]. \tag{A.29}
\end{aligned}
$$

The expression for $\text{Var}(\hat{P}_{A\bar{B}C})$ can be derived in exactly the same way as that of $\text{Var}(\hat{P}_{AB})$ in section 3.3. The expression for $\text{Var}(\hat{P}_A\hat{P}_{\bar{B}}\hat{P}_C)$ was obtained in section A.2. Turning to the third term, note that

$$\text{Cov}(M_{A\bar{B}C}, N_A N_{\bar{B}} N_C) = \text{E}\left(M_{A\bar{B}C} N_A N_{\bar{B}} N_C\right) - \text{E}\left(M_{A\bar{B}C}\right)\text{E}\left(N_A N_{\bar{B}} N_C\right),$$

where

$$\text{E}\left(M_{A\bar{B}C} N_A N_{\bar{B}} N_C\right) = \text{E}\left[M_{A\bar{B}C} N_{A\bar{B}C}\right] + \text{E}\left[M_{A\bar{B}C} N'_{A\bar{B}C}\right].$$

Now,

$$E[M_{A\bar{B}C}N_{A\bar{B}C}] = E\left[M_{A\bar{B}C}\left(M_{A\bar{B}C} + \sum_{j=1}^{M_{A\bar{B}C}} R_j\right)\right]$$

$$= E\left[M_{A\bar{B}C}^2\right] + E\left[M_{A\bar{B}C}\sum_{j=1}^{M_{A\bar{B}C}} R_j\right]$$

$$= E\left[M_{A\bar{B}C}^2(1 + kP_{AB})\right] \tag{A.30}$$

and

$$E[M_{A\bar{B}C}N'_{A\bar{B}C}] = E\left[\sum_{t^*=1}^{L-k_1-k_2} I_A(t^*)I_{\bar{B}}(t^* + k_1)I_C(t^* + k_1 + k_2)\right.$$

$$\times \sum_{t=1}^{L-k_1-k_2} \sum_{s \neq t} \sum_{r \neq s \neq t} I_A(t)I_{\bar{B}}(s + k_1)I_C(r + k_1 + k_2)\bigg]$$

$$= \frac{(L - k_1 - k_2)P_{A\bar{B}C}}{(1 + (k_1 + k_2)P_{A\bar{B}C})}\bigg[(L - k_1 - k_2 - 1)(P_A + P_{\bar{B}} + P_C)$$

$$+ (L - k_1 - k_2 - 1)(L - k_1 - k_2 - 2)$$

$$\times (P_A P_{\bar{B}} + P_A P_C + P_{\bar{B}} P_C)$$

$$+ (L - k_1 - k_2 - 1)(L - k_1 - k_2 - 2)(L - k_1 - k_2 - 3)$$

$$\times (P_A P_{\bar{B}} P_C)\bigg]. \tag{A.31}$$

Here t^* refers to the nonoverlapped ts, and there are $M_{A\bar{B}C}$ of them. We get the final result by combining these equations.

The variance of $\hat{\eta}_M$ in equation 4.4 can be obtained similarly. We omit the details and provide only the expressions:

$$\text{Var}\left(\hat{\eta}_M\right) = \text{Var}\left(\hat{P}_{\bar{A}BC}\right) + \text{Var}\left(\hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right) - 2\text{Cov}\left(\hat{P}_{\bar{A}BC}, \hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right).$$

where

$$\text{Var}\left(\hat{P}_{\bar{A}BC}\right) \approx \frac{(1 + k_1 + k_2 P_{\bar{A}BC})P_{\bar{A}BC}(1 - P_{\bar{A}BC})}{(L - k_1 - k_2)},$$

$$\text{Var}\left(\hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right) = \frac{1}{\left(L - k_1 - k_2\right)^3}P_A P_{\bar{A}} P_B P_{\bar{B}} P_C P_{\bar{C}}$$

$$+ \frac{1}{(L - k_1 - k_2)^2} P_A P_{\bar{A}} P_B P_{\bar{B}} P_C^2$$

$$+ \frac{1}{(L - k_1 - k_2)^2} P_A P_{\bar{A}} P_B^2 P_C P_{\bar{C}}$$

$$+ \frac{1}{(L - k_1 - k_2)^2} P_{\bar{A}}^2 P_B P_{\bar{B}} P_C P_{\bar{C}} + \frac{1}{(L - k_1 - k_2)} P_{\bar{A}}^2 P_B^2 P_C P_{\bar{C}}$$

$$+ \frac{1}{(L - k_1 - k_2)} P_{\bar{A}}^2 P_B P_{\bar{B}} P_C^2 + \frac{1}{(L - k_1 - k_2)} P_A P_{\bar{A}} P_B^2 P_C^2$$

and

$$\mathrm{Cov}\left(\hat{P}_{\bar{A}BC}, \hat{P}_{\bar{A}}\hat{P}_B\hat{P}_C\right) \approx \frac{P_{\bar{A}BC}\left(1 + (k_1 + k_2)P_{\bar{A}BC}\right)}{(L - k_1 - k_2)^2}$$

$$\times \Big[(1 - P_{\bar{A}BC})/(1 + (k_1 + k_2)P_{\bar{A}BC})$$
$$- (P_{\bar{A}} + P_B + P_C - 2P_A P_B - 2P_{\bar{A}}P_C - 2P_B P_C$$
$$+ 6P_{\bar{A}}P_B P_C) + (L - k_1 - k_2)(P_{\bar{A}} + P_B + P_C - 3P_{\bar{A}}P_B$$
$$- 3P_{\bar{A}}P_C - 3P_B P_C + 11P_{\bar{A}}P_B P_C)$$
$$+ (L - k_1 - k_2)^2(P_{\bar{A}}P_B + P_{\bar{A}}P_C + P_B P_C - 6P_{\bar{A}}P_B P_C)\Big].$$

References

Abeles, M., & Gerstein, G. L. (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *Journal of Neurophysiology*, *60*, 909–924.

Achar, A., Laxman, S., & Sastry, P. S. (2011). A unified view of the a priori-based algorithms for frequent episode discovery. *Knowl. Inf. Syst.*, *31*, 223–250.

Amari, S., Nakahara, H., Wu, S., & Sakai, Y. (2003). Synchronous firing and higher-order interactions in neuron pool. *Neural Computation*, *15*, 127–142.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B (Methodological)*, *57*, 289–300.

Brillinger, D. R. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, *59*, 189–200.

Brown, E. N., Kass, R. E., & Mitra, P. M. (2004). Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nature Neuroscience*, *7*, 456–461.

Chornoboy, E., Schramm, L., & Karr, A. (1988). Maximum likelihood identification of neuronal point process systems. *Biological Cybernetics*, *59*, 265–275.

Destexhe, A., & Sejnowski, T. (2001). *Thalamocortical assemblies: How ion channels, single neurons, and large-scale networks organize sleep oscillations*. New York: Oxford University Press.

Diekman, C. O., Sastry, P. S., & Unnikrishnan, K. P. (2009). Statistical significance of sequential firing patterns in multi-neuronal spike trains. *J. Neurosci. Methods*, *182*, 279–284.

Fleishman, A. L. (1978). A method for simulating non-normal distributions. *Psychometrika*, *43*, 521–532.

Gerhard, F., Pipa, G., Lima, B., Neuenschwander, S., & Gerstner, W. (2011). Extraction of network topology from multi-electrode recordings: Is there a small-world effect? *Frontiers in Computational Neuroscience*, *5*. doi:10.3389/fncom.2011.00004

Grün, S., Diesmann, M., & Aertsen, A. (2001). Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Computation*, *14*, 43–80.

Ito, S., Hansen, M. E., Heiland, R., Lumsdaine, A., Litke, A. M., & Beggs, J. M. (2011). Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. *PLoS One*, *6*. doi:10.1371/journal.pone.0027431

Kass, R. E., & Ventura, V. (2001). A spike-train probability model. *Neural Computation*, *13*, 1713–1720.

Kass, R. E., Ventura, V., & Brown, E. N. (2005). Statistical issues in the analysis of neuronal data. *J. Neurophysiol.*, *94*, 8–25.

Kim, S., Putrino, D., Ghosh, S., & Brown, E. N. (2011). A Granger causality measure for point process models of ensemble neural spiking activity. *PLOS Comput Biol.*, *7*. doi:10.1371/journal.pcbi.1001110

Kispersky, T., Gutierrez, G. J., & Marder, E. (2011). Functional connectivity in a rhythmic inhibitory circuit using Granger causality. *Neural Systems and Circuits*, *1*. doi:10.1186/2042-1001-1-9

Krumin, M., & Shoham, S. (2010). Multivariate autoregressive modeling and granger causality analysis of multiple spike trains. *Comput. Intell. Neurosci*. doi:10.1155/2010/752428

Laxman, S., Sastry, P. S., & Unnikrishnan, K. P. (2005). Discovering frequent episodes and learning hidden Markov models: A formal connection. *IEEE Transactions on Knowledge and Engineering*, *17*, 1505–1517.

Laxman, S., Sastry, P. S., & Unnikrishnan, K. P. (2007). A fast algorithm for finding frequent episodes in event streams. In *Proceedings of the 13th ACM SIGKDD*

International Conference on Knowledge Discovery and Data Mining. New York: ACM Press.

Mannila, H., Toivonen, H., & Verkamo, A. (1997). Discovery of frequent episodes in event sequences. *Data Mining and Knowledge Discovery*, *1*, 259–289.

Marbach, D., Prill, R. J., Schaffter, T., Mattiussi, C., Floreano, D., & Stolovitsky, G. (2010). Revealing strengths and weaknesses of methods for gene network inference. *Proc. Natl. Acad. Sci. USA*, *107*, 6286–6291.

Martignon, L., Deco, G., Laskey, K., Diamond, M., Freiwald, W., & Vaadia, E. (2000). Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Computation*, *12*, 2621–2653.

Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., & Alon, U. (2002). Network motifs: Simple building blocks of complex networks. *Science*, *298*, 824–827.

Olson, C. R., Gettner, S. N., Ventura, V., Carta, R., & Kass, R. E. (2000). Neuronal activity in macaque supplementary eye field during planning of saccades in response to pattern and spatial cues. *J. Neurophysiol.*, *84*, 1369–1384.

Patnaik, D., Sastry, P. S., & Unnikrishnan, K. P. (2008). Inferring neuronal network connectivity from spike data: A temporal datamining approach. *Scientific Programming*, *16*, 49–77.

Patnaik, D., Laxman, S., & Ramakrishnan, N. (2011). Discovering excitatory relationships using dynamic Bayesian networks. *Knowl. Inf. Syst.*, *29*, 273–303.

Pipa, G., Wheeler, D. W., Singer, W., & Nikolic, D. (2008). NeuroXidence: Reliable and efficient analysis of an excess or deficiency of joint spike events. *J. Comput. Neurosci.*, *25*, 64–88.

Platt, S. R. (2007). The role of glutamate in central nervous system health and disease: A review. *Vet. J.*, *173*, 278–286.

Quinn, C. J., Coleman, T. P., Kiyavash, N., & Hatsopoulos, N. G. (2011). Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *J. Comput. Neurosci.*, *30*, 17–44.

Rigat, F., Gunst, M., & Pelt, J. (2006). Bayesian modelling and analysis of spatiotemporal neuronal networks. *Bayesian Analysis*, *1*, 733–764.

Rolston, J., Wagenaar, D. A., & Potter, S. M. (2007). Precisely timed spatiotemporal patterns of neural activity in dissociated cell cultures. *Neuroscience*, *148*, 294–303.

Sastry, P. S., & Unnikrishnan, K. P. (2010). Conditional probability-based significance tests for sequential patterns in multineuronal spike trains. *Neural Computation*, *22*, 1025–1059.

Schneidman, E., Berry, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, *440*, 1007–1012.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci.*, *18*, 3870–3896.

Shaffer, J. P. (1995). Multiple hypothesis testing. *Ann. Rev. Psych.*, *46*, 561–584.

Tetko, I. V., & Villa, A. E. P. (2001). A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains: 1. Detection of repeated patterns. *J. Neurosci. Methods*, *105*, 1–14.

Truccolo, W., Eden, U., Fellows, M., Donoghue, J., & Brown, E. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *J. Neurophysiol.*, *93*, 1074–1089.

Viswanathan, R., Sastry, P. S., & Unnikrishnan, K. P. (2010). *Efficient discovery of large synchronous events in neural spike streams.* arXiv:1006.1543v1

Wagenaar, D. A., Pine, J., & Potter, S. M. (2006). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci.*, *7.* doi:10.1186/1471-2202-7-11

Zhao, Y. G., Ono, T., & Ishii, K. (2002). Monte Carlo simulation using moments of random variables. *Journal of Asian Architecture and Building Engineering*, *1*, 13–20.