

## Network inference with confidence from multivariate time series

Mark A. Kramer,<sup>1,\*</sup> Uri T. Eden,<sup>1</sup> Sydney S. Cash,<sup>2</sup> and Eric D. Kolaczyk<sup>1</sup>

<sup>1</sup>*Department of Mathematics and Statistics, Boston University, Boston, Massachusetts 02215, USA*

<sup>2</sup>*Department of Neurology, Epilepsy Service, Harvard Medical School, ACC 835, Massachusetts General Hospital, 55 Fruit Street, Boston, Massachusetts 02114, USA*

(Received 9 March 2009; revised manuscript received 14 May 2009; published 11 June 2009)

Networks—collections of interacting elements or nodes—abound in the natural and manmade worlds. For many networks, complex spatiotemporal dynamics stem from patterns of physical interactions unknown to us. To infer these interactions, it is common to include edges between those nodes whose time series exhibit sufficient functional connectivity, typically defined as a measure of coupling exceeding a predetermined threshold. However, when uncertainty exists in the original network measurements, uncertainty in the inferred network is likely, and hence a statistical propagation of error is needed. In this manuscript, we describe a principled and systematic procedure for the inference of functional connectivity networks from multivariate time series data. Our procedure yields as output both the inferred network and a quantification of uncertainty of the most fundamental interest: uncertainty in the number of edges. To illustrate this approach, we apply a measure of linear coupling to simulated data and electrocorticogram data recorded from a human subject during an epileptic seizure. We demonstrate that the procedure is accurate and robust in both the determination of edges and the reporting of uncertainty associated with that determination.

DOI: [10.1103/PhysRevE.79.061916](https://doi.org/10.1103/PhysRevE.79.061916)

PACS number(s): 87.19.L-, 64.60.aq, 34.80.Pa, 89.75.-k

### I. INTRODUCTION

Many examples of natural and fabricated networks exist in the world, including airline networks, computer networks, and neural networks. To define a network is, in principle, straightforward: we simply identify a collection of nodes and edges [1–3]. A node (or vertex) represents a participant or actor in a network, while an edge represents a link or association between two nodes (Fig. 1). For example, in an airline network, individual airports constitute nodes and direct connections between airports identify edges. In a neural network, individual neurons and the physical connections between neurons determine the nodes and edges of the network, respectively. Having established network representations of these complex systems, we may then address pertinent issues, such as the worldwide spread of infectious disease through the airline network [4] or the effect of cortical lesions on brain dynamics [5].

The decision to link two nodes with an edge varies in difficulty [6]. In some cases, a known physical connection exists between two nodes, and the choice to include an edge is then obvious. Does an airline connect two cities or not [7]? Do two actors collaborate on a film or not [8,9]? Does a physical connection exist between two brain regions or not [10,11]? In these cases, the decision to include a link between nodes is simple and based on the known association or physical connection between two nodes.

In other cases, the interactions between nodes are obscure. For example, we may only observe the dynamic activity at individual nodes and have no access to the physical connections between nodes. In these cases, we may apply coupling measures to multivariate time series data associated with the node dynamics and attempt to infer their associations, with-

out explicit knowledge of their structural connections [12–14]. This approach has proven useful in, for example, climate studies [15,16] and also human brain studies, in which the structural connections between brain regions remain difficult to classify (although perhaps not for long [17]).

Viewed from a statistical perspective, two key challenges are inherent in this task of network inference: (i) appropriate interpretation of the coupling results in declaring network edges and (ii) accurate quantification of the uncertainty associated with the resulting network. The simplest—and, indeed, most common—method to interpret the coupling results and declare network edges involves comparison of the coupling strength to a threshold value [16,18–24]. If the coupling strength between two nodes exceeds this threshold, then we connect these nodes with an edge; otherwise, we leave the nodes unconnected. The number of edges in the resulting network depends critically on the choice of coupling threshold (as we illustrate schematically in Fig. 1). Fur-




FIG. 1. The number of edges in this four-node network depends upon our choice of coupling threshold. Each node (gray circle) represents a spatial location from which we record time series data. Applying a coupling measure to the dynamic activity recorded at node pairs, we obtain the normalized (between 0 and 1) coupling values shown in (a). If we choose the coupling threshold too low (0.05) we include edges between all nodes as in (b). If we choose the coupling threshold too high (0.95) we obtain no edges in the network as in (d). An intermediate choice of coupling threshold [0.65 in (c)] yields a different network.

\*mak@bu.edu

thermore, for a given choice of threshold, we expect a certain rate of error in (mis)declaring the connectivity status between pairs of nodes. This network uncertainty—intimately tied to choice of threshold—is often overlooked.

How do we choose such a threshold? One strategy is to apply a variety of different thresholds and examine the resulting collection of networks for robustness as a function of threshold [19,22,24]. This procedure of redundant analysis—which some of these authors recently employed [23]—is both time consuming and unsatisfying. Instead, a threshold should be chosen in a principled way, for example, one that links the choice of threshold with the achievement of a pre-specified level of network uncertainty. Such is the goal of this paper. In referring to “network uncertainty” many aspects of the network structure might be of interest (e.g., connectivity, degree distribution, or clustering). Here we will focus on the most basic aspect of network confidence: the presence or absence of network edges. Our particular goal is to equip the process of threshold-based inference of a network with a number quantifying the expected rate of falsely declared network edges. This number serves as a natural measure of network confidence.

In what follows, we adopt a statistical hypothesis testing paradigm to analyze multivariate time series data and create network representations of functional connectivity. The general paradigm involves three steps: (1) calculate the strength of coupling between time series data recorded at node pairs, (2) threshold each coupling measure through the use of a formal statistical hypothesis test, and (3) control the rate of falsely declared edges through the use of statistical multiple testing procedures. In Sec. II, we present a high-level outline of this general protocol, while in Sec. III we develop the procedure in detail, making specific choices of methodology for each step. We apply the protocol to three data sets in Sec. IV and show, in particular, that appropriate handling of the significance tests is vital. Fundamentally, the proposed protocol is a way of constructing functional networks that, rather than emerging as the result of some arbitrarily chosen coupling threshold, are composed of edges selected to achieve a guaranteed level of overall network accuracy. That is, it is a way of constructing networks with confidence.

## II. GENERAL PARADIGM

In this paper we are interested in the inference of networks (or, more precisely, graphs)  $G=(V,E)$  in which edges  $\{i,j\} \in E$  indicate a coupling (perhaps at nonzero lag) between time series  $x_i[t]$  and  $x_j[t]$  observed at  $N$  nodes  $i,j \in V$ . Our primary motivation is the desire to infer networks reflecting the functional (as opposed to structural) topology of neural systems [25]. This goal is reflected in our terminology, as well as in the numerical illustrations we present in Sec. IV, and prohibits our use of standard Gaussian graphical model techniques which assume independent measurements in time. However, the methods we propose—and the underlying principles upon which we base the methods—have quite general applicability.

We implement a procedure to create functional topologies from multivariate time series data that involves three general

steps. First, a coupling measure is specified and applied to the data, yielding a noisy indication of the functional connectivity between all nodes pairs. In the neurological data described below, this measure captures the extent of interactions between activities recorded simultaneously at separate spatial locations of the brain. Second, we develop significance tests appropriate for our choice of coupling measure, and associate a statistical  $p$ -value with each coupling result. Third, we analyze the resulting  $p$ -values using principles of statistical multiple testing to construct a network representation of the functional connectivity. In the course of this last step, we determine a number controlling the proportion of falsely inferred edges. We present this number as a basic and natural measure of network uncertainty.

Of course, the first step above implements a version of the standard approach to constructing such functional networks. In neuroscience, for example, investigators (including some of these authors) typically specify a measure of coupling and then assign edges between node pairs whose coupling is judged to be sufficiently strong. However, determination of just how strong is strong enough is invariably *ad hoc* or, at best, driven by “expert judgment.” As a result, there is no way to annotate the resulting networks with any indication of their inherent (in)accuracy. The subsequent steps in the proposed approach, therefore, are critically important to produce networks accompanied by accurate characterizations of their uncertainty. Put another way, we are interested here in the propagation of uncertainty in network inference, from the original time series data  $x_i[t]$  to a final assessment of network uncertainty. Our statistical hypothesis testing procedure, described above in three steps, achieves this goal. Furthermore, our numerical results indicate that it in fact does so in a robust fashion.

We achieve our goal primarily through careful attention to the interdependency among each of the three steps. In so doing, we also demonstrate how lack of such attention can lead to nonsensical network uncertainty statements. For example, the choice of coupling measure in the first step affects the hypotheses tested in the second step (i.e., the null hypothesis  $H_0$ : *No Coupling*, versus  $H_1$ : *Coupling*). The declaration of either edge or nonedge for each pair of nodes  $i,j \in V$  corresponds to either rejection of the null hypothesis or a failure to do so, respectively. If rejection is determined by comparison of the observed coupling values to a threshold, clearly the choice of threshold will affect the network results. But if we also wish to propagate uncertainty—from the original time series data, through the testing procedure, to the final network inferred—it is necessary to construct accurate probabilistic statements appropriate for the particular coupling measure we choose.

In constructing functional networks, we must consider the collection of individual hypothesis tests as a whole and the so-called “multiple testing problem” in statistics. To maintain the statistical power with which we can detect edges, we focus upon controlling the *rate* of falsely declared edges. Conditional on at least one edge being declared, the expected proportion of falsely declared edges here is equivalent to what is called the *false discovery rate* (FDR) in the statistical literature. The control of FDR in multiple testing situations, ranging from signal and image processing to genome-wide

association testing, has become a *de facto* standard technique for addressing the multiple testing problems (see [26], for example). In fact, the use of FDR control occurs with increasing frequency in the network literature as well (e.g., [27,28]). However, there is little evidence in this literature that the rates quoted are necessarily being achieved. We show later that it is quite easy, using a seemingly reasonable significance test, to end up with rates that are completely unrepresentative.

**III. IMPLEMENTATION OF THE GENERAL PARADIGM**

We described in the previous section three general steps to create a functional topology from multivariate time series data. In this section, we (1) define our coupling measure, the cross correlation, (2) develop appropriate significance tests for this measure, and (3) integrate these with a technique to account for multiple significance tests. In Sec. IV we apply these specific protocols to simulated and observed data, and show that the choice of the significance test is critical.

**A. Step 1: Choice of coupling measure**

The choice of coupling measure between pairs of time series permits many alternatives [29]. We may select a simple measure of linear coupling (e.g., the cross correlation [15,16,23,30,31] or the coherence [32–34]) or more sophisticated coupling measures (e.g., synchronization likelihood [20], wavelet coherence [35], Granger causality, and the related directed transfer function [36]). In this manuscript, we choose to focus on a simple measure of coupling based on the cross correlation; we discuss the application of nonlinear association measures to human voltage data in Sec. IV E. We also note that the general statistical hypothesis testing paradigm we adopt here can in principle be applied to any choice of coupling measure.

For a pair of time series  $x_i[t]$  and  $x_j[t]$  of lengths  $n$ , the sample cross correlation at lag  $\tau$  is defined as

$$C_{ij}[\tau] = \frac{1}{\hat{\sigma}_i \hat{\sigma}_j (n - 2\tau)} \sum_{t=1}^{n-\tau} (x_i[t] - \bar{x}_i)(x_j[t + \tau] - \bar{x}_j), \quad (1)$$

where  $\bar{x}_i$  and  $\bar{x}_j$  are the averages and  $\hat{\sigma}_i$  and  $\hat{\sigma}_j$  are the standard deviations of  $x_i[t]$  and  $x_j[t]$ , respectively, estimated from the data. This quantity can be computed efficiently over a range of lags using Fourier transform methods for convolutions. In our applications, we first transform the time series at each network node to have zero mean and unit variance, after which we compute the Fourier transforms  $\tilde{x}_i[\omega]$  and  $\tilde{x}_j[\omega]$ , multiply the first by the complex conjugate of the second, and take the inverse Fourier transform of the resulting product. For all of the data considered below, we compute the cross correlation for  $\tau$  ranging over indices between  $-100$  and  $100$  (the range of  $\tau$  in milliseconds depends on the sampling rate of the data, as we describe below).

Our formal measure of coupling will be the *maximal* cross correlation, i.e.,  $s_{ij} = \max_{\tau} |C_{ij}[\tau]|$ , the maximum of the absolute value of  $C_{ij}[\tau]$  over  $\tau$ . This measure will serve as our statistic for testing whether or not to assign an edge between nodes  $i$  and  $j$ , for each such pair of nodes.

**B. Step 2: Significance test**

Having chosen the test statistic  $s_{ij}$ , the maximal cross correlation between  $x_i[t]$  and  $x_j[t]$ , do we include a network edge between nodes  $i$  and  $j$ ? To answer this, we will use  $s_{ij}$  to test the null hypothesis that  $x_i[t]$  and  $x_j[t]$  are uncorrelated against the alternative that they are correlated. Rather than focusing on testing at a preassigned significance level, we will instead concentrate first on computing an appropriate  $p$ -value for each edge. Accurate evaluation of the  $p$ -values is critical to successful use of the false discovery rate principles we employ for the network inference problem here, as we discuss in Sec. III C. We compute the  $p$ -value in two different ways that make different assumptions about the coupling results. The first method is an analytic measure and specifically designed for our choice of coupling measure. The second is more general but computationally expensive. We define the measures below and, in the next section, apply each to simulated and observed time series data.

**1. Analytic method**

In this section we propose an analytic method. Frequently such methods involve comparison of a test statistic to a normal distribution. Following this approach, we would scale  $s_{ij}$  by an estimated variance, and then compare this scaled quantity to the standard normal distribution (i.e., with mean 0 and variance 1) to calculate a  $p$ -value. Here, however, this would be naive.

More specifically, under the stated null hypothesis of no coupling the statistic  $s_{ij}$  should have mean zero. A reasonable estimate of the variance of  $C_{ij}[l]$  under the null, motivated by a result of Bartlett [37,38], is given by

$$\widehat{\text{var}}(l) = \frac{1}{n-l} \sum_{\tau=-n}^n C_{ii}[\tau] C_{jj}[\tau], \quad (2)$$

where the  $C_{kk}[\tau]$  are the autocorrelations of time series  $k$  at lag  $\tau$ . This estimate takes a nontrivial form because the cross correlation will depend on the statistical properties of the underlying time series, and in particular on the autocovariance structure. Spurious cross correlations between the two times series are expected even if they are uncoupled [39], and this variance formula accounts for that.

Intuitively we might think to use Eq. (2) to define  $z_{ij} = s_{ij} / \sqrt{\widehat{\text{var}}(\hat{l}_{ij})}$ , where  $\hat{l}_{ij}$  is the lag corresponding to  $s_{ij}$  (i.e., the lag at which the maximum of the absolute value of the cross correlation occurs) and test the significance of the value  $z_{ij}$  by comparing it to the standard normal distribution. Unfortunately, although standardizing  $C_{ij}[l]$  by the estimated variance in Eq. (2) is sensible for any fixed  $l$ , use of the standard normal distribution with  $z_{ij}$  is not appropriate here, as we explain and illustrate below.

Two potential problems exist in using this naive method to determine the significance of  $s_{ij}$ . First, the distribution of the cross correlation  $C_{ij}[\tau]$ —strictly speaking—is normal only in the asymptotic case of large sample size  $n$ . In finite samples this approximation can be poor, particularly since the cross correlations are bounded between  $-1$  and  $1$ , while the normal distribution varies over an unbounded range. Sec-

ond, we choose  $s_{ij}$  as an extremum of the cross correlation; therefore, we must account for this choice when testing the significance of this statistic. That is, even in cases where the distributions of the cross correlations  $C_{ij}[\tau]$  are well approximated by the normal distribution, their extrema will not be normally distributed, and so  $p$ -values calculated using this distribution will be inaccurate.

To address both of these issues, we propose a more appropriate analytic method: the extremum method. We start by applying the well-known Fisher transformation [40] to each  $C_{ij}[\tau]$ , yielding

$$C_{ij}^F[\tau] = \frac{1}{2} \ln \frac{1 + C_{ij}[\tau]}{1 - C_{ij}[\tau]}, \quad (3)$$

which should more closely follow a normal distribution than the original  $C_{ij}[\tau]$ . Since this transformation is monotone and symmetric about zero, the lag  $\hat{l}_{ij}$  maximizing  $|C_{ij}[\tau]|$  will also maximize  $|C_{ij}^F[\tau]|$ . Let  $s_{ij}^F$  be the Fisher transformation of  $s_{ij}$ , which we propose to use instead of  $s_{ij}$ .

Next, we use results from extreme value theory to approximate the distribution of our new test statistic. We scale the values  $C_{ij}^F[\tau]$  over  $\tau$  by their empirical standard deviations  $(\widehat{\text{var}}(C_{ij}^F))^{1/2}$  so that the resulting scaled values should approximately follow a standard normal distribution. If there were no dependency within the time series  $x_i[t]$ , and instead we observed independent and identically distributed sequences at each node  $i$ , then the appropriate standard deviation is known to be  $(n-3)^{-1/2}$  [40]. But given the dependency in our time series data, we expect that the true standard deviations may differ from this value, and so we choose to estimate them empirically.

The scaled value  $z_{ij}^F = s_{ij}^F / (\widehat{\text{var}}(C_{ij}^F))^{1/2}$  can be expected to behave like the maximum of the absolute values of a sequence of standard normal random variables. Using established results for statistics of this form, we obtain therefore that

$$P[z] \approx \exp\{-2 \exp[-a_n(z - b_n)]\}, \quad (4)$$

where  $P[z] = \text{Pr}\{z_{ij}^F \leq z\}$ ,  $a_n = \sqrt{2 \ln n}$ , and  $b_n = a_n - (2a_n)^{-1}(\ln \ln n + \ln 4\pi)$ . A derivation of Eq. (4), which holds in the asymptotic sense of large  $n$ , is provided in Appendix A. For the case of  $n=201$ , as in all of our numerical results below,  $a_n=3.2568$  and  $b_n=2.6121$ . Using the approximation above, it is straightforward to calculate  $p$ -values for the rescaled test statistics  $z_{ij}^F$ .

Intuitively, the extremum method accounts for our choice of a maximum cross correlation. By virtue of the Fisher transformation, the values  $C_{ij}^F[\tau]$  will be approximately normally distributed. But because we have chosen  $s_{ij}^F$  as the *maximum* of the absolute value of the transformed correlations, we expect its value to be skewed toward the tail of the normal distribution. If we had chosen instead any other lag than that maximizing the cross correlation, then the corresponding value  $C_{ij}$  (and hence  $C_{ij}^F$ ) would be smaller. Therefore, our definition of  $s_{ij}^F$  produces  $p$ -values that, if computed from the normal distribution, are biased in the sense of being inappropriately small. From the perspective of our network inference task, this means that—for any given choice of

threshold—we will be more liberal in our assignment of edges than we should be. The distribution in Eq. (4) essentially corrects for this bias, by explicitly accounting for our use of the maximum.

## 2. Frequency domain bootstrap method

The previous method provides an analytic formula for testing the significance of  $s_{ij}$ . In utilizing this formula, we make specific asymptotic distributional assumptions about the test statistic—the maximal cross correlation. These assumptions are likely to be only approximate idealizations of the correlation results emerging from, for example, a complicated physical system like the human brain. A method to test the significance of  $s_{ij}$  that requires fewer assumptions is desirable. The second method we introduce—the frequency domain bootstrap—satisfies this desire, but is computationally expensive. We note that more sophisticated surrogate procedures exist, for example to preserve nonlinear properties of the original time series [41]. Determining the most appropriate surrogate procedure for a particular coupling measure (e.g., the cross correlation or the nonlinear correlation coefficient discussed below) is beyond the scope of this work.

As the name indicates, the method consists of applying the bootstrap principle (e.g., [42]), but in the spectral domain; methods of this sort were first proposed in [43]. We calculate our frequency domain bootstrap through the following steps. First, we compute the power spectrum (Hanning tapered) of each time series in the network. We then average these power spectra from all time series, and smooth the resulting average spectrum. We use this spectrum estimate ( $P[\omega]$ ) to compute the standardized and whitened residuals for each time series  $x_i[t]$ ,

$$e_i[t] = \mathcal{F}^{-1}(\tilde{x}_i[\omega] / \sqrt{P[\omega]}). \quad (5)$$

Here  $\tilde{x}_i[\omega] = \mathcal{F}(x_i[t])$  is the Fourier transform of the original time series  $x_i[t]$  and  $\mathcal{F}^{-1}(\star)$  is the inverse Fourier transform of  $\star$ . Finally, for each bootstrap replicate, we resample the values  $e_i[t]$  with replacement and compute the surrogate data,

$$\hat{x}_i[t] = \mathcal{F}^{-1}(\tilde{e}_i[\omega] \cdot \sqrt{P[\omega]}), \quad (6)$$

where  $\tilde{e}_i[\omega]$  is the Fourier transform of the residuals  $e_i[t]$  resampled with replacement. This last step ensures that the spectral characteristics of the original data are preserved in the surrogate data.

We compute  $N_s$  instances of these surrogate data, and for each instance we calculate the test statistic  $\hat{s}_{ij}$  for each pair of nodes  $i$  and  $j$ , (i.e., we calculate the maximum of the absolute value of the cross correlation between the surrogate data  $\hat{x}_i[t]$  and  $\hat{x}_j[t]$ ). The  $N_s$  instances of  $\hat{s}_{ij}$  form a bootstrap distribution of maximum cross correlation values to which we compare  $s_{ij}$  observed in the original data and assign a  $p$ -value.

Constructing the bootstrap distribution of  $\hat{s}_{ij}$  values for all node pairs is computationally expensive. If our network contains 100 nodes, then we would like to compute a bootstrap distribution (and test the significance) of each of the 100

$\times 99/2=4950$  values  $s_{ij}$ . If each bootstrap distribution requires  $N_s=10\,000$  surrogates, a standard choice in the literature, then we construct surrogate data and compute the cross correlation over  $10^7$  times. We reduce this expensive operation in the following way: instead of computing bootstrap distributions for all electrode pairs, we compute the bootstrap distribution (with  $N_s$  surrogates) for only a subset of node pairs. We then define the *merged distribution* as the combined distribution for the entire subset of node pairs. We use the merged distribution to test the significance of  $s_{ij}$  for all node pairs (even pairs not used in calculating the merged distribution). Note, however, that in doing so we assume that the null distribution of  $s_{ij}$  is the same for all node pairs.

### C. Step 3: Control of the false detection rate

To test the statistical significance of the values  $s_{ij}$ , we may apply either of the two methods described above (or even, practically speaking, the naive method as well). Networks of, say, 100 nodes will consist of  $100 \times 99/2=4950$   $s_{ij}$  values, each with an associated  $p$ -value. Clearly, multiple testing is an important concern. If we simply choose a standard  $p$ -value cutoff for assessing significance (such as  $p < 0.05$ ), then we expect the number of network edges incorrectly declared present to scale proportionally (i.e., roughly 250 such edges, for an 0.05 cutoff). To control for this abundance of false positives, we could define a stricter cutoff; for example, we could use the Bonferroni correction and divide the  $p$ -value threshold by the number of node pairs (i.e.,  $p < 0.05/4950=10^{-5}$ ). This conservative control of the family-wise error rate (the probability of making one or more false discoveries) is likely too strict for data in which we expect relatively few significant edges *a priori*, i.e., for sparse networks.

Instead, we employ the less conservative FDR to control for multiple testing. The FDR is defined as the expected proportion of erroneously rejected null hypotheses among the rejected ones [44,45], and various procedures exist for controlling the FDR in practice [26]. Generally speaking, the notion of FDR control guarantees that the expected proportion of falsely declared edges in our inferred networks is no more than a prespecified fraction  $q \in (0, 1)$ . However, in order for this guarantee to hold, two assumptions must be true, namely, that (i) statistical  $p$ -values associated with each test are computed accurately and (ii) tests are independent. Of these assumptions, the first is critical, while the second is less so. Dependency among tests is often relatively benign and, when it is not, can be dealt with using adjusted FDR procedures (e.g., [45]). However, inaccurate calculation of  $p$ -values is known to be disastrous to FDR procedures. Our analyses presented below confirm this in the context of our network inference problem, and the majority of our efforts focus around this point, as we described in Sec. III B.

Here we implement the linear step-up FDR controlling procedure of Benjamini and Hochberg [44], which is computed as follows. First, order the  $m=N(N-1)/2$   $p$ -values  $p_1 \leq p_2 \leq \dots \leq p_m$ . Then, choose a desired FDR level  $q$ . Finally, compare each  $p_i$  to the critical value  $q \cdot i/m$  and find the maximum  $i$  (call it  $k$ ) such that  $p_k \leq q \cdot k/m$  [and therefore

$p_{k+1} > q \cdot (k+1)/m$ ]. We reject the null hypothesis that time series  $x_i[t]$  and  $x_j[t]$  are uncoupled for  $p_1 \leq \dots \leq p_k$ .

The choice of  $q$  determines a threshold  $p$ -value, for which we declare all features significant up through that threshold [46]. The value  $q$  represents an upper bound on the expected proportion of false positives among all declared edges in our inferred network (i.e., among all node pairs for which  $s_{ij}$  was declared to be significant). For example, if we fix  $q=0.05$  and find 100 significant values of  $s_{ij}$ , then we expect  $0.05 \cdot 100=5$  false positives (i.e., five false edges in the 100 edge network).

## IV. RESULTS

We analyze three data sets using the procedure defined in Sec. III. Two data sets we create with specific (known) structural topologies, to which we compare the functional topologies extracted through analysis of the dynamic data. The third we observe from a human epileptic subject undergoing invasive electrical monitoring of the cortex.

### A. Pink noise data

Many time series data produced by natural systems possesses a  $1/f^\alpha$  power spectrum [47–49]. To mimic this behavior, we create a nine-node network, first by generating 500 ms (sampling interval 1 ms) of independent colored noise ( $\alpha=0.33$ ) data  $w_i[t]$  at each node  $i$ . We then connect node  $i$  to  $j$  by adding pointwise to  $w_j[t]$  the data  $w_i[t]$  scaled by a factor of 0.4. For example, we connect node No. 1 to No. 2 by adding to  $w_2[t]$  the time series  $0.4 \cdot w_1[t]$  for each time point  $t$  to create  $x_2[t]=w_2[t]+0.4 \cdot w_1[t]$ , the time series associated with node No. 2. In Fig. 2(a) we illustrate the topology of the constructed network (a total of nine directed connections exist) and the power spectrum averaged over the data from all of the nodes.

Having established the network topology, we now attempt to recover it directly from the time series data. To do so, we apply our coupling measure ( $s_{ij}$ , the maximum of the absolute value of the cross correlation) pairwise to all  $m=9 \cdot 8/2=36$  electrode pairs in the network. We then test the significance of each  $s_{ij}$  and compute a  $p$ -value using the analytic and computational procedures defined above. We begin with the naive method, whose  $p$ -values we plot as asterisks in the lower portion of Fig. 2(b). Plotted in increasing order, these  $p$ -values range from  $\sim 10^{-5}$  to 0.3. We fix  $q=0.10$  and also plot in the lower portion of Fig. 2(b) the line of slope  $q/m=0.10/36=0.0028$  and zero intercept. Following the linear step-up FDR procedure, we reject the null hypothesis of no coupling for those (nine) electrode pairs with  $p$ -values below this line. We plot in the upper portion of Fig. 2(b) the (nine) “significant edges” corresponding to the significant  $p$ -values. Our confidence in this nine-node network—derived from the time series data—is high; from the FDR procedure we expect  $0.10 \cdot 9 \sim 1$  false positive edge (i.e., one spurious edge between uncoupled nodes). In this case, we find exact agreement between the known network topology [Fig. 2(a)] and the derived topology. We note that, for sake of clarity, we chose a simple coupling measure that

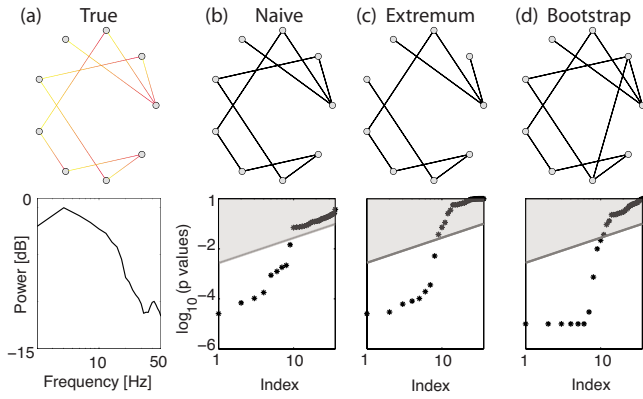


FIG. 2. (Color online) For the colored noise network all three measures perform equally well and detect the underlying structural topology. In the upper row, each circle indicates a node. [(a), upper] The true connectivity of the network. With colored (or shaded) lines we indicate directed connections between nodes; connections initiate from the red (dark) line end and terminate at the yellow (light) line end. [(a), lower] The power spectrum averaged over the data from all nine nodes follows an approximate  $1/f^\alpha$  behavior. [(b)–(d)] The  $p$ -values (lower) and corresponding functional network topologies (upper) derived from (b) the naive method, (c) extremum method, and (d) bootstrap method. The dark gray line in the lower figures indicates the threshold for the linear step-up FDR procedure; we consider  $p$ -values below this line—in the unshaded region—significant. All three significance tests capture the functional topology equally well.

does not determine edge direction. More sophisticated coupling measures that indicate edge direction may be employed following the general paradigm outlined above, as we discuss in Sec. V.

In Figs. 2(c) and 2(d), we show the topology derived using the extremum and bootstrap methods, respectively. In both cases, we follow the linear step-up FDR procedure with  $q=0.10$  to identify significant edges. For the extremum method [Fig. 2(c)] we identify eight significant edges, one less than expected and anticipate  $0.10 \cdot 8 \sim 1$  false positive edge.

To compute the frequency domain bootstrap, we first calculate the average power spectrum of all (nine) nodes. We then create a merged distribution using a subset of ten electrode pairs (of the possible 36) and  $N_s=10\,000$  for each surrogate distribution. The resulting merged distribution contains  $10 \cdot N_s=10^5$  points; therefore, the smallest  $p$ -value we can compute through this method is  $10^{-5}$ . We find, in this case, six  $p$ -values at this detection limit. Using the bootstrap method [Fig. 2(d)], we identify ten network edges, one more than expected. We do expect  $0.10 \cdot 10=1$  false positive edge in the network although given only the time series data we could not identify which of the ten edges is spurious.

These simulation results suggest that all three measures of edge significance perform equally well. This is surprising, especially for the naive method in which we neither Fisher transform the maximal correlation values (to induce normality), nor account for our choice of an extremum (the maximum of the absolute value of the cross correlation). The naive method succeeds, in this case, because the two omissions appear to balance. Omitting the Fisher transformation

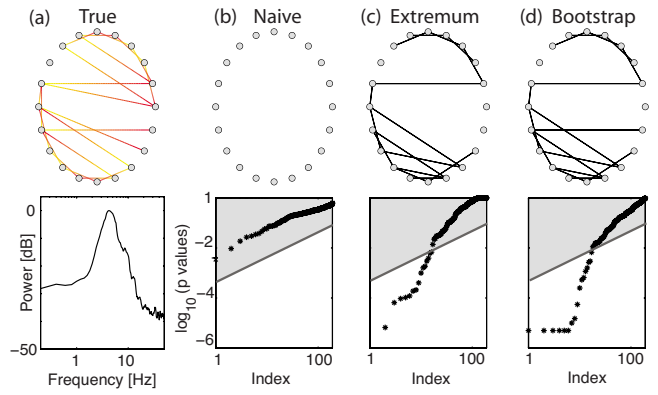


FIG. 3. (Color online) For the simulated neural data the choice of statistical test is vital to construct an appropriate network. [(a), upper] The data consist of twenty cell groups (gray circles) and 22 connections between cell groups. Directed connections proceed from the red (dark) to yellow (light) end of each line. [(a), lower] The average power spectrum of the 22 nodes exhibits a prominent peak near 5 Hz. [(b)–(d)] The functional networks deduced. The naive method (b) identifies no significant  $p$ -values; with  $q=0.10$  in the linear step-up FDR procedure, none of the  $p$ -values lie below the (dark gray) line  $(q/m) \cdot i$ . The extremum method (c) identifies 17 significant edges (of which we expect two are false positives); 14 match the structural network in (a). The bootstrap method (d) detects 18 edges, of which we expect two false positives. This procedure detects 15 (of the 22) true edges.

increases the  $p$ -values we observe, while utilizing the normal distribution with zero mean—not the extremum distribution—decreases the  $p$ -values. One omission compensates the other so that, in this case, the resulting  $p$ -values are approximately correct. Unfortunately, we cannot rely on this delicate balance to always succeed as we illustrate in the next example.

### B. Simulated neural data

In the previous model, we simulated colored noise activity possessing a  $1/f^\alpha$  falloff of the power spectrum. We now consider a more realistic model of interacting neural populations. We provide a brief description of the model here; more details may be found in Appendix B. The model consists of 1000 neurons divided into twenty groups of 50 cells. Within each group we include strong connections (excitatory synapses) between randomly chosen neurons; activity initiated by a few neurons in a group quickly spreads to the other neurons of the same group. Between cell groups, we establish only weak (excitatory synaptic) connections joining individual neurons of specific cell groups. We illustrate the topology of these weak connections between cell groups in Fig. 3(a). In this figure, each gray circle represents a cell group (of 50 strongly connected neurons) and lines represent connections between cell groups. With this connectivity in place, we simulate the neural dynamics and compute the average population activity of each group. The local populations tend to oscillate at 5 Hz; in Fig. 3(a) we plot the power spectrum averaged over all groups. We then employ the general paradigm described above to the resulting neural activi-

ties and compare the measured functional connectivity (i.e., the pattern of connectivity inferred from the simulated neural dynamics) to the known structural connectivity between nodes shown in Fig. 3(a). The results, as we show below, depend upon the significance test we employ.

We apply the coupling measure pairwise to all  $m = 20 \cdot 19 / 2 = 190$  possible group pairs in the network and test the significance of each result by computing a  $p$ -value using one of the three procedures defined above. We begin with the naive method, whose  $p$ -values we plot as asterisks in the lower portion of Fig. 3(b). With  $q=0.10$  in the linear step-up FDR procedure, we find no significant values of maximal cross correlation; none of the  $p$ -values lie below the line  $(q/m) \cdot i$ . The resulting (trivial) network—shown in the upper portion of Fig. 3(b)—contains no edges. The other two significance tests produce nontrivial networks. Using the extremum method and linear step-up FDR procedure (with  $q=0.10$ ) we identify 17 significant edges. The resulting network, shown in Fig. 3(c), correctly identifies 14 edges and possesses three erroneous edges (i.e., edges we identify in the functional network that do not exist in the structural network). We expect from the FDR procedure  $q \cdot 17 \sim 2$  false positives, in approximate agreement with the three erroneous edges observed. Finally, we show in Fig. 3(d) the  $p$ -values and network determined using the bootstrap method. In this case, we detect 18 edges (and expect two false positives). This procedure detects 15 (of the 22) true structural edges and produces three erroneous edges, again in approximate agreement with the number of false positives expected.

In all three cases, the functional topology derived from the mean dynamics fails to capture exactly the true structural topology of the network. The naive method detects no significant edges and performs most poorly. This is not surprising; we expect that the un-normalized  $p$ -values and incorrect distribution of maximal correlation values will compromise the naive method. The extremum and bootstrap methods produce similar functional networks that approximate the true structural network. Although both measures make mistakes, the FDR procedure provides an estimate for the number of erroneous edges to expect. We conclude that for these simulated data, the extremum and bootstrap methods outperform the naive method and qualitatively reproduce many (but not all) of the network edges.

### C. Human ECoG data

In the previous two examples, we applied the coupling analysis to networks with known structural topology. This allowed us to compare the derived functional topology with the true structural topology and determine each method's performance. As a last illustration of the methods, we consider voltage activity recorded directly from the cortical surface (electrocorticogram or ECoG data) of an epileptic human subject for clinical purposes (Appendix C). We focus on a short interval (1 s) of data recorded from 97 electrodes while the subject experienced a seizure. The data within the 1 s interval are approximately weakly stationary (as a rough indication we find that the mean, standard deviation, and power spectra [Fig. 4(a)] are not significantly different in the

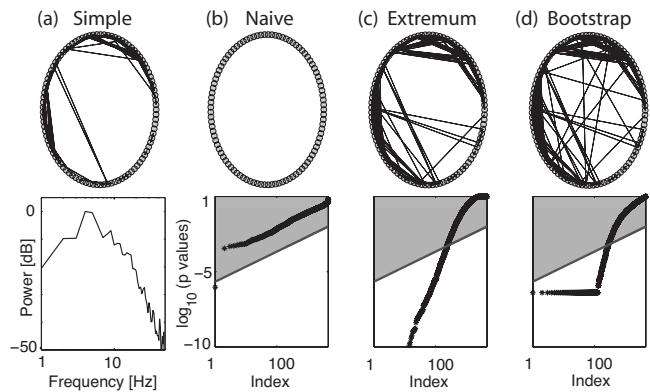


FIG. 4. Functional networks constructed from 1 s of ECoG data recorded at 97 electrodes during a seizure depend upon the statistical test we perform. [(a), upper] A simple threshold network with edges (black lines) drawn between nodes pairs exhibiting sufficient functional coupling,  $s_{ij} > 0.75$ . [(a), lower] The power spectrum averaged over all 97 electrodes exhibits a broad peak from 5–10 Hz and an approximate  $1/f^\alpha$  falloff from 10–50 Hz. [(b)–(d), lower] The 4656  $p$ -values calculated from (b) the naive method, (c) extremum method, and (d) bootstrap method. For each method, we fix  $q=0.01$  in the FDR procedure. [(b)–(d), upper] The corresponding functional networks. The naive method (b) detects only one significant edge and the corresponding network is nearly trivial. The network created from the extremum method (c) contains 238 edges, and from the bootstrap method (d) 279 edges.

first and second halves of the interval). We apply all three methods to the data and compare the resulting (functional) networks. In this case, the structural connectivity is unknown. We find that, as before, the extremum and bootstrap methods produce networks consistent with one another.

We show the deduced functional networks in Figs. 4(b)–4(d). In each case, we test the significance of  $m = 97 \cdot 96 / 2 = 4656$  maximal cross correlation values, and use a linear step-up FDR procedure with  $q=0.01$  to define significant  $p$ -values. For the naive method [Fig. 4(b)] we find one significant  $p$ -value and the corresponding (nearly trivial) network contains only one edge. We note that the node locations in Fig. 4 do not correspond to their physical locations on the human cortex. Instead, we simply arrange the nodes in a circle.

From the extremum and bootstrap methods we create similar networks. For the former, we identify 238 significant edges (of which we expect three false positives) as drawn in Fig. 4(c). For the latter, we select 500 electrode pairs (of the possible 4656 pairs) to compute surrogate distributions, each distribution containing  $N_s=5000$  realizations. The smallest  $p$ -value detectable in the resulting merged distribution is  $4 \times 10^{-7}$ . Using this method we find the 279 significant edges drawn in Fig. 4(d), of which we expect three false positives.

Comparing the functional networks deduced from the extremum and bootstrap methods, we find that the two are similar. Moreover, we show in Fig. 4(a) a fourth functional network constructed using a simple threshold procedure; we include edges only between those node pairs with  $s_{ij} > 0.75$ . Remarkably, all three networks are qualitatively similar (notice the concentration of edges near the left and upper nodes) although we use different techniques to construct each net-

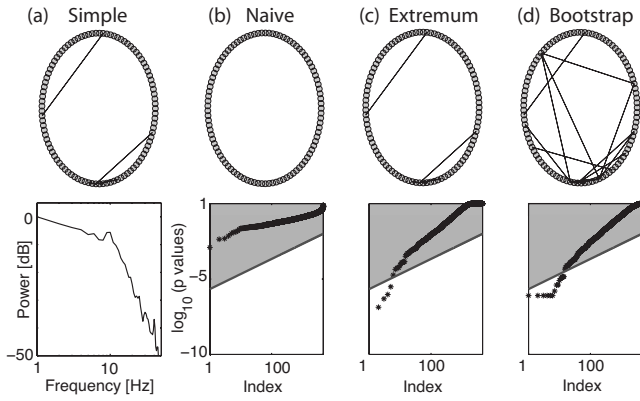


FIG. 5. By shuffling the ECoG data, we reduce coupling between the time series and detect no (or few) edges. [(a), upper] A simple threshold network with edges drawn between node pairs with sufficient functional coupling ( $s_{ij} > 0.75$ ). [(a), lower] The power spectrum average over all electrodes. [(b)–(c)] The  $p$ -values (lower) and corresponding networks (upper) derived from the (b) naive, (c) extremum, and (d) bootstrap methods. Only the latter two methods detect edges, five and 14 respectively, of we expect one false positive.

work. Of course the simple threshold network does not indicate our confidence in the network: how many edges in Fig. 4(a) are false positives? In addition, we note that the bootstrap method is computationally expensive; constructing the surrogate distribution requires approximately 90 min on a 2 GHz Core Duo processor and therefore at least 45 h to construct the networks for 30 s of seizure activity. The extremum method, designed for our particular choice of coupling measure, identifies a network similar to the bootstrap method in a computationally efficient way.

#### D. Human ECoG data: Shuffled

For the human ECoG data, we do not know the structural network (i.e., we do not know the topology of chemical and electrical connections between neurons in these cortical regions). Therefore, we cannot validate the functional networks shown in Fig. 4 by comparison with anatomical connections. However, we can manipulate the ECoG data to disrupt functional connections and verify that our significance tests detect no coupling. To do so, we create a new data set: we assign to each electrode 1 s of data chosen at random from a 120 s interval that includes 60 s of pre-seizure and 60 s of seizure activity. For example, electrode No. 1 may contain ECoG data from  $t=[8.2, 9.2]$ , electrode No. 2 data from  $t=[97.0, 98.0]$ , electrode No. 3 from  $t=[110.4, 111.4]$ , and so on. With the data chosen in this way, we expect only weak associations between electrode pairs.

We follow the procedure described above to analyze these “shuffled” data. We compute the maximal cross correlation for each electrode pair, and show the corresponding  $p$ -values and functional networks in Figs. 5(b) and 5(c). With  $q = 0.01$ , we find no significant coupling using the naive method. We do detect five significant edges with the extremum method and 14 with the bootstrap method (and expect one false positive in each case). Some of these significant

edges match those determined using a simple threshold procedure ( $s_{ij} > 0.75$ ) whose network we show in Fig. 5(a). We conclude that the three significance tests behave as expected for the shuffled data; if we disrupt the coupling in the data, we expect nearly trivial functional networks.

#### E. Human ECoG data: Nonlinear correlation

In this work we described a general statistical paradigm for establishing confident network topologies. We chose to use a simple linear measures of association (the cross correlation) and found in the simulated data that this measure captures much of the association structure within the networks. For some types of data, further coupling structure could be discovered with more sophisticated higher-order association measures. The general paradigm we propose applies to any choice of coupling measure, including measures that detect nonlinear interactions.

The observational data of interest here—the ECoG recorded during a seizure—arise from a complicated dynamical system of nonlinear elements (i.e., populations of coupled neurons). Therefore, we might expect nonlinear measures of coupling to be more appropriate for the analysis of these data. Yet, linear and nonlinear measures of coupling appear to perform equally well when applied to simulated and ECoG data [50–55]. One reason for this may be that, although both linear and nonlinear couplings occur between ECoG data, the linear effects dominate. To explore the appropriateness of nonlinear methods here, we plot in Fig. 6 two association matrices generated for the 1 s of ECoG data analyzed in Sec. IV C. In Fig. 6(a) we plot  $s_{ij}$ —the linear correlation measure—and in Fig. 6(b) we compute the nonlinear correlation coefficient ( $h_{ij}^2$ ) to construct the association matrix [56,57]. Briefly, we compute  $h_{ij}^2$  for each electrode pair over delays of  $\pm 200$  ms and select the maximum of  $h_{ij}^2$  to plot in Fig. 6(b). Visual inspection suggests that both measures produce qualitatively similar association matrices. To quantify this, we compute the two-dimensional correlation between the association matrices in Figs. 6(a) and 6(b) and find a value of 0.83. Repeating this calculation for sixty separate 1 s intervals extracted from the ECoG data, we find correlation values of  $0.78 \pm 0.07$  between the association matrices.

To construct a confident network from the association matrix in Fig. 6(b), we must test the significance of each nonlinear correlation coefficient  $h_{ij}^2$ . We do not possess an analytic expression for the significance of  $h_{ij}^2$ , but can utilize the same frequency domain bootstrap paradigm described in Sec. III. Briefly, we first make the association matrix symmetric by defining  $h_{ij}^2 = h_{ji}^2 \equiv \max(h_{ij}^2, h_{ji}^2)$  so that the resulting network is undirected. Then we compare each  $h_{ij}^2$  to a bootstrap distribution of  $10^6$  surrogates calculated following the procedure in Sec. III. We fix  $q=0.01$  and plot the 231 edges identified as significant in Fig. 6(d). For comparison, we plot in Fig. 6(c) the network constructed from the cross correlation using the bootstrap procedure. The resulting networks are qualitatively similar. In agreement with earlier reports, it is evident that the linear and nonlinear coupling measures produce similar association matrices and networks [52].



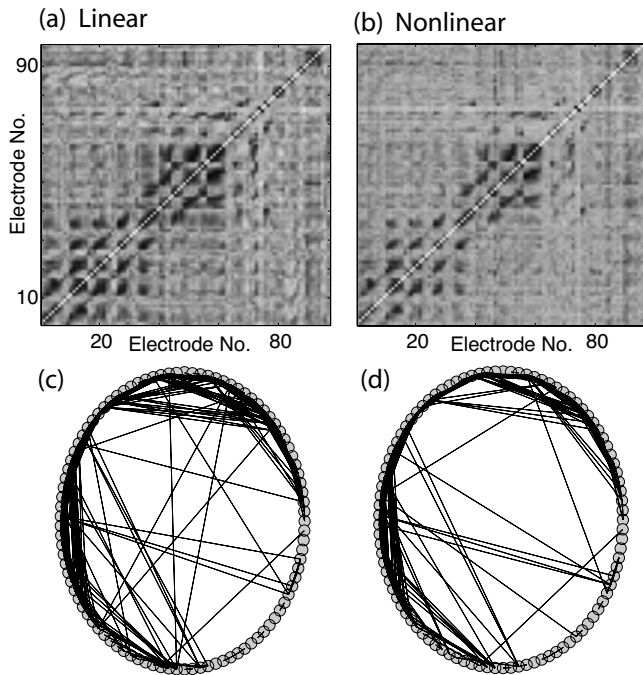


FIG. 6. For the ECoG data, linear and nonlinear association measures produce similar association matrices and networks. In (a) we plot  $s_{ij}$  for each electrode pair, while in (b) we plot the maximum of the nonlinear correlation coefficient for each electrode pair. Both measures range between 0 (weak coupling in white) and 1 (strong coupling in black). Visual inspection suggests the two measures produce comparable association matrices. [(c)–(d)] The networks constructed from the (c) linear and (d) nonlinear association matrices. We employ the bootstrap method and fix  $q=0.01$  in the FDR procedure. The resulting networks are quite similar. We note that (c) is identical to the network in Fig. 4(d).

The general paradigm we describe in Sec. II applies for any choice of coupling measure, linear or nonlinear. In this work, we chose the cross correlation—a linear association measure—to construct the functional networks. There are also two important practical reasons to employ this simple measure. First, computing the cross correlation requires only one user-specified parameter, the lag  $\tau$ . Nonlinear measures require, for example, the choice of frequency band for phase synchronization or the choice of embedding dimension and delay time for generalized synchronization. Second, for the cross correlation measure we can derive analytic expressions for the significance test, making this technique computationally efficient compared to surrogate procedures. This is particularly important for studying dynamically evolving networks in which we construct many networks each with associated levels of confidence. Although a linear association measure suffices here, subtle changes in coupling of ECoG data (perhaps resulting from behavioral changes less dramatic than seizure) may require more sophisticated nonlinear measures.

## V. DISCUSSION

Our increased ability to collect multivariate spatiotemporal data (e.g., from high-density electrode arrays) necessitates

the construction and analysis of complex functional networks. In this manuscript, we adopted a statistical hypothesis testing paradigm for constructing such functional networks. This paradigm involved three steps: (1) choice of an association measure, (2) definition of a significance test, and (3) accounting for multiple significance tests. Although the paradigm itself is quite general, the details accompanying each step are problem specific.

Here we developed this general paradigm for multivariate time series data. For the association measure we chose the maximum of the absolute value of the cross correlation. We defined two approaches to significance testing (one analytic and the other computational), and employed a linear step-up FDR procedure to account for multiple tests. Applying these techniques to three data sets, we showed that the choice of significance test was critical. Without accurate  $p$ -values for each network edge, we lack confidence in the resulting network.

The general paradigm outlined in Sec. II applies to any choice of association measure. In this work we focused on a simple cross correlation measure for four reasons. First, the measure is computationally efficient. Second, the cross correlation requires only one user-specified parameter (the lag  $\tau$ ). Third, analytic expressions exist (or can be derived) to test the significance of each cross correlation result. Fourth, linear coupling measures appear to perform just as well as nonlinear coupling measures applied to macroscopic voltage data [50–55]. More appropriate coupling measures exist [29] that may perhaps improve the network results we present here. In particular, we do not distinguish between direct and indirect interactions with the cross correlation measure (e.g., if region  $A$  drives regions  $B$  and  $C$ , then we will detect correlations between  $B$  and  $C$ , although these two regions may possess no direct interactions). Measures that distinguish direct from indirect interactions and incorporate the flow of information [58,59] would be of use. However, choosing a more sophisticated association measure does not guarantee more accurate functional networks. The coupling measure must also include an accurate significance test; without precise  $p$ -values for each network edge, we weaken our measures of network confidence.

Researchers in various other contexts have followed a similar strategy of associating  $p$ -values with each network edge and accounting for multiple significance tests (e.g., [60–62]). Our numerical results illustrate how the choice of an appropriate significance test associated with a specific coupling measure is critical. That a measure possesses a significance test does not guarantee accurate  $p$ -values; typically significance tests make specific assumptions about the data. For example, we found that the naive method—although perhaps intuitively appealing—was inappropriate because we did not account for taking the *maximum* of the absolute value of the cross correlation, and thus produced inaccurate  $p$ -values and inaccurate networks. Therefore, we utilized two additional complimentary measures. By testing the paradigm on simulated data with known physical connectivity we deduced appropriate significance tests for the association measure implemented here.

We note that trivial networks (e.g., networks without edges as in Fig. 5) rarely appear in practice. Upon finding a

trivial (or nearly trivial) network, a common response is to adjust the network threshold to include more edges, perhaps until the network becomes strongly connected. To follow a similar strategy here we increase the value of  $q$  in the linear step-up FDR procedure. If, for example, we set  $q=0.5$  (instead of  $q=0.01$ ) we may detect new significant network edges. But by increasing  $q$  we decrease our confidence in the network; with  $q=0.5$ , we expect half of the network edges declared significant are false positives. Thus, through our choice of  $q$ , we balance the number of edges detected with our confidence in the network.

The typical approach to construct functional networks from multivariate time series data involves thresholding an association measure. For example, we may define edges between nodes whose maximal cross correlation exceeds 0.75, as in [23]. This procedure for constructing a network suffers from numerous inadequacies. First, we lack a measure of confidence in the resulting network. With this choice of 0.75 as threshold how many spurious edges do we expect, and does this number change as we vary the threshold? Second, we expect the choice of threshold may depend on the particular instance of data observed. For example, in constructing functional networks of ECoG data recorded during a seizure, the threshold may vary from patient to patient, depending on mechanisms intrinsic to each individual. Finally, a more robust approach to constructing functional networks must propagate error in the association measure to uncertainty in network measures (e.g., to uncertainty in measures of degree or betweenness).

We propose that choosing a threshold value of  $q$ , rather than a threshold value of an association measure, constitutes a more rigorous procedure for establishing functional networks. By choosing the threshold through the use of formal statistical hypothesis tests, we create functional networks with specified levels of network uncertainty that may be calibrated across a population of multivariate data. In the future, we will use this approach to study how uncertainty in the association measure affects uncertainty in network characteristics, how to adopt these procedures for weighted (rather than binary) networks, and how networks derived from non-stationary ECoG data evolve in time. Combined with biophysical models, robust techniques to create functional networks will perhaps illuminate the mechanisms that produce the observed activity and, when necessary, suggest how to alter this activity.

**ACKNOWLEDGMENTS**

M.A.K. was supported by a Career Award at the Scientific Interface from the Burroughs Wellcome Fund. U.T.E. was supported by the NSF under Grant No. NSF-IIS-0643995. S.S.C. was supported by the Rappaport Foundation and the Executive Committee of Research of Partners Healthcare. E.D.K. was supported by the ONR through Award No. N00014-06-1-0096.

**APPENDIX A: DERIVATION OF EQ. (4)**

Suppose that  $Z_1, \dots, Z_n$  are independent and identically distributed normal random variables, with mean 0 and vari-

ance 1. Define  $M_n = \max_i(Z_i)$  and  $m_n = \min_i(Z_i)$ . Then

$$\begin{aligned} \Pr(\max_i |Z_i| \leq z) &= \Pr(M_n \leq z, \quad m_n \geq -z) \\ &\approx \Pr(M_n \leq z)\Pr(m_n \geq -z) \\ &= [\Pr(M_n \leq z)]^2. \end{aligned}$$

The approximate equality in the second line follows by asymptotic independence of the max and min (e.g., Theorem 1.8.2 of [63]), while the equality in the last line follows by symmetry of the normal distribution. Now by, for example, Theorem 1.5.3 of [63] we have

$$\Pr(a_n(M_n - b_n) \leq z) \approx \exp(-e^{-z}), \quad (A1)$$

with equality holding asymptotically in  $n$ . As a result, we obtain the expression in Eq. (4) in the ideal case that the standardized statistic  $z_{ij}^F$  derives from cross correlations  $C_{ij}[\tau]$  that are independent. Although these cross correlations will clearly be dependent, approximation (4) nevertheless can be expected to hold fairly generally, as the basic limiting extreme value distribution used here is quite robust for sequences of normal random variables under a range of dependency conditions, for both stationary and even nonstationary cases. See, for example, chapters 4 and 6 of [63].

**APPENDIX B: NEURAL MODEL**

We model the dynamics of each neuron with two ordinary differential equations, one to represent the membrane voltage, and the other a membrane recovery variable [64]. We choose the model parameters so that each neuron generates regular spiking activity (i.e.,  $a=0.02$ ,  $b=0.2$ ,  $c=-65.0$ , and  $d=8.0$  in [64]). We then connect the neurons with excitatory synaptic connections to establish two connectivity patterns: strong-local connectivity and weak-distant connectivity. In both cases, we divide the neurons into groups of 50 cells numbered sequentially (i.e., group 1 contains cells {1–50}, group 2 cells {51–100}, group 3 cells {101–150}, and so on.) Within each local group of cells, we create 1200 directed excitatory synapses (of the possible  $50 \times 49 = 2450$  directed pairs with no self synapses). Each synapse is assigned a uniform random conduction delay between 0 and 10 ms and synaptic strength chosen uniformly between 0 and 15. These synapses establish the strong-local connectivity within a cell group and define the twenty cell groups in the network; see Fig. 7.

We also create weaker synaptic connections between the local cell populations. To do so, we select two groups (e.g., 1 and 8) and create 550 excitatory synapses from neurons in one group (e.g., 1) to neurons in another (e.g., 8). These “distant” synapses are weaker than the local connections; we assign the synaptic strengths smaller random values (chosen uniformly between 0 and 5) and uniform random conduction delay between 0 and 10 ms. We illustrate the 22 distant connections between the twenty cell groups in Fig. 3(a). Each gray circle represents a local cell group (i.e., a subset of 50 neurons). The colored (shaded) lines represent the distant synaptic connections between groups. In addition to the local and distant synaptic inputs, we also include strong synaptic




FIG. 7. (Color online) A cartoon representation of the network used to simulate the neural data. Each group of neurons contains 50 cells, represented as filled triangles. Within a group, we include many strong synaptic connections (terminating at large circles, orange). Between groups, we include few weak synaptic connections (terminating at small circles, blue). We show examples of the average activity of each group (from which we construct the functional networks) in the lower portion of the figure.

input (from the thalamus, say) to one randomly chosen neuron each millisecond, causing this neuron to generate an action potential. We follow the algorithm in [64] to simulate the neural population for 5000 steps (or 5 s) with a sampling interval of 1 ms. The model is similar to recent simulations [65,66], except that we introduce here connectivity with a particular structural topology.

In the human ECoG recordings described in Sec. IV C, we observe the dynamics of postsynaptic potentials produced by large neural populations, not the spiking activity of individual neurons [67,68]. To mimic these population dynamics, we construct the mean activity of the local neural groups in the following way. First, we define  $I[t]$  as the total current input to each neuron at time  $t$ . In the model equations we simulate here, these current inputs change instantaneously [64]. In reality, current inputs follow the opening and closing of channels and evolve more slowly [69]. To approximate these slow dynamics at postsynaptic neuron  $j$ , we use the following equation:

$$\dot{s}_j = I_j[t](1 - s_j) - s_j/20, \quad (\text{B1})$$

where  $s_j$  represents the state of the synapse at neuron  $j$ , and  $I_j[t]$  represents the total excitatory input current to the neuron  $j$  at time  $t$ . We note that, for simplicity, we approximate the total synaptic input to neuron  $j$  as a single synapse with dynamics driven by  $I_j[t]$ , the activity of all neurons presynaptic to  $j$ . When  $I_j[t]$  is large, excitatory current enters neuron  $j$  and  $s_j \rightarrow 1$ . When  $I_j[t]$  is small,  $s_j \rightarrow 0$  with a decay time constant of 20 ms, and neuron  $j$  approaches its resting potential. We define the mean activity of a neural group as the average of  $s_j$  over the local (50) cell population. For example, we compute the mean activity of population 1 as the average value of  $s_j$  for neurons  $j = \{1, 2, \dots, 50\}$ . We only use  $s_j$  to define the mean population activity; these synaptic dynamics do not impact the voltage dynamics of the model neurons.

With the population activity defined in this way (i.e., as the mean synaptic dynamics within a neural group), we simulate 5 s of neural dynamics and record the average activity of each group. We then scale the group activity to have zero mean and unit variance and add Gaussian noise (zero mean and 0.55 variance) to each sample of each time series. Finally, we downsample the group traces by a factor of five,




FIG. 8. (Color online) Examples of the ECoG data recorded from the human subject during seizure. In the upper trace we show 160 s of data recorded from a single electrode; the seizure begins at the transition from low amplitude to high-amplitude fluctuations denoted by the arrow. We indicate the one-second interval analyzed with the vertical red line. The lower ten traces illustrate the voltage activity recorded from ten (of the 97) electrodes during this one second interval; we apply the coupling measure to all pairs of these data.

reducing the sampling rate (from 1000 to 200 Hz) to decrease subsequent computational time; we show examples of the resulting time series data in Fig. 7.

### APPENDIX C: HUMAN SUBJECT DATA

The ECoG data were recorded from a 37 year old male patient with medically refractory epilepsy whose seizures began at age 3. Following the failure of seven antiseizure medications and a vagal nerve stimulator, and upon recommendation of his clinical team consisting of epileptologists and neurosurgeons, the decision was made to pursue resection of the tissue from which the seizures arose. To this end, subdural grids of electrodes were implanted. The goal of this procedure was to identify the epileptogenic zone—the region of the brain producing recurrent seizures—and surgically remove it [70].

The ECoG recordings consisted of 100 electrodes placed directly on the cortical surface (92 electrodes over the left frontal and temporal lobes) and within deep brain regions (eight electrodes within the temporal lobe). Following electrode implantation, the subject was admitted to a specialized monitoring unit and data recorded continuously at 500 Hz for ten days. During this time, four seizures were observed; to illustrate the measures, we analyze only the second seizure here. Analysis of the collected data was approved through the Partners Health Care Human Research Committee and the Charles River Campus Institutional Review Board.

We apply our coupling analysis to simultaneous recordings from 97 electrodes; three electrodes—suffering from extreme artifacts—were discarded. Before beginning the coupling analysis, we process the ECoG data in the following way. First, we low-pass filter the data (two-way least-squares finite impulse response filter) below 55 Hz to isolate the low-frequency components. We therefore ignore higher-

frequency activity that may delineate seizure onset [71,72] and instead focus on the high-amplitude low-frequency oscillations that characterize unequivocal clinical seizures [73]. We then choose a 1 s interval of the ECoG data during the seizure. We chose this short interval to balance two competing needs: stationarity and sufficient data. For the former, we must choose an interval in which the voltage dynamics at each electrode remain approximately consistent (i.e., exhibit

oscillations of the same approximate character). For the latter, we must choose an interval that contains enough data to calculate the coupling measure (e.g., a 50 ms interval would fail to capture some slow oscillations characteristic of a seizure). Finally, we normalize the data from each electrode within the 1 s interval to have zero mean and unit variance. We show examples of the ECoG data employed in the analysis in Fig. 8.

- 
- [1] S. Wasserman and K. Faust, *Social Network Analysis: Methods and Applications* (Cambridge University Press, Cambridge, 1994).
- [2] M. Newman, *SIAM Rev.* **45**, 167 (2003).
- [3] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D. U. Hwang, *Phys. Rep.* **424**, 175 (2006).
- [4] L. Hufnagel, D. Brockmann, and T. Geisel, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 15124 (2004).
- [5] C. J. Honey and O. Sporns, *Hum. Brain Mapp.* **29**, 802 (2008).
- [6] E. D. Kolaczyk, *Statistical Analysis of Network Data: Methods and Models* (Springer, New York, 2009), Chaps. 3 and 7.
- [7] L. A. N. Amaral, A. Scala, M. Barthelemy, and H. E. Stanley, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 11149 (2000).
- [8] D. J. Watts and S. H. Strogatz, *Nature (London)* **393**, 440 (1998).
- [9] M. E. J. Newman, S. H. Strogatz, and D. J. Watts, *Phys. Rev. E* **64**, 026118 (2001).
- [10] J. White, E. Southgate, J. Thomson, and S. Brenner, *Philos. Trans. R. Soc. London, Ser. B* **314**, 1 (1986).
- [11] E. Marder and D. Bucher, *Annu. Rev. Physiol.* **69**, 291 (2007).
- [12] D. Yu, M. Righero, and L. Kocarev, *Phys. Rev. Lett.* **97**, 188701 (2006).
- [13] M. Timme, *Phys. Rev. Lett.* **98**, 224101 (2007).
- [14] D. Napolitano and T. D. Sauer, *Phys. Rev. E* **77**, 026103 (2008).
- [15] A. Gozolchiani, K. Yamasaki, O. Gazit, and S. Havlin, *EPL* **83**, 28005 (2008).
- [16] K. Yamasaki, A. Gozolchiani, and S. Havlin, *Phys. Rev. Lett.* **100**, 228501 (2008).
- [17] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. Honey, V. Wedeen, and O. Sporns, *PLoS Biol.* **6**, e159 (2008).
- [18] C. J. Stam, *Neurosci. Lett.* **355**, 25 (2004).
- [19] S. Micheloyannis, E. Pachou, C. J. Stam, M. Vourkas, S. Erimaki, and V. Tsirka, *Neurosci. Lett.* **402**, 273 (2006).
- [20] S. C. Ponten, F. Bartolomei, and C. J. Stam, *Clin. Neurophysiol.* **118**, 918 (2007).
- [21] K. Srinivas, R. Jain, S. Saurav, and S. Sikdar, *Eur. J. Neurosci.* **25**, 3276 (2007).
- [22] C. J. Stam, B. F. Jones, G. Nolte, M. Breakspear, and P. Scheltens, *Cereb. Cortex* **17**, 92 (2006).
- [23] M. A. Kramer, E. D. Kolaczyk, and H. E. Kirsch, *Epilepsy Res.* **79**, 173 (2008).
- [24] K. Supekar, V. Menon, D. Rubin, M. Musen, and M. Greicius, *PLOS Comput. Biol.* **4**, e1000100 (2008).
- [25] E. Bullmore and O. Sporns, *Nat. Rev. Neurosci.* **10**, 186 (2009).
- [26] S. Dudoit and M. J. van der Laan, *Multiple Testing Procedures with Applications to Genomics* (Springer, New York, 2008).
- [27] J. J. Faith, B. Hayete, J. T. Thaden, I. Mogno, J. Wierzbowski, G. Cottarel, S. Kasif, J. J. Collins, and T. S. Gardner, *PLoS Biol.* **5**, e8 (2007).
- [28] N. Nariyai, E. D. Kolaczyk, and S. Kasif, *PLoS ONE* **2**, e337 (2007).
- [29] E. Pereda, R. Q. Quiroga, and J. Bhattacharya, *Prog. Neurobiol.* **77**, 1 (2005).
- [30] M. A. Braizer, in *Epilepsy, its Phenomenon in Man*, edited by M. A. Braizer (Academic, New York, 1973).
- [31] V. Menon, W. Freeman, B. Cutillo, J. Desmond, M. Ward, S. Bressler, K. Laxer, N. Barbaro, and A. Gevins, *Electroencephalogr. Clin. Neurophysiol.* **98**, 89 (1996).
- [32] B. R. Tharp and W. Gersch, *Comput. Biomed. Res.* **8**, 503 (1975).
- [33] P. Nunez, R. Srinivasan, A. Westdorp, R. Wijesinghe, D. Tucker, R. Silberstein, and P. Cadusch, *Electroencephalogr. Clin. Neurophysiol.* **103**, 499 (1997).
- [34] V. L. Towle, R. K. Carder, L. Khorasani, and D. Lindberg, *J. Clin. Neurophysiol.* **16**, 528 (1999).
- [35] J.-P. Lachaux, A. Lutz, D. Rudrauf, D. Cosmelli, M. L. V. Quyen, J. Martinerie, and F. Varela, *Clin. Neurophysiol.* **32**, 157 (2002).
- [36] M. Kamiński, M. Ding, W. A. Truccolo, and S. L. Bressler, *Biol. Cybern.* **85**, 145 (2001).
- [37] M. S. Bartlett, *J. R. Stat. Soc. Ser. B (Methodol.)* **8**, 24 (1946).
- [38] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control* (Holden-Day, San Francisco, 1970).
- [39] T. I. Netoff and S. J. Schiff, *J. Neurosci.* **22**, 7297 (2002).
- [40] R. A. Fisher, *Biometrika* **10**, 507 (1915).
- [41] M. Thiel, M. Romano, J. Kurths, M. Rolf, and R. Kliegl, *Europhys. Lett.* **75**, 535 (2006).
- [42] B. Efron and R. Tibshirani, *An Introduction to the Bootstrap* (CRC, Boca Raton, Florida, 1993).
- [43] J. Franke and W. Hardle, *Ann. Stat.* **20**, 121 (1992).
- [44] Y. Benjamini and Y. Hochberg, *J. R. Stat. Soc. Ser. B (Methodol.)* **57**, 289 (1995).
- [45] A. Reiner, D. Yekutieli, and Y. Benjamini, *Bioinformatics* **19**, 368 (2003).
- [46] J. D. Storey and R. Tibshirani, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 9440 (2003).
- [47] P. Dutta and P. Horn, *Rev. Mod. Phys.* **53**, 497 (1981).
- [48] T. Gisiger, *Biol. Rev. Cambridge Philos. Soc.* **76**, 161 (2001).
- [49] K. Linkenkaer-Hansen, V. V. Nikouline, J. M. Palva, and R. J. Ilmoniemi, *J. Neurosci.* **21**, 1370 (2001).
- [50] A. Aarabi, F. Wallois, and R. Grebe, *Brain Res.* **1188**, 207 (2008).

- [51] K. Ansari-Asl, L. Senhadji, J.-J. Bellanger, and F. Wendling, *Phys. Rev. E* **74**, 031916 (2006).
- [52] T. Kreuz, F. Mormann, R. Andrzejak, A. Kraskov, K. Lehnertz, and P. Grassberger, *Physica D* **225**, 29 (2007).
- [53] F. Mormann, T. Kreuz, C. Rieke, R. Andrzejak, A. Kraskov, P. David, C. Elger, and K. Lehnertz, *Clin. Neurophysiol.* **116**, 569 (2005).
- [54] H. Osterhage, F. Mormann, M. Staniek, and K. Lehnertz, *Int. J. Bifurcation Chaos Appl. Sci. Eng.* **17**, 3539 (2007).
- [55] R. Quiñero, A. Kraskov, T. Kreuz, and P. Grassberger, *Phys. Rev. E* **65**, 041903 (2002).
- [56] J. P. M. Pijn, P. C. M. Vijn, F. H. Lopes da Silva, W. Van Emde Boas, and W. Blanes, *Adv. Epileptology* **17**, 272 (1989).
- [57] H. K. M. Meeren, J. P. M. Pijn, E. L. J. M. Van Luijckelaar, A. M. L. Coenen, and F. H. Lopes da Silva, *J. Neurosci.* **22**, 1480 (2002).
- [58] A. Seth, *Network Comput. Neural Syst.* **16**, 35 (2005).
- [59] B. Schelter, M. Winterhalder, R. Dahlhaus, J. Kurths, and J. Timmer, *Phys. Rev. Lett.* **96**, 208103 (2006).
- [60] S. Achard, R. Salvador, B. Whitcher, J. Suckling, and E. Bullmore, *J. Neurosci.* **26**, 63 (2006).
- [61] M. Valencia, J. Martinerie, S. Dupont, and M. Chavez, *Phys. Rev. E* **77**, 050905(R) (2008).
- [62] S. Pajevic and D. Plenz, *PLOS Comput. Biol.* **5**, e1000271 (2009).
- [63] M. R. Leadbetter, G. Lindgren, and H. Rootzén, *Extremes and Related Properties of Random Sequences and Processes* (Springer-Verlag, New York, 1983).
- [64] E. M. Izhikevich, *IEEE Trans. Neural Netw.* **14**, 1569 (2003).
- [65] E. M. Izhikevich, *Neural Comput.* **18**, 245 (2006).
- [66] E. V. Lubenov and A. G. Siapas, *Neuron* **58**, 118 (2008).
- [67] U. Mitzdorf, *Physiol. Rev.* **65**, 37 (1985).
- [68] P. Lakatos, A. Shah, K. Knuth, I. Ulbert, G. Karmos, and C. Schroeder, *J. Neurophysiol.* **94**, 1904 (2005).
- [69] C. Koch, *Biophysics of Computation* (Oxford University Press, New York, 1999).
- [70] F. Rosenow and H. Lüders, *Brain* **124**, 1683 (2001).
- [71] A. Bragin, J. Engel, C. L. Wilson, I. Fried, and G. Buzsáki, *Hippocampus* **9**, 137 (1999).
- [72] R. D. Traub, M. A. Whittington, E. H. Buhl, F. E. N. LeBeau, A. Bibbig, S. Boyd, H. Cross, and T. Baldeweg, *Epilepsia* **42**, 153 (2001).
- [73] B. Litt and J. Echauz, *Lancet Neurol.* **1**, 22 (2002).