

Non-parametric test for connectivity detection in multivariate autoregressive networks and application to multiunit activity data

M Gilson, A Tauste Campo, X Chen, A Thiele, G Deco

March 17, 2017

Abstract

Directed connectivity inference has become a cornerstone in neuroscience following the recent progress in neuroimaging and electrophysiological techniques to characterize anatomical and functional networks. This paper focuses on the detection of existing connections from the observed activity in networks of 50 to 150 nodes with linear feedback in discrete time. Through the variation of multiple network parameters, our numerical results indicate that directed connections - in the time domain - are more accurately estimated based on the coefficients obtained from multivariate autoregressive (MVAR) than Granger causality analysis, which is based on the error residuals of the same MVAR linear regression. Based on these findings, we propose a non-parametric significance test for connectivity detection, which achieves a good control of false positives (type 1 error) and is robust to various network topologies. When generating surrogate distributions, we compare the effects of circular shifts, random permutations and phase randomization of the observed time series, each breaking down covariances in a specific manner: the MVAR estimates from those shuffled covariances build a null-hypothesis distribution for each connection, from which the original connectivity estimate can be compared. We apply our method to multiunit activity data recorded from Utah electrode arrays in monkey and examine the detected interactions between 25 channels for a proof of concept. The results unravel a non-trivial underlying connectivity structure, which differentiates the effect of incoming and outgoing connections.

1 Introduction

In recent years, there has been a growing interest in developing multivariate techniques to infer causal relations among time series. The initial formulation of the problem goes back to the seminal work by Granger in 1960's [20] motivated by the analysis of the pairwise influence between economic time series. In this work, Granger decomposes the cross-spectrum of two autoregressive time series into two directional components that account for the potential causal influences between each other. A general solution of the problem in multivariate scenarios was developed a decade later by the introduction of multivariate autoregressive (MVAR) processes, which allow the estimation of causal relationships between nodes in networks with feedback based on their observed activity [2, 18, 19, 27]. The MVAR was further combined with spectral analysis to develop the directed transfer entropy function [22, 23], which has been employed to analyze connectivity patterns in neurobiological systems [3, 49]. Estimating directed networks has been regarded in a number of scientific disciplines, gaining central relevance in the field of genetics, in which the inference of gene regulatory networks is a cornerstone. In this field, the growth of inference methods has been specially notorious over the last decade [8, 52].

In neuroscience, this inference problem has been transposed to analyse interactions between neuronal populations from spiking activity or neuroimaging measurements such as fMRI, EEG and MEG [43, 34, 31, 30, 40, 26]. Two types of estimation procedures may be distinguished: measures relying on an underlying interaction model such as Granger causality analysis [13] and dynamic causal modeling (DCM) [16] on the one hand; and model-free measures such as transfer entropy [38] and directed information [29], which make minimal model assumptions on the other hand. Although model-free approaches have proven useful to describe neural propagation at spike-train level [41, 46], certain assumptions are required when estimating interactions at the neuronal population level, in which broader spatial and temporal scales contribute to shaping the signals. Motivated by data-driven practical problems, methodological refinements of Granger causality analysis (or MVAR-based methods) have considered additive noise [48] or measurement noise via state-space models [15, 5].

In this article we compare connectivity estimation techniques for the MVAR dynamic model in the time domain (not frequency), under the specific angle of the application to brain network activity. As a generative model, the MVAR process is canonically related to Granger causality analysis: the linear regression in the upper right inset of Fig. 1A provides both coefficients and residuals, the latter being viewed as the remaining uncertainty in the prediction of the target time series by its source(s). By comparing the residuals of two linear regressions - one involving a supposed source node and one without it - in a log ratio, Granger causality estimates the effective contribution of one node onto another in a network [6]. Since these log ratios asymptotically converge to known distributions, parametric statistical tests have been developed to assess the significance of Granger's estimates [4]. Instead, our proposed method directly assesses the significance of the MVAR coefficients to infer the existence of network connections. To do so, we propose a non-parametric significance test in the regression coefficients space. Previous literature on non-parametric testing for Granger causality has resorted to surrogate data generated by trial shufflings [10, 32], bootstrap procedures [11] or by phase randomization in frequency-domain measures [12, 25]. Here we focus on within-trial surrogate tests for time-domain coefficients and compare them across standard techniques [39, 14, 50]. In the context of recurrent networks, we provide numerical evidence that these tests can achieve a good control of the false-alarm rate and might improve the miss rate by properly adapting the null distribution to each connection. The focus of the present analysis is on the case where we observe more time samples (a few thousands per node) than the network size (about a hundred nodes). Within this regime, we test the robustness of the detection method for a broad range of network parameters and various topologies.

2 Methods: multivariate autoregressive model and interaction estimation

The activity in the MVAR process - a.k.a. noise-diffusion discrete-time network - is described by the following equation:

$$x^t = Ax^{t-1} + \zeta^t, \quad (1)$$

where the matrix A describes the interactions between coordinates of the vector $x^t = (x_i^t)$ with time t being an integer and node index $1 \leq i \leq N$. Here we constrain our study to the case where ζ^t is Gaussian (possibly cross-correlated noise), whose realizations are time independent for successive time steps. Without loss of generality, we assume that all variables ζ^t have zero means, giving the same property for all x_i . We only consider MVAR processes of order 1 in a first place, but will extend the work to the case of order 2 in a later section.

2.1 Granger causality analysis

Granger causality analysis is usually presented using time series and the estimation of non-zero coefficients in A from observed activity over a period $1 \leq t \leq T$ relies on the linear regression of the activity x_i^t of a given node i at time t by the past activity of a subset \mathcal{S} of network nodes:

$$x_i^t = \sum_{j \in \mathcal{S}} A_{ij} x_j^{t-1} + \epsilon^t \quad (2)$$

for $2 \leq t \leq T$. When T is large, the coefficients a_{ij} converge toward A_{ij} . We define the residual ϵ as the standard deviation of the ϵ^t for the ordinary least-square (OLS) regression in Eq. (2)

$$\epsilon(x_i^{2 \leq t \leq T} | x_i^{1 \leq t \leq T-1}) = \sqrt{\sum_t (\epsilon^t)^2}, \quad (3)$$

with a notation similar to conditional probabilities; the superscript t indicates the considered time range and the subscripts indicate the nodes involved. To detect the existence of connection $j \rightarrow i$ in a network, two types of Granger causality analysis exist: ‘unconditional’ and ‘conditional’ [18, 19]., they consider the comparisons of the following residuals:

$$\begin{aligned} \text{GRu}(x_j \rightarrow x_i) &= \ln \left[\frac{\epsilon(x_i^{2 \leq t \leq T} | x_i^{1 \leq t \leq T-1})}{\epsilon(x_i^{2 \leq t \leq T} | x_{i,j}^{1 \leq t \leq T-1})} \right]; \\ \text{GRc}(x_j \rightarrow x_i) &= \ln \left[\frac{\epsilon(x_i^{2 \leq t \leq T} | x_{1, \dots, j-1, j+1, \dots, N}^{1 \leq t \leq T-1})}{\epsilon(x_i^{2 \leq t \leq T} | x_{1, \dots, N}^{1 \leq t \leq T-1})} \right]. \end{aligned} \quad (4)$$

For both GRu and GRc, which have a univariate target node x_i , the usual parametric test for significance relies on the F statistics, which performs better for small number of samples [4]. The null hypothesis of no interaction for GRu($x_j \rightarrow x_i$) corresponds to $m = T$, $p = 1$, $n_x = 1$ and $n_y = 2$ using the notation in [4]

$$\frac{\epsilon(x_i^{2 \leq t \leq T} | x_i^{1 \leq t \leq T-1}) - \epsilon(x_i^{2 \leq t \leq T} | x_{i,j}^{1 \leq t \leq T-1})}{\epsilon(x_i^{2 \leq t \leq T} | x_{i,j}^{1 \leq t \leq T-1})} = [\exp(\text{GRu}_{ij}) - 1] > \frac{\phi(\alpha, 1, T-3)}{T-3} \quad (5)$$

with α the desired sensitivity and ϕ the inverse survival function of the F-distribution [51]. The equivalent for GRc corresponds to $n_y = N$, yielding

$$\frac{\epsilon(x_i^{2 \leq t \leq T} | x_{1, \dots, N}^{1 \leq t \leq T-1}) - \epsilon(x_i^{2 \leq t \leq T} | x_{1, \dots, j-1, j+1, \dots, N}^{1 \leq t \leq T-1})}{\epsilon(x_i^{2 \leq t \leq T} | x_{1, \dots, j-1, j+1, \dots, N}^{1 \leq t \leq T-1})} > \frac{\phi(\alpha, 1, T-N-1)}{T-N-1}. \quad (6)$$

We also use non-parametric tests for GRc by performing a circular shift (see details below in Section ‘Generation of surrogate time series’) on the observed driver (source) time series for all tested connections. This provides a null distribution for the log ratio, with which the actual estimated log ratio can be compared.

2.2 Multivariate autoregressive (MVAR) estimation

To detect the existence of connections $A_{ij} > 0$, another possibility is to estimate the coefficients themselves, which can be done using the covariances of the observed activity variables x^t [27]:

$$\hat{Q}_{ij}^\tau = \frac{1}{T - \tau_{\max} - 1} \sum_{1 \leq t \leq T - \tau_{\max}} (x_i^{t+\tau} - \bar{x}_i)(x_j^t - \bar{x}_j), \quad (7)$$

where T denotes the number of successive samples indexed by t , $\tau \in \{0, 1\}$ is the time shift (here $\tau_{\max} = 1$) and the observed mean activity for each node is $\bar{x}_i = \frac{1}{T} \sum_t x_i^t$. The Yule-Walker equation gives a consistency equation for the theoretical covariance matrices (without hat) in terms of the connectivity A in the dynamics described by Eq.(1):

$$Q^1 = A Q^0 . \quad (8)$$

The estimation of network interactions relies on evaluating A from Eq. (8) for the empirical covariance matrices defined as Eq. (7) and calculated for a given time series:

$$A = \widehat{Q}^1 (\widehat{Q}^0)^{-1} . \quad (9)$$

Note that this OLS estimate corresponds to the linear regression related to $\epsilon(x_{1,\dots,N}^{2 \leq t \leq T} | x_{1,\dots,N}^{1 \leq t \leq T-1})$ and is also to the linear model with maximum likelihood under the assumption that the observed process is Gaussian.

2.3 MVAR of order 2

Eq. (1) can be extended to the case where the activity vector x^t is determined by the two previous time steps:

$$x^t = A^1 x^{t-1} + A^2 x^{t-2} + \zeta^t . \quad (10)$$

For the second order, we use $\tau_{\max} = 2$ in Eq. (7) and the estimation of A^1 and A^2 via the Yule -Walker equation is given by [27, p. 86]

$$\tilde{A} = \tilde{Q}^1 (\tilde{Q}^0)^{-1}, \quad (11)$$

with the block matrices

$$\begin{aligned} \tilde{A} &= \begin{pmatrix} A^1 & A^2 \end{pmatrix} , \\ \tilde{Q}^0 &= \begin{pmatrix} \widehat{Q}^0 & \widehat{Q}^1 \\ (\widehat{Q}^1)^\dagger & \widehat{Q}^0 \end{pmatrix} , \\ \tilde{Q}^1 &= \begin{pmatrix} \widehat{Q}^1 & \widehat{Q}^2 \end{pmatrix} . \end{aligned} \quad (12)$$

The coefficients of A^1 and A^2 can thus be estimated using a matrix multiplication and an inversion involving the covariances, as with the first-order case in Eq. (9).

2.4 Generation of surrogate time series

In this paper, we consider circular shifts (CS), random permutations (RP) and phase randomization (PR) to shuffle the time points of the observed time series. From the original x_i^t with $1 \leq t \leq T$,

- CS draws a random integer $t_0 \in \{1, \dots, T\}$ and returns $(x_i^{t_0}, \dots, x_i^T, x_i^1, \dots, x_i^{t_0-1})$;
- RP draws a random permutation σ of $\{1, \dots, T\}$ such that each integer appears once (and only once) and returns $x_i^{\sigma(t)}$;
- PR calculates the discrete Fourier transform $\mathcal{F}(x_i^t)$ of the original x_i^t , then multiplies each of the T coefficients of $\mathcal{F}(x_i^t)$ by $\exp(2\pi i \phi^t)$ with ϕ^t randomly chosen in $[0, 2\pi]$, and performs the inverse transform.

Importantly, these operations are applied to each time series independently of the others.

In addition, we consider the replacement of all time series in the network by T normally distributed variables with a standard deviation equal to the mean of the standard deviations of x_i^t along the time axis, then averaged for all nodes. We refer to these surrogates as STD.

2.5 Experimental setup and processing of electrode measurements to extract MUAe activity

All procedures were carried out in accordance with the European Communities Council Directive RL 2010/63/EC, the US National Institutes of Health Guidelines for the Care and Use of Animals for Experimental Procedures, and the UK Animals Scientific Procedures Act. Two male macaque monkeys (5 - 14 years of age) were used in the experiment; only the data for the first one is used here. A surgical operation was performed under sterile conditions, in which a custom-made head post (Peek, Tecapeek) was embedded into a dental acrylic head stage. Details of surgical procedures and post-operative care have been published previously [47]. During the surgery microelectrode chronic Utah arrays (5*5 grids), attached to a CerePort base (Blackrock Microsystems) were implanted into V1. Electrodes were 1 mm in length in line with procedures described in [45]. Stimulus presentation was controlled using CORTEX software (Laboratory of Neuropsychology, NIMH, <http://dally.nimh.nih.gov/index.html>) on a computer with an Intel Core i3-540 processor. Stimuli were displayed at a viewing distance of 0.54 m, on a 25" Sony Trinitron CRT monitor with a resolution of 1280 by 1024 pixels, yielding a resolution of 31.5 pixels / degree of visual angle (dva). The monitor refresh rate was 85 Hz for monkey 1, and 75 Hz for monkey 2. A gamma correction was used to linearize the monitor output, and the gratings had 50% contrast. Monkeys performed a passive viewing task where they fixated centrally while stationary sinusoidal grating of either horizontal or vertical orientation and 2 cycle per degree spatial frequency, were presented in a location that covered all receptive fields recorded from the 25 electrode tips. Stimuli were presented 500 ms after fixation onset for 150 ms. Raw data were acquired at a sampling frequency of 32556 Hz using a 64-channel Digital Lynx 16SX Data Acquisition System (Neuralynx, Inc.). Following each recording session, the raw data were processed offline using commercial (Neuralynx, Inc.). Signals were extracted using Cheetah 5 Data Acquisition Software, with bandpass filtering set to allow for spike extraction (600-9000 Hz) and saved at 16-bit resolution.

In the present study, we focus on the period starting 200 ms before and finishing 200 ms after the stimulus onset, for 4 conditions (vertical gradings with pre/post cue in the receptor/opposite field) that will not be compared in details. The electrode recordings is firstly down-sampled from 32556 Hz to 1000 Hz. A high-pass filter above 400 Hz is then applied - 3rd-order Butterworth filter at 0.8 of the Nyquist frequency [51] - followed by a smoothing of 4 ms to extract the envelope of the resulting signal, by retaining the 250 time points of 1000-ms period surrounding the stimulus onset.

3 Benchmark of detection performance for synthetic data

The workflow of the benchmark for the estimation procedure is schematically represented in Fig.1A. We first consider a MVAR process defined by Eq. (1) with given connectivity matrix A and input covariances (obtained by mixing independent Gaussian processes) to generate the activity of the network. From the observed activity over a period of duration T , we estimate the coefficients matrix A using the covariances as described in Eq. (9). We also perform the linear regressions of each node activity over the past activity of given subsets of nodes corresponding to the unconditional (GRu) and conditional (GRc) Granger causality analysis, from which we calculate the ratios of residuals in Eq. (4). Actually, these estimates correspond to the same OLS regression (top right in Fig. 1A) and the difference resides in the spaces where they lie: coefficients versus residuals.

For each method, the level of predictability can be measured by the relationship between the estimated values and the original connectivity values, as illustrated in Fig. 1B (left). To discriminate between existing and absent connections, one can apply a common threshold for all connections (top thread); by sliding this threshold, we obtain the ROC curve with

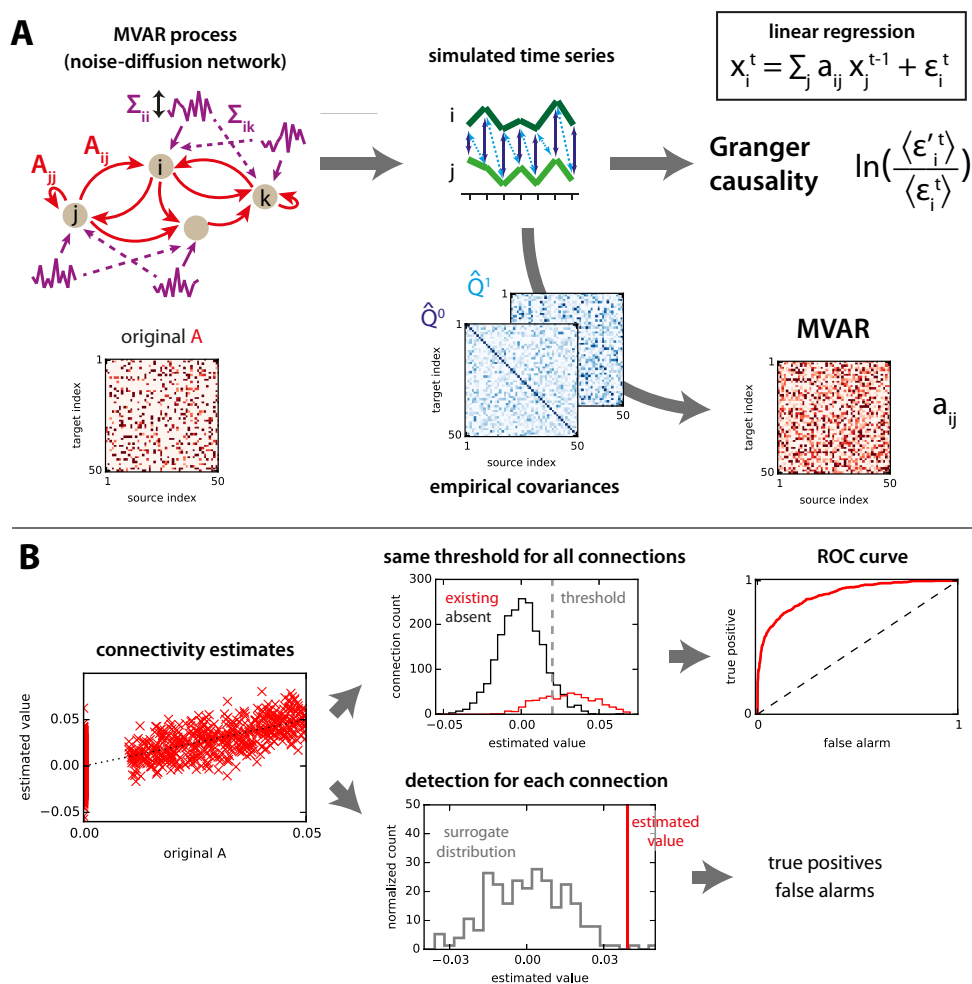


Figure 1: Network model and connectivity estimation. A: For a given directed connectivity A and input covariances Σ (left), the network activity (middle) is simulated using Eq. (1). From the observed time series, the existing interactions in the original connectivity can be estimated (right): Granger causality analysis uses the residuals of linear regressions (ϵ in the upper right equation; see Methods for details about the residuals used in the log ratio), whereas MVAR corresponds to the coefficients. Note that MVAR can be obtained using the empirical covariance matrices \hat{Q}^0 and \hat{Q}^1 , see in Eq. (7) for $\tau = 0$ and 1. **B:** The left panel compares the estimated values to the original values for all connections in the network. The upper thread displays the distributions of estimated values for existing and non-existing connections in the original network. Using a sliding threshold (vertical dashed gray line) on the estimated values, one can calculate the ROC curve (right). The lower thread compares the estimated value for a single connection to a null distribution. From this, the choice whether the connection exists or not is made for each individual connection, yielding a single pair of true-positive and false-alarm rates.

the rate of false alarms on the x-axis and true positives on the y-axis. The area under the curve indicates in a single value how well the ranking of estimated value performs for the detection of connections in the original connectivity. Alternatively, an individual test can be made for each connection in the network, for example by comparing the estimated value to a null distribution (bottom thread). Here again, we obtain two rates of false positives and negatives.

3.1 Coefficients from linear regression potentially predict better existing connections than residuals

We start with the comparison between the predictability of coefficients and residuals for MV, GRu and GRc for all connections in a given network. To do so, we simulate 500 randomly connected graphs, which are simulated with different sizes ($N = 50$ to 150 nodes), density and connectivity weights (uniformly drawn in a randomly chosen range $[w_{\min}, w_{\max}]$); here inputs are *not* correlated: the ζ_i are independent across node indices in Eq. (1). For each network configuration, we evaluate the accuracy for connection detection via the area under the ROC curve (see the upper thread in Fig. 1B). Fig. 2A displays this ROC-based accuracy as a function of the number of observed time samples (x-axis) represented by violin plots for 500 randomly connected networks. When considering many samples (10^4), all methods perform well. However, for smaller sample sets, the MVAR method exhibits superior performance than GRu: as measured by the Mann-Whitney test, $p < 10^{-45}$, $p < 10^{-19}$ and $p < 10^{-5}$ for the three values of observed samples, respectively.

Although error residuals ratios are in a different space from the true coefficients in A , one expects some degree of correlation between them, such that Granger causality analysis effectively detects connections. In Fig. 2B, both GRu and GRc estimates have a ranking similar to the original A weights (as measured by the Spearman correlation) for $T = 10000$ observed time samples, but this weakens dramatically for $T \leq 3000$. In contrast, the ranking for estimated MVAR coefficients reflects much better the original A for $T \leq 3000$. In the studied networks, GRu performs slightly better than GRc. As analyzed in previous studies, this can be consequence of the balance between redundant and synergistic activity exhibited by the simulated network nodes [44]. To shed light into the effect of the network structure, we next examine how the ROC-based performance in Fig. 2A depends on the controlled network parameters. The four panels in Fig. 2C display the trends of the values for the 500 networks as a function of the network size N , the network density, the minimum weight in the original network (w_{\min} mentioned above) and the mean sum of incoming weights per node. For illustration purpose, the 500 networks are grouped in quartiles for each parameter. Not surprisingly, the estimation accuracy of all methods decreases as a function of the network size N and density, and increases as a function of the minimum connectivity weight and the mean incoming weight per node. More interestingly, in challenging configurations with small weights, MVAR consistently shows a superior performance by a larger gap compared to Granger causality analyses. These findings support the use of coefficients to robustly detect connections in recurrently connected networks. Note that GRu performs on average slightly better than GRc here: the discrepancy decreases as a function of the network density, which may follow from lower redundancy in the recurrent network [44].

3.2 A robust non-parametric significance test for MVAR

We have so far examined the performance of different estimation methods based on the ROC curve, which considers a single threshold for all connections in a network and integrates the information about false alarms and true detection over the whole range of estimated values. However, in the context of real data, the decision for the existence of a connection typically relies on comparing the value of the connection estimate with a given statistical threshold. For GRu and GRc, such parametric tests have been developed, for example, based on the

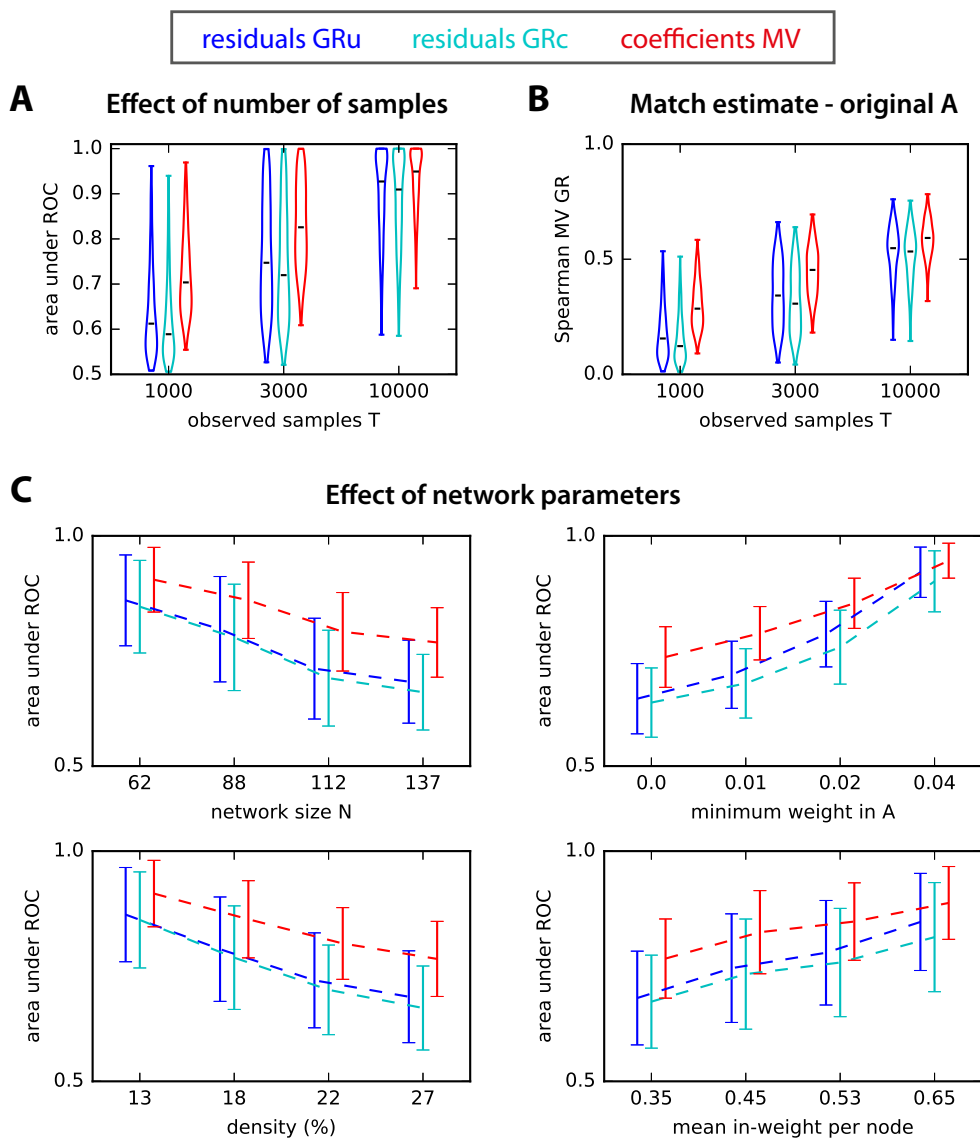


Figure 2: **ROC accuracy and computation time.** **A:** Area under ROC for estimated A obtain from Granger (unconditional for GRu and conditional for GRc) and MVAR. The x-axis indicates three sample size T for the observed network activity. The boxplots indicate the median (horizontal bar in the middle of the box) as well as the upper and lower tail spreads (standard deviation) over 500 simulated networks of various sizes and connectivity strengths. **B:** Match of the ranking between GRu, GRc and MVAR estimates and the original connectivity weights A , as measured by the Spearman correlation. The plotted values correspond to the 500 networks in A and the x-axis indicates the sample size. **C:** Effect of network parameters on ROC-based detection performance. Influence of network size N , connectivity density, sum of recurrent connectivity strengths, minimum weight w_{\min} in A , mean noise on the diagonal of Σ and mean off-diagonal noise in Σ on the ROC-based accuracy in Fig. 2A. In each graph, the network configurations have been grouped in quartiles according to the parameter plotted on the x-axis, and the corresponding group mean and standard deviations are indicated; the curves are displaced horizontally to favor the legibility.

F statistics [4]. Equivalently, it is sufficient to know how the values of the estimates for absent connections are distributed, in order to select a desired rate of false alarms (type-1 error). In this section, we develop a significance test for the MVAR estimates by providing a null-hypothesis distribution for absent connections.

Our approach relies on the fact that covariances reflect the underlying connectivity: we thus construct the null distribution estimates by performing a random permutation for each of the observed time samples, as described in Fig. 3A; other methods will be tested in a later section. From the resulting covariance matrices, we evaluate a surrogate connectivity matrix. This procedure destroys the information in both surrogate covariance matrices, except for the variances on the diagonal of \hat{Q}^0 alone. The core result underlying our surrogate approach is illustrated in Fig. 3B: the distribution of surrogate estimates (thick black line) is compared against the distribution of existing (red) and absent (blue) connections in a simulated random network model: the surrogate distribution in black provides a good approximation for the distribution of estimates for non-existing connections in blue.

We consider two options - corresponding to the two threads in Fig. 1B - to test the existence of a connection from an MVAR estimate while keeping the false alarm rate under control.

- The *global* test relies on the null distribution corresponding to the black histogram in Fig. 3C, which is obtained by grouping together all SN^2 matrix elements of all matrices for $S = 200$ surrogates. From that surrogate distribution, we perform a detection test by setting a threshold corresponding to a percentage of the right tail equal to the desired false-alarm rate (here 2%), as illustrated by the vertical gray dashed line.
- Instead, the *local* test uses for each connection the surrogate distribution of S values, corresponding to the same matrix element in each of the S surrogates. From that distribution in red in Fig. 3C, the detection threshold is defined similarly (vertical dark red dashed line).

The rationale behind these two choices lies in the trade-off between taking into account spatial heterogeneity in the network and gaining larger sample size, as illustrated by the distinct thresholds in Fig. 3C. Note also that the F statistical test for Granger causality analysis corresponds to a global threshold on the log ratio values. When varying the desired false-alarm rate, the two tests perform well, as illustrated in Fig. 3D by their location close to the ROC curve (circles and triangles for local and global, respectively).

To assess the effect of the small variability observed in Fig. 3D over the randomness of network configurations, we simulate 500 randomly connected networks with the same parameters as in Fig. 2, except for the size $50 \leq N \leq 90$ and the presence of input cross-correlations. Note that, from Fig. 2, the chosen size N corresponds to a situation where Granger causality analysis performs relatively well as compared to MVAR. The control of the false-alarm rate is displayed in Fig. 3E for both local and global tests with various numbers S of surrogates. The control of false-alarm rates is close to perfect across various values for all S and both tests (local and global), demonstrating the robustness of the proposed method for randomly connected networks. Following, we fix the desired false-alarm rate to 2% and evaluate the miss rate (true negatives) of both methods depending on the actual weight strength: in Fig. 3F, connections are grouped in terciles for each network configuration. Interestingly, the local test improves with the number S of surrogates (right panel), whereas the global test exhibits a constant performance for all S (left panel). Note that the advantage of the local test over the global test particularly concerns connections with small weights, which are difficult to detect, in line with Fig. 2C (see influence of the minimum original weight).

Because GRu does not take all nodes into account, the presence of spatially correlated noise (indicated by the purple dashed arrows in Fig. 1A) dramatically affects the false-alarm rate when using the parametric significance F-test [4], as shown in Fig. 4A by the dark

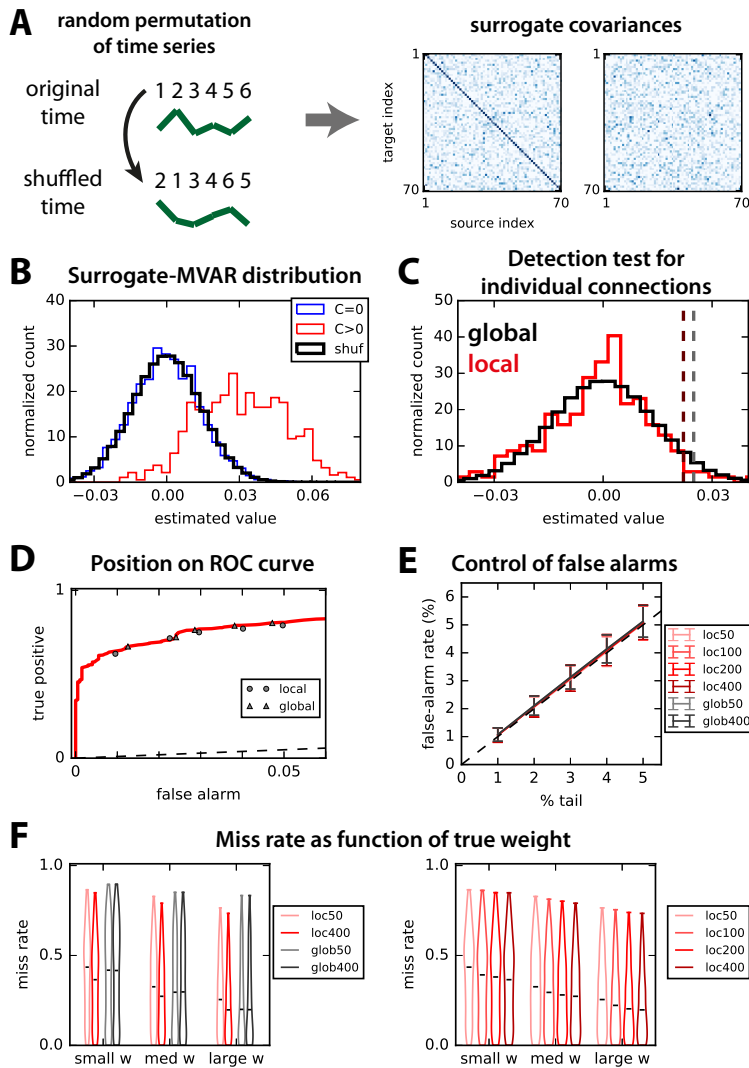


Figure 3: Non-parametric test to assess significance based on random permutations. **A:** Schematic illustration of time rolling applied on the observed time series to generate the surrogate covariances. **B:** Pooled distributions of estimated weights for the existing (in red) and absent (blue) connections. The thick black curve indicates the distribution of connections over 100 surrogates, which closely matches the blue distribution. **C:** For a given connection, we compare two methods: for ‘local’ in red, the null distribution corresponds to the matrix elements for the same connection in 200 surrogates; for ‘global’ in black, the null distribution is the pooled distribution for all N^2 elements of the 200 surrogate matrices (same as in B). The dark red and gray dashed lines indicate the detection thresholds corresponding to the 4% tail for those two options. **D:** The performance of the two non-parametric methods for the thresholds described in B and C is displayed on the ROC curve for a desired false-alarm rate ranging from 1 to 5%. Triangles indicate the local test and circles the global test. **E:** Comparison of the desired (% of the tail of null distribution) and actual rate of false alarms for the local and global tests when varying the the number S of surrogates (see figure in legend). Error bars indicate one standard deviation for 500 random networks; importantly, inputs for these networks have cross-correlation, unlike Fig. 2. **F:** Influence of the strength of original weight on the detection performance for a desired false-alarm rate set to 2%. In both panels, lighter colors indicate smaller numbers of surrogates S , in red for the local test and gray for the global test (see legends).

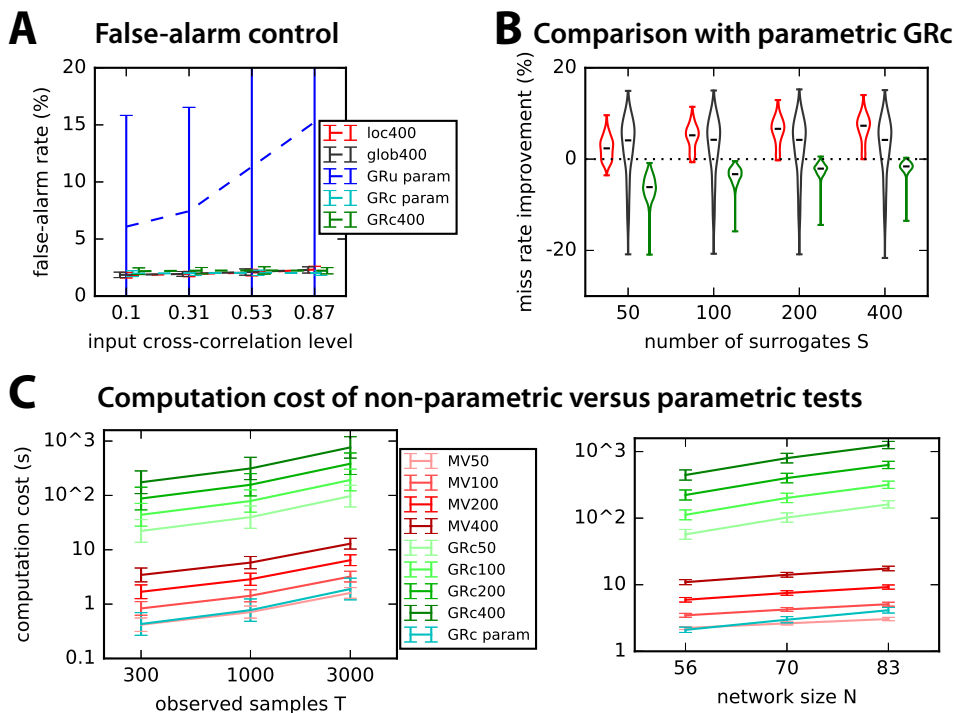


Figure 4: **Comparison of our surrogate-based method with Granger causality analysis.** **A:** Comparison of the parametric tests for GRu (blue curve) and GRc (cyan) with the non-parametric methods for GRc (green) and MVAR (red for local test and gray for global). The x-axis indicates the strength of input correlations (i.e., pink noise) in the simulated network. The desired false-alarm rate is set to 2% as in Fig. 3F and the number of observed time samples is $T = 3000$. Error bars indicate one standard deviation over the 500 random networks as in Fig. 3E. **B:** Comparison of the miss rate improvement with respect to GRc for the 500 networks in A as a function of the number S of surrogates (x-axis). Red indicates the local test, gray the global test and green GRc. **C:** Comparison of the computational cost for the surrogate-based method and parametric tests as a function of the number T of observed samples (left) and network size (right).

blue dashed curve. This is solved by the “complete” linear regression in GRc, achieving a quasi perfect control irrespective of the input correlation level for both parametric and non-parametric tests (cyan and green dashed curves, respectively), as our non-parametric tests do (Fig. 3E). However, the miss rate improves (i.e., decreases) for both local and global tests on average as compared to the parametric GRc over the 500 considered random networks in Fig. 4B. For $S \geq 200$, it even becomes better in all cases. Note that the seemingly small false-alarm improvements of about 7% actually corresponds to more than 50 existing connections per network here. However, the non-parametric test for GRc performs worse than the parametric test for GRc in this case.

From Figs. 3F and Fig. 4B, we conclude that the local test is preferable to the global test provided $S \geq 200$ surrogates are generated. However, the computational cost increases with S , as illustrated in Fig. 4C by the red curves. Note that GRc (in cyan) takes the same time to calculate as $S = 50$ surrogates. However, our non-parametric method scales better than GRc when the network size increases. The non-parametric test for GRc takes longer time to compute, but further optimization of the calculations could be made that were not incorporated here.

3.3 Comparison of generation methods for surrogates for non-parametric MVAR

The fact that the OLS MVAR estimates can be obtained via the two covariance matrices (with and without time shift, see Fig. 1A) hints at possible methods to generate surrogate by destroying the information in these covariances. Methods to generate surrogate time series have been widely used in the past: circular shifts of the time series [14], random permutation [50] and phase randomization [39] to generate a null distribution for the ratios in Eq. (4); they are referred to here as CS, RP and PR, respectively. We thus consider these three methods (cf. box in Fig. 5), as well as surrogate time series that only preserve the mean standard deviation averaged over the network (STD), so as test to which extent it is important to preserve the spatial heterogeneity of the nodes' activity. See Methods for details about the calculations.

The control of false alarms for local tests in Fig. 5A and B is better for CS and RP, whereas the detection of true connections is similar for the four methods over 500 random networks of size $N = 70$. However, CS fails to detect self connections (Fig. 5C). The reason is that, because CS surrogates preserve the autocovariances in the time-shifted covariance, they fail to build a proper null distribution for self-connections. The influence of the number of samples used in the estimation is similar for all methods, as illustrated in Fig. 5D. The comparison with STD (purple), which averages the covariance statistics over the whole network, suggests that the local test makes a good use of the heterogeneous information across nodes. As a conclusion, we retain RP as the best option.

3.4 Influence of network topology

In this part, we test and compare the robustness of global and local surrogate-based detection tests to specific connections and topological configurations. Here, $T = 3000$ observed samples and we compare the local and global tests with $S = 400$ surrogates for 500 networks of each type. In all cases, the simulated networks have the same size $N = 70$, but vary in connectivity density, distribution of recurrent weights and level of input cross-correlation. We compare the miss rate for unidirectional, reciprocal and self connections in the random networks examined until now (and a desired 2% of false alarms). Fig. 6A shows that the miss rate is similar in unidirectional and reciprocal connections with the local test, which performs slightly better than the global test (as in Fig. 3F).

Now we consider more elaborate network topologies than the random connectivity (Erdős-Rényi) considered so far, namely modular and hierarchical networks. In Fig. 6B, we simulate 500 modular networks with two groups (green and blue) linked by hubs (red, about 5 to 15% of the nodes). Interestingly, intra-group and hub-group connections have a similar miss rate with regard to using local and global surrogates. In Fig. 6C, we simulate hierarchical networks of three layers, for which connections either link the center and an intermediate node, or link an intermediate node and a leaf, or are self connections. This network type is much sparser than the two types in A and B, yielding a quasi perfect detection performance for all types of connections (miss rate < 0.1 in Fig. 6C). In all cases, the local test performs better than the global test. However, the control of false-alarm rate is similar for both tests with all topologies, as can be seen in Fig. 6D.

Finally, we consider a network with both excitatory and inhibitory connections (with a inhibitory ratio equal to 5 to 50% of all) and perform the test by defining a threshold on both tails of the null distributions. As can be seen in Fig. 6E, the positive/negative nature of the connection weights affect neither the false-alarm nor the miss rate. However, the performance is poorer than with excitatory connections only.

We conclude that, in those networks with spatial heterogeneity as with randomly connected networks, the local test with an individual null distribution per connection performs better than the global test. Recall that an improvement of the miss rate by 1% in a network

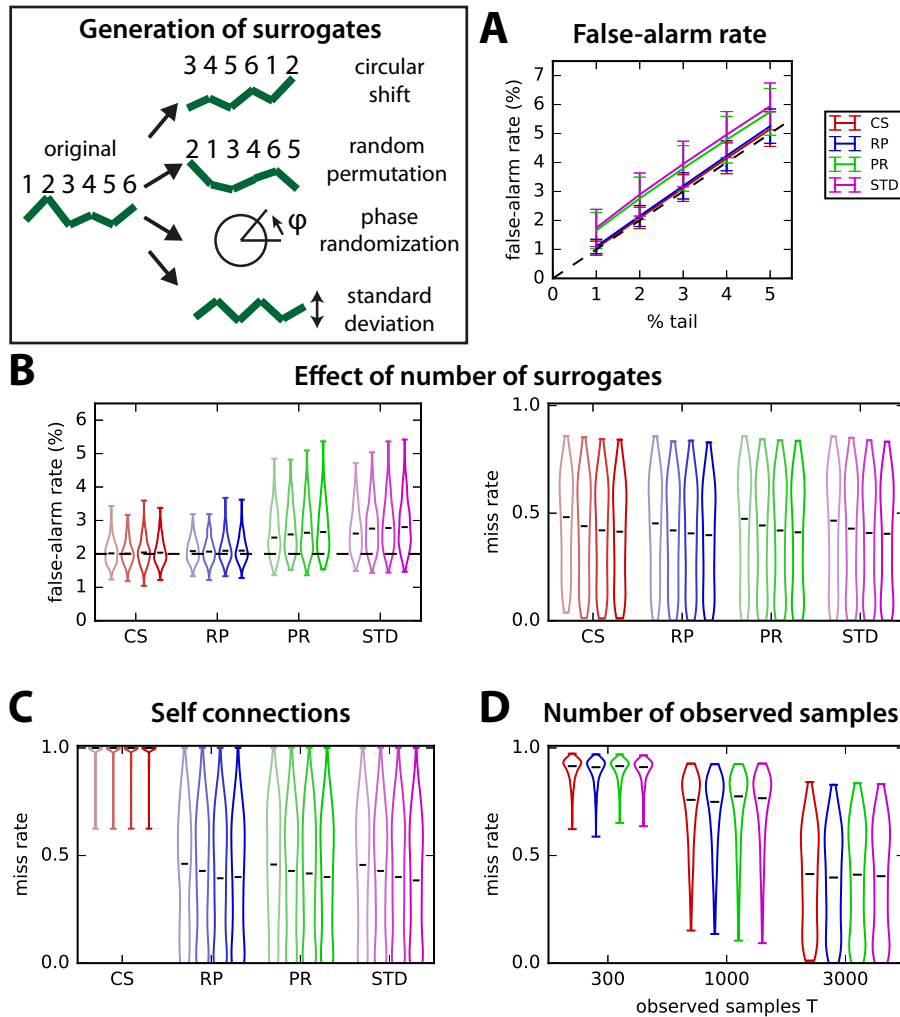


Figure 5: Comparison of the local test for 4 surrogate generation methods. Surrogates are generated by independently performing for each original time series 1) random permutations, 2) circular shifts, 3) phase randomization and 4) replacing the original by a new random time series with the mean standard deviation averaged over all of the original time series in the network. **A:** Comparison of the control of the false-alarm rate for various thresholds on the tail of the distributions of 400 surrogates (% indicated on the x-axis). The error bars correspond to the variability over 500 random networks similar to Fig. 3 with $T = 3000$ observed time samples and the local test. **B:** Influence of the number S of surrogates on the detection performance for the 4 surrogate methods (x-axis). The violin plots indicate the distribution of false-alarm rates (left panel) and miss rates (right) for the 500 networks in A with a desired false-alarm rate set to 2% (dashed line in the left panel). Lighter to darker colors correspond to 50, 100, 200 and 400 surrogates, respectively. **C:** Same as the miss rate in B, but only for self connections. **D:** Influence of the number T of observed time samples on the miss rate for $S = 400$ surrogates.

with a density of 20% actually corresponds to $N^2 0.2/100 \simeq 10$ existing connections here, so the plotted improvements concern about 50 connections.

3.5 Applicability to second-order MVAR processes

As explained in Methods, an MVAR process whose state depends on the two previous time steps can be estimated with the covariances with time shifts $\tau = 0, 1$ and 2 ; see Eqs. (11) and (12) for details. Here we simply focus on random connectivity for the two corresponding matrices A^1 and A^2 , with size N that is randomly drawn between 30 and 80; we construct A^1 and A^2 such that a connection $j \rightarrow i$ cannot be in both matrices, but at most in one. The existing connections are detected with the non-parametric local test relying on RP surrogates for each matrix separately, as a proof of concept. The control of false alarms in Fig. 7A and the overall detection performance in Fig. 7B suggest that our surrogate method can be extended satisfactorily to higher-order MVAR processes. Note that the improvement by generating more surrogates is rather weak here. Importantly, there is no difference between the detection in A^1 and A^2 , as demonstrated in Fig. 7C. Last, the network size worsens the miss rate in Fig. 7D, which affects more dramatically the global test as compared to the local test.

4 Application to experimental data

4.1 Multiunit activity data obtained from Utah electrode array in monkey

Now we consider data recorded from a monkey performing a visual task, where the stimulus corresponds to vertical gratings covering all recorded V1 receptive fields from the Utah arrays (see Methods for details). We aim to provide a proof of concept for the connectivity analysis for this type of data, so as to complement the more classical analysis based on the activity of individual channels; therefore we do not focus on comparing the 4 stimulus conditions or additional details.

The multiunit activity envelope (MUAe) is obtained as described in Methods. In Fig. 8A, the resulting MUAe is represented for two out of the 26 channels (red and purple) for two trials in the top and middle panels, 400 ms before and 600 ms after the stimulus onset. The typical analysis of MUAe activity consists of averaging over 200 trials, which exhibits a peak immediately after the stimulus for the two channels in the bottom panel. Among the 26 channels, about a third show a large increase in activity after the stimulus onset as compared to before (namely, a post-stimulus mean activity larger by more than three standard deviations compared to the pre-stimulus activity); almost all channels show a moderate increase of one standard deviation. One channel is discarded for a much larger activity (by 5 times) than all others.

To investigate further the temporal information conveyed by MUAe jointly for pairs of channels, we calculate the pairwise covariances between them, after centering the MUAe activity individually for each trial. We firstly verify that the model can be applied to these data, by examining the MUAe autocovariances in Fig. 8B, which exhibit a profile corresponding to an exponential decay up to two time shifts (i.e., 8 ms for the downsampling every 4 ms), namely a straight line in the log plot. This suits an autoregressive model with large positive values on the diagonal of the connectivity matrix A , which generates this type of profiles.

The two connectivity matrices for the 25 channels estimated using the MVAR method before and after the stimulus are illustrated in Fig. 8C for condition 1: we find larger off-diagonal values for the period after the stimulus than before. This is actually true for all conditions, as indicated by the more spread distributions in red as compared to gray in

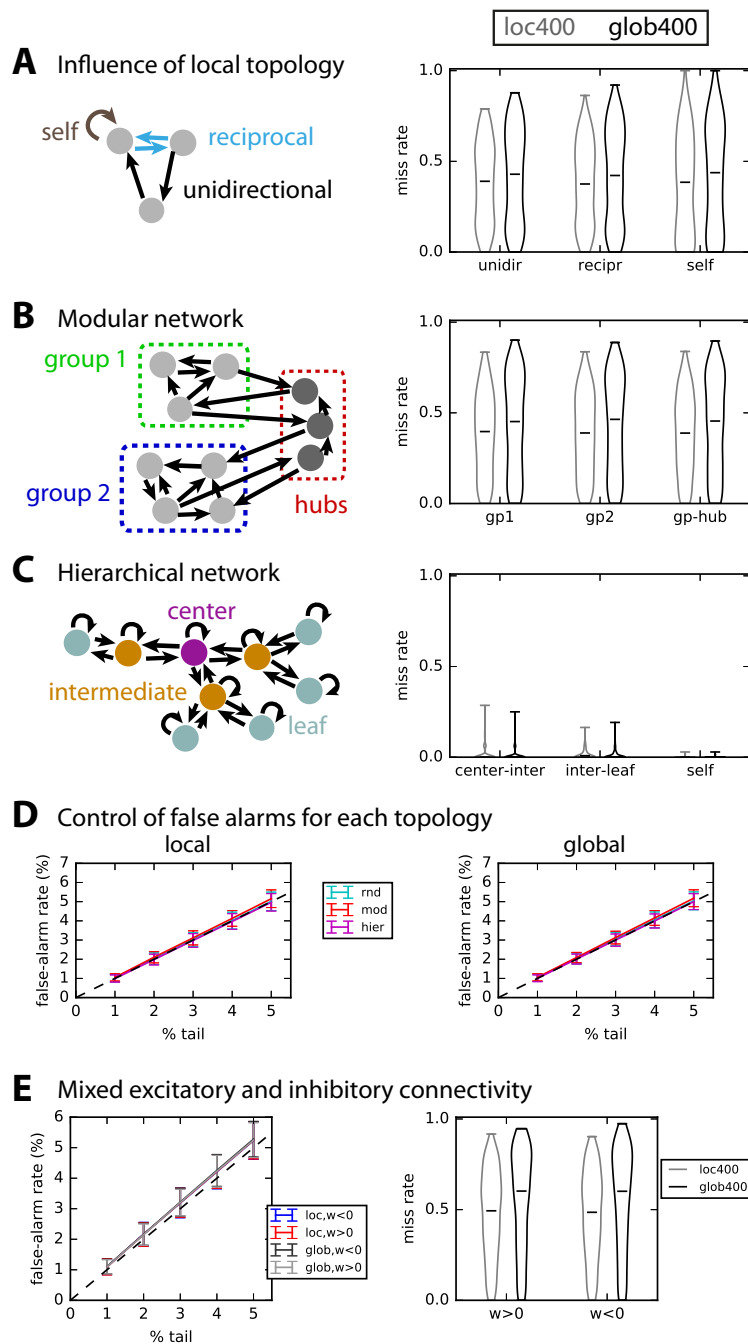


Figure 6: **Robustness to network topology.** **A:** Detection performance for unidirectional, reciprocal and self connections in 500 the randomly connected networks used so far in Fig. 3. **B:** Detection performance for modular topology schematically represented in the left: each of the 500 networks comprises of two groups connected by hubs. The connections are separated depending on whether the connect group nodes or hubs, as indicated by the diagram on the left. **C:** Similar to B with a hierarchical topology, where connections are grouped in 3 subsets: from center to intermediate nodes; from intermediate nodes to leaves; self connections. The results concern 500 such networks, which all have very low density. **D:** Control of false-alarm rate for the local (left) and global (right) significance tests and the three network topologies; the plot is similar to Fig. 3E. **E:** Control of false-alarm rate and miss rate for networks with both excitatory and inhibitory connections.

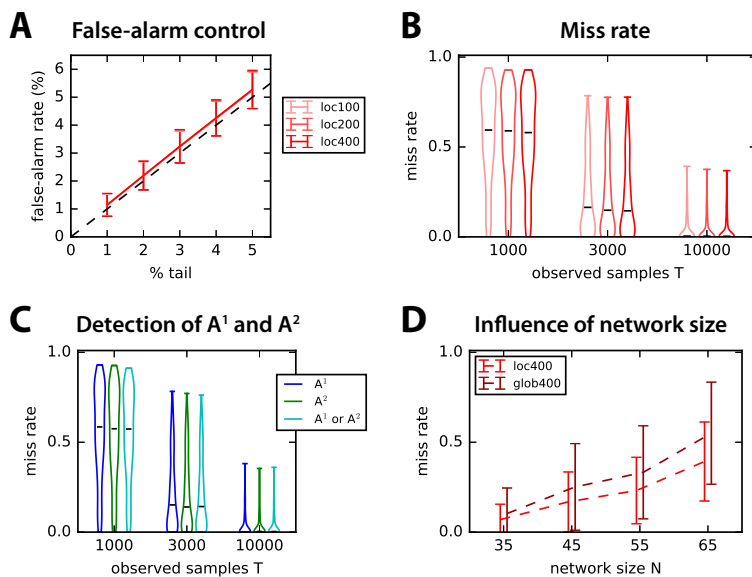


Figure 7: **Connectivity detection for second-order MVAR process.** **A:** Control of false-alarm rate for the local test with $S = 100, 200$ and 400 in both connectivity matrices A^1 and A^2 , corresponding to each time step. Error bars correspond to the variability over 500 network configurations with random connectivity and $T = 3000$ observed samples. **B:** Influence of the number T of observed samples (x-axis) on the miss rate for the 500 networks in A. Lighter to darker red indicates the number of surrogates S . **C:** Details of the detection performance for A^1 and A^2 separately, as well as connections in either A^1 or A^2 . The number of observed samples is indicated on the x-axis as in B. **D:** Influence of the network size N on the detection performance in C for the local and global tests with $S = 400$ surrogates.

Fig. 8D. The channels appear to be coordinated at the considered time scale of 4 ms and their collective interaction scheme is affected by the stimulus presentation.

4.2 Significance test for real data: interactions related to stimulus presentation

We then use the local and global tests based on 1000 surrogates (with random permutation) to retain only significant interactions from the estimates in Fig. 8C: this leaves a few interactions for the pre-stimulus period in Fig. 9A (left panel), 8 out of 650, which is of the order of the desired false-alarm rate set to 1% (namely, the extreme 0.5% of each tail). In contrast, many more post-stimulus interactions survive the significance tests in the right panel: interestingly, almost all these connections are unidirectional. The counterpart for circular shift for Fig. 9B involve 24 connections in common with Fig. 9A. Over all conditions, common post-stimulus interactions number 22 to be compared with 7 for the pre-stimulus period (both with a standard deviation of 4); this corresponds to 3.5% of all possible connections. Almost all detected interactions are unidirectional, as illustrated in Fig. 9C for both local and global tests for the post-stimulus period. Varying the threshold on the tail of the null distributions, we see that the number of detected interactions is close to the desired false-alarm rate for the pre-stimulus period in Fig. 9D (dark red and black curves, respectively). In contrast, post-stimulus interactions are many more for both local and global tests (light red and grey). The global test detects fewer interactions than the local test, indicating the necessity to take into account the disparities across channels. Around 57% of post-stimulus interactions detected by the global test (largest values in absolute value) are found by the local test.

Finally, we check the relationship between the strengths of significant interactions and the increase of average MUAe observed in Fig. 8A. In Fig. 9E, the plotted dots correspond to the sum of incoming (left panel) and outgoing (right panel) significant interactions for each node in the post-stimulus period. The summed interaction values positively correlate with the MUAe difference (post minus pre) only for the outgoing connections: $p < 10^{-2}$ with a coefficient of 0.28. In contrast, incoming connections exhibit a negative correlation, but with $p > 0.1$. This suggests a gain increase of the neuronal populations due to the stimulus, which may drive propagating activity through their anatomical connections and induce effective interactions between channels. If the detected interactions were a result of a noisy detection test, there would likely be no difference between incoming and outgoing interactions, nor significant relationship. Our results thus suggest a non-trivial connectivity that does not reflect observation noise.

5 Discussion

5.1 Beyond Granger causality analysis: properly integrating the network feedback in the connectivity detection

In this work we analyzed the advantage of detecting directional connections in MVAR-like models by using the OLS autoregressive coefficients (instead of the error residuals ratios) and performing non-parametric significance tests based on time-series randomizations. To our understanding the improvement of the method presented here is due to several facts. First, the detection in Fig. 2 - with a common sliding threshold for all connections - is more robust for the regression coefficients than Granger causality analysis based on residual log ratios for recurrent networks with relatively large density (0.1 – 0.3%), which may imply many redundancy and convergence patterns of connections [13, 44]. Second, MVAR-based detection performs better than conditional Granger causality analysis with either the parametric F statistics or non-parametric test based on circular shifts of the observed time series (Fig. 4B). Finally, the use of connection-specific significance testing achieves higher accuracy than para-

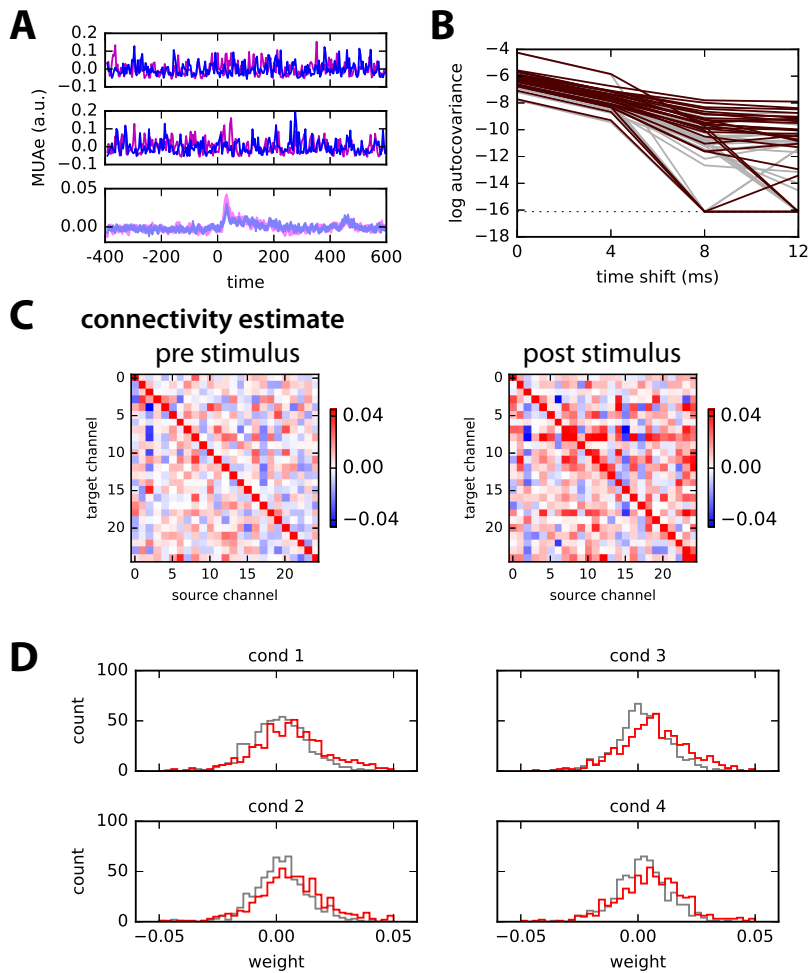


Figure 8: **Application to multiunit activity (MUAe) data.** **A:** Example of two trials (top and middle panels) of multiunit activity envelope (MUAe) for two channels of recordings using Utah electrode array in the primary visual cortex of a monkey (in arbitrary units; see text for further details). The bottom panel represents the average over 200 trials (with standard-error mean for the thickness of the curve). The stimulus is presented at time point 100 (actually 400 ms, since the smoothing corresponds to a smoothing window of 4 ms). **B:** Autocovariances of MUAe signals for time shifts up to 12 ms averaged over 200 trials plotted with a log y-axis: comparison of signals before (gray) and after (red) the stimulus presentation. **C:** MVAR estimates of the connectivity between the 25 channels for the MUAe activity 200 ms before and after stimulus presentation (i.e., 50 time points each), averaged over 200 trials. The scaling has been optimized to enhance the legibility off-diagonal elements. **D:** Comparison of the distribution of connectivity weights (off-diagonal elements in D) for the 4 conditions. Gray and red indicate before and after the stimulus, respectively.

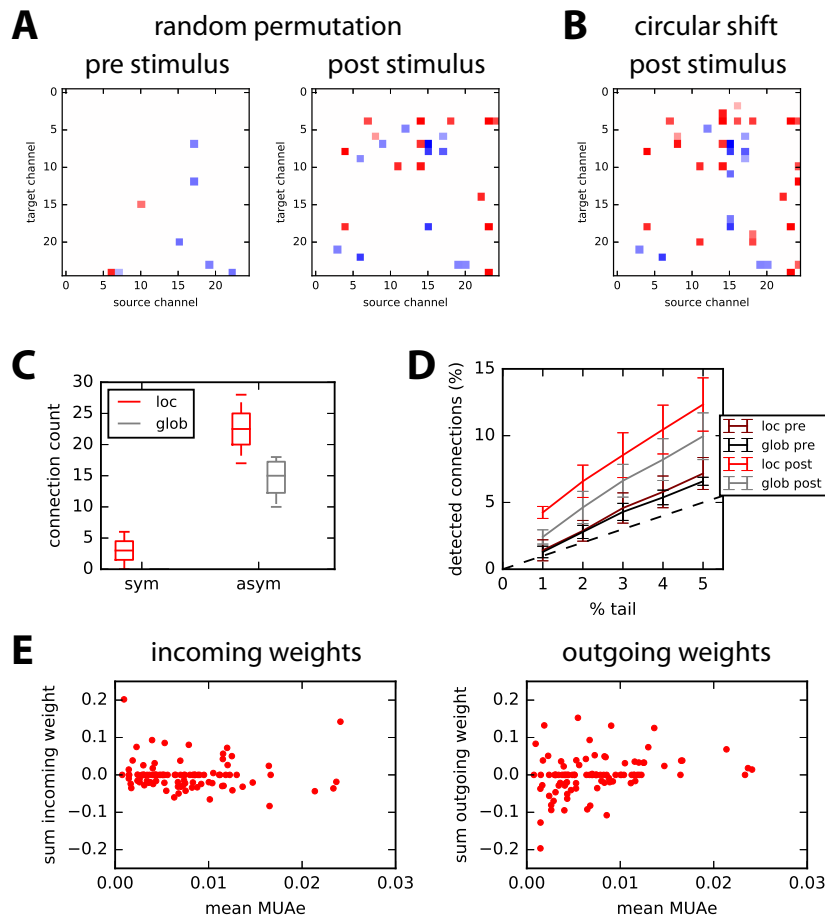


Figure 9: **Detection of significant connections.** **A:** Examples of significant connections with the time-rolling surrogate method; left and right panels correspond to pre- and post-stimulus periods, respectively. The p-value corresponds to the upper and lower 0.5% tails of 1000 surrogates (local test). Many more connections are found for post- than pre-stimulus period. **B:** Matrix of post-stimulus interactions where green pixels indicate symmetric and purple asymmetric significant connections; other are left blank. **C:** Comparison of number of asymmetric connections versus symmetric connections for the local (red) and global (gray) tests over the 4 conditions. **D:** Ratios of detected connections for pre- and post-stimulus periods with the desired false-alarm rate equal to 1 to 5% corresponding to both local and global tests. **E:** Sum of incoming (left) and outgoing (right) significant weights for each channel (dots) for the pre- and post-stimulus period in gray and red, respectively, plotted against the mean MUAe for the 4 conditions.

metric testing, especially when asymptotic assumptions do not hold (e.g., small number of time samples), as illustrated by the local non-parametric test in Figs. 3F and 4B. Note that performing connectivity estimation in the regression coefficients space brings an additional advantage for network studies: the estimated connectivity weights can be interpreted and compared across the whole network, for example using graph theory [42].

Our approach for generating surrogate distributions can be encompassed in the family of constrained randomization methods [39, 37]. Here we have shown that random permutation provides a good estimation of all types of connections (Fig. 5) in false-alarm and miss rates. Comparatively, the circular-shift method performs as well except for self connections that are not detected at all; this also holds for non-random topologies (results not shown). Hence, preserving the autocovariance structure in the generation of surrogates does not provide a substantial advantage here (Fig. 5B-C). Both methods show a good control for the false-alarm rate in comparison to phase randomization [39] and a control Gaussian approximation over the whole network (STD), which lead to an excess of about 1% of false alarms (i.e., ~ 50 connections for a network of 70 nodes). These results show the importance of choosing a surrogate method adapted to the detection problem. For distinct dynamics governing the nodal activity such as nonlinearities, conclusions may differ and further research along these lines is necessary.

As mentioned earlier, the use of an individual null distribution for each connection (local test) gives better results for the miss rate (by a few %) than lumping together all matrix elements of all surrogates (global test), provided sufficiently many surrogates are generated. When computation time is not an issue (as in the present study, see Fig. 4C), these results support the choice of the local test over the global test to attain between accuracy in the true-positive detection. This may be especially true for specific topologies or networks with both excitatory and inhibitory connections, see Fig. 6. In other words, the local test incorporates to a better extent the network heterogeneities in order to build the null distribution for each connection. The present study was limited to ordinary least-square (OLS) estimates for MVAR, but there exist alternative estimators such as the locally weighted least-square regression [36] that may perform better for particular network topologies. The extension of the presented surrogate techniques to the case where observations are sparser than connections - implying that the covariance matrix is not invertible - is another interesting direction to explore [9].

The problem of multiple comparison is intrinsic to brain connectivity detection as the number of testable connections across brain regions is massive [35]. In this context, different approaches have been developed to control the family-wise error rate in the weak sense. For instance, many studies on neuroimaging data [17, 33] or electrophysiology [24] have resorted to procedures that control the false discovery rate (FDR) [7], namely, the expected number of falsely declared connections among the total number of detections. These methods make decisions on single connections relying on the entire sequence of p-values computed for each connection and yield substantial statistical power gains over more conservative methods such as Šidák-Bonferroni [1]. With the ever growing application of graph theory to brain connectivity, new methods have been proposed that exploit the clustered structure of the the declared connections [53, 21] to propose cluster-based statistical tests [28] that attain similar performance to FDR methods. The present work can therefore be understood as a primary step before performing any or several multiple-correction procedures. By defining an accurate null model of inexistent connections, p-value estimates per connection are improved and cluster-based surrogate distributions can be better approximated, which is expected to empower the overall control of false positive rates in network connectivity analysis.

5.2 Applications to real electrophysiological and neuroimaging data

Here we have applied the method to MUAe recorded from macaque area V1. Our point was to verify the applicability of our network connectivity analysis to this type of data rather

than addressing biological questions, which is beyond the scope of the current manuscript. Looking at individual trials Fig. 8A, it is rather surprising that the variable MUAe activity convey temporal information due to causal interactions that can be detected. This variability is therefore not an absolute limitation to temporal coordination, although the latter only becomes apparent over multiple trial repetitions. The research of optimal preprocessing - in particular the filtering to obtain MUAe - to obtain a robust detection of interactions is left for a later study.

More generally, we target adequate preprocessing of multivariate time series - activity aggregation over hundreds of voxels for fMRI and MUAe for electrode recordings - such that the autocovariance profiles match the exponential decay of the dynamic model with linear feedback that underlies the connectivity analysis: provided it resembles an exponential decay over a few time shifts as in Fig. 8B, we expect our approach to be applicable. In theory, stationarity of the time series remains a critical issue as we need sufficiently many observed samples to obtain precise covariances from which we evaluate the connectivity estimates. On the other hand, despite the non-stationary nature of the MUAe signals illustrated in Fig. 8A, we obtain significant interactions that preferentially arise for channels exhibiting an increase of MUAe, but are not merely a consequence of it: the method provides information about the directionality of the interactions in Fig. 9B to E.

Acknowledgments

MG acknowledges funding from the Marie Skłodowska-Curie Action (grant H2020-MSCA-656547). MG and GD were supported by the Human Brain Project (grant FP7-FET-ICT-604102 and H2020-720270 HBP SGA1). GD and ATC were supported by the European Research Council Advanced Grant DYSTRUCTURE 588 (Grant 295129). The authors are grateful to Robert Castelo and Inma Tur for constructive discussions.

References

- [1] H. Abdi. The bonferonni and šidák corrections for multiple comparisons. *Encyclopedia of measurement and statistics*, 3:103–107, 2007.
- [2] T. Amemiya. Multivariate regression and simultaneous equation models when the dependent variables are truncated normal. *Econometrica: Journal of the Econometric Society*, pages 999–1012, 1974.
- [3] F. Babiloni, F. Cincotti, C. Babiloni, F. Carducci, D. Mattia, L. Astolfi, A. Basilisco, P. Rossini, L. Ding, Y. Ni, J. Cheng, K. Christine, J. Sweeney, and B. He. Estimation of the cortical functional connectivity with the multimodal integration of high-resolution eeg and fmri data by directed transfer function. *Neuroimage*, 24(1):118–131, 2005.
- [4] L. Barnett and A. Seth. The MVGC multivariate granger causality toolbox: A new approach to granger-causal inference. *J Neurosci Methods*, 223:50–68, 2014.
- [5] L. Barnett and A. Seth. Granger causality for state-space models. *Phys Rev E*, 91:040101, 2015.
- [6] A. Barrett and L. Barnett. Granger causality is designed to measure effect, not mechanism. *Front Neuroinform*, 7:6, 2013.
- [7] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300, 1995.

- [8] N. Bing and I. Hoeschele. Genetical genomics analysis of a yeast segregant population for transcription network inference. *Genetics*, 170(2):533–542, 2005.
- [9] R. Castelo and A. Roverato. A robust procedure for Gaussian graphical model search from microarray data with p larger than n . *The Journal of Machine Learning Research*, 7:2621–2650, 2006.
- [10] M. Dhamala, G. Rangarajan, and M. Ding. Analyzing information flow in brain networks with nonparametric granger causality. *Neuroimage*, 41(2):354–362, 2008.
- [11] C. Diks and J. DeGoede. A general nonparametric bootstrap test for granger causality. pages 391–403. Bristol, UK: Institute of Physics Publishing, 2001.
- [12] L. Ding, G. Worrell, T. Lagerlund, and B. He. Ictal source analysis: localization and imaging of causal interactions in humans. *Neuroimage*, 34(2):575–586, 2007.
- [13] M. Ding, Y. Chen, and S. Bressler. Granger causality: basic theory and application to neuroscience. In B. Schelter, M. Winterhalder, and J. Timmer, editors, *Handbook of Time Series Analysis: Recent Theoretical Developments and Applications*, ISBN: 978-3-527-40623-4. Wiley-VCH Verlage, arXiv preprint q-bio/0608035, 2006.
- [14] L. Faes, D. Marinazzo, A. Montalto, and G. Nollo. Lag-specific transfer entropy as a tool to assess cardiovascular and cardiorespiratory information transfer. *IEEE Trans Biomed Eng*, 61:2556–68, 2014.
- [15] K. Friston, A. Bastos, A. Oswal, B. van Wijk, C. Richter, and V. Litvak. Granger causality revisited. *Neuroimage*, 101:796–808, 2014.
- [16] K. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302, 2003.
- [17] C. Genovese, N. Lazar, and T. Nichols. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15(4):870–878, 2002.
- [18] J. Geweke. Measurement of linear dependence and feedback between multiple time series. *J Am Stat Assoc*, 77:304–313, 1982.
- [19] J. Geweke. Measures of conditional linear dependence and feedback between time series. *J Am Stat Assoc*, 79:907–915, 1984.
- [20] C. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- [21] C. Han, S. Yoo, S. Seo, D. Na, and J.-K. Seong. Cluster-based statistics for brain connectivity in correlation with behavioral measures. *PLoS one*, 8(8):e72332, 2013.
- [22] M. Kaminski and K. Blinowska. A new method of the description of the information flow in the brain structures. *Biological cybernetics*, 65(3):203–210, 1991.
- [23] M. Kamiński, M. Ding, W. Truccolo, and S. Bressler. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. *Biological cybernetics*, 85(2):145–157, 2001.
- [24] A. Lage-Castellanos, E. Martínez-Montes, J. Hernández-Cabrera, and L. Galán. False discovery rate and permutation test: an evaluation in erp data analysis. *Statistics in medicine*, 29(1):63–74, 2010.
- [25] Y. Li, X. Ye, Q. Liu, J. Mao, P. Liang, J. Xu, and P. Zhang. Localization of epileptogenic zone based on graph analysis of stereo-eeg. *Epilepsy Res*, 128:149–157, 2016.

- [26] B. Lusch, P. Maia, and J. Kutz. Physical review e covering statistical, nonlinear, biological, and soft matter physics highlights recent accepted authors referees search press about inferring connectivity in networked dynamical systems: Challenges using granger causality. *Phys Rev E*, 94:032220, 2016.
- [27] H. Lütkepohl. *New introduction to multiple time series analysis*. Springer Science & Business Media, 2005.
- [28] E. Maris and R. Oostenveld. Nonparametric statistical testing of eeg-and meg-data. *Journal of neuroscience methods*, 164(1):177–190, 2007.
- [29] J. Massey. Causality, feedback and directed information. In *Proc. Int. Symp. Inf. Theory Applic. (ISITA-90)*, pages 303–305. Citeseer, 1990.
- [30] A. Messé, D. Rudrauf, H. Benali, and G. Marrelec. Relating structure and function in the human brain: relative contributions of anatomy, stationary dynamics, and non-stationarities. *PLoS Comput Biol*, 10:e1003530, 2014.
- [31] G. Michalareas, J. Schoffelen, G. Paterson, and J. Gross. Investigating causality between interacting brain areas with multivariate autoregressive models of meg sensor data. *Human Brain Mapping*, 34(4):890–913, 2013.
- [32] A. Nedungadi, G. Rangarajan, N. Jain, and M. Ding. Analyzing multiple spike trains with nonparametric granger causality. *J Comp Neu*, 27(1):55–64, 2009.
- [33] T. Nichols and S. Hayasaka. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Statistical methods in medical research*, 12(5):419–446, 2003.
- [34] B. Rogers, S. Katwal, V. Morgan, C. Asplund, and J. Gore. Functional mri and multivariate autoregressive models. *Magnetic Resonance Imaging*, 28(8):1058–1065, 2016/11/05 2010.
- [35] M. Rubinov and O. Sporns. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.
- [36] D. Ruppert and M. Wand. Multivariate locally weighted least squares regression. *Ann Statist*, 22:1346–1370, 1994.
- [37] T. Schreiber. Constrained randomization of time series data. *Phys Rev Lett*, 80:2105, 1998.
- [38] T. Schreiber. Measuring information transfer. *Phys Rev Lett*, 85:461, 2000.
- [39] T. Schreiber and A. Schmitz. Improved surrogate data for nonlinearity tests. *Phys Rev Lett*, 77:635, 1996.
- [40] A. Seth, A. Barrett, and L. Barnett. Granger causality analysis in neuroscience and neuroimaging. *J Neurosci*, 35:3293–3297, 2015.
- [41] K. So, A. Koralek, K. Ganguly, M. Gastpar, and J. Carmena. Assessing functional connectivity of neural ensembles using directed information. *J Neu Eng*, 9(2):026004, 2012.
- [42] O. Sporns. Making sense of brain network data. *Nat Methods*, 10:491–493, 2013.
- [43] A. Storkey, E. Simonotto, H. Whalley, S. Lawrie, L. Murray, and D. Mcgonigle. Learning structural equation models for fmri. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1329–1336. MIT Press, 2007.

- [44] S. Stramaglia, J. Cortes, and D. Marinazzo. Synergy and redundancy in the granger causal analysis of dynamical networks. *New Journal of Physics*, 16:105003, 2014.
- [45] H. Supèr and P. Roelfsema. Chronic multiunit recordings in behaving animals: advantages and limitations. *Prog Brain Res*, 147:263–282, 2005.
- [46] A. Tauste Campo, M. Martinez-Garcia, V. Nácher, R. Luna, R. Romo, and G. Deco. Task-driven intra-and interarea communications in primate cerebral cortex. *Proc Natl Acad Sci U.S.A.*, 112(15):4761–4766, 2015.
- [47] A. Thiele, L. Delicato, M. Roberts, and M. Gieselmann. A novel electrode-pipette design for simultaneous recording of extracellular spikes and iontophoretic drug application in awake behaving monkeys. *J Neurosci Methods*, 158:207–211, 2006.
- [48] M. Vinck, L. Huurdeman, C. Bosman, P. Fries, F. Battaglia, C. Pennartz, and P. Tiesinga. How to detect the granger-causal flow direction in the presence of additive noise? *Neuroimage*, 108:301–18, 2015.
- [49] C. Wilke, L. Ding, and B. He. Estimation of time-varying connectivity patterns through the use of an adaptive directed transfer function. *IEEE Trans Biomed Eng*, 55(11):2557–2564, 2008.
- [50] A. Winkler, G. Ridgway, M. Webster, S. Smith, and T. Nichols. Permutation inference for the general linear model. *Neuroimage*, 92:381–97, 2014.
- [51] www.scipy.org. the scientific python library.
- [52] Y. Yuan, C. Li, and O. Windram. Directed Partial Correlation: Inferring Large-Scale Gene Regulatory Network through Induced Topology Disruptions. *PLoS ONE*, 6(4), 2011.
- [53] A. Zalesky, A. Fornito, and E. Bullmore. Network-based statistic: identifying differences in brain networks. *Neuroimage*, 53(4):1197–1207, 2010.