



# A MATLAB toolbox for Granger causal connectivity analysis

Anil K. Seth\*

Sackler Centre for Consciousness Science and School of Informatics, University of Sussex, Brighton, BN1 9QJ, UK

## ARTICLE INFO

### Article history:

Received 23 September 2009  
Received in revised form  
21 November 2009  
Accepted 23 November 2009

### Keywords:

MATLAB  
Granger causality  
Toolbox  
Network theory  
Causal density

## ABSTRACT

Assessing directed functional connectivity from time series data is a key challenge in neuroscience. One approach to this problem leverages a combination of Granger causality analysis and network theory. This article describes a freely available MATLAB toolbox – ‘Granger causal connectivity analysis’ (GCCA) – which provides a core set of methods for performing this analysis on a variety of neuroscience data types including neuroelectric, neuromagnetic, functional MRI, and other neural signals. The toolbox includes core functions for Granger causality analysis of multivariate steady-state and event-related data, functions to preprocess data, assess statistical significance and validate results, and to compute and display network-level indices of causal connectivity including ‘causal density’ and ‘causal flow’. The toolbox is deliberately small, enabling its easy assimilation into the repertoire of researchers. It is however readily extensible given proficiency with the MATLAB language.

© 2009 Elsevier B.V. All rights reserved.

## 1. Introduction

The accelerating availability of neuroscience data is placing increasing demands on analysis methods. These data are generated at multiple levels of description, from spike trains to local field potentials to functional MRI (fMRI) BOLD signals. A key challenge when analyzing such data is to determine the functional connectivity of the underlying mechanisms. Many methods for functional connectivity analysis identify *undirected* connectivity; examples include synchrony (Engel and Singer, 2001) and phase coherence (Nunez et al., 2001; Doesburg et al., 2009). However, a satisfactory understanding of neural mechanisms is likely to require identification of *directed* functional connectivity. A powerful technique for extracting such connectivity from data is *Granger causality* (G-causality) (Granger, 1969; Ding et al., 2006; Seth, 2007). According to G-causality, a variable  $X_1$  ‘Granger causes’ a variable  $X_2$  if information in the past of  $X_1$  helps predict the future of  $X_2$  with better accuracy than is possible when considering only information in the past of  $X_2$  itself (Granger, 1969; Seth, 2007). This paper describes a freely available software toolbox, programmed in the MATLAB (Natick, MA) environment, which allows application of a range of G-causality analyses to neuroscience data broadly construed. Taken together, the analysis methods incorporated in the toolbox form ‘Granger causal connectivity analysis’ (GCCA). A first version of the toolbox was released in 2005 (Seth, 2005); the present paper describes the first significant revision and extension of the software,

as well as the theoretical infrastructure underpinning its functionality.

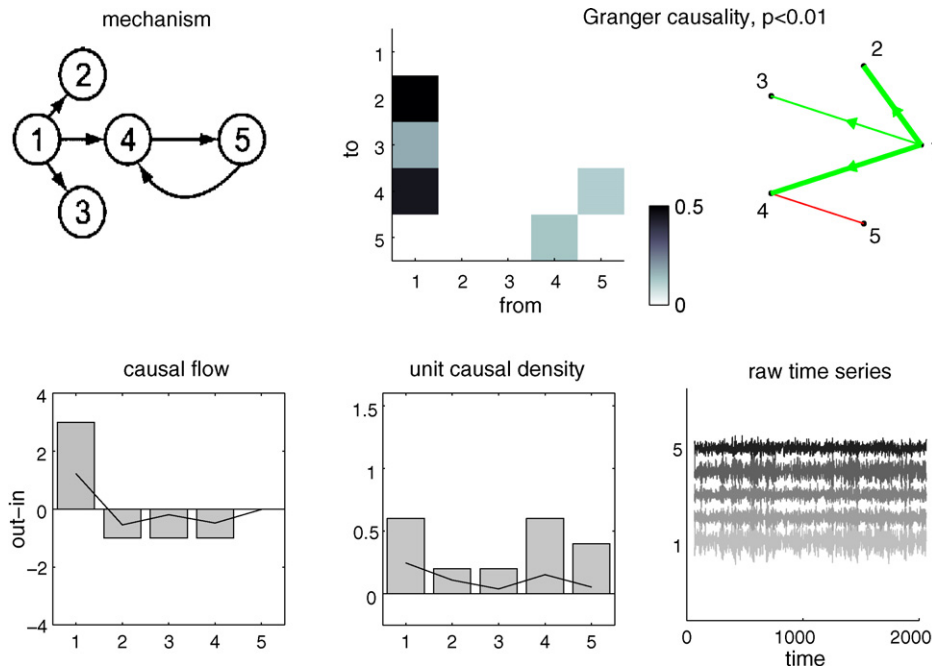
It should be noted that identifying directed functional connectivity is *not* equivalent to identifying physically instantiated causal interactions in systems. Although the two descriptions are intimately related (Seth and Edelman, 2007; Cadotte et al., 2008), physically instantiated causal structure can only be unambiguously identified by perturbing a system and observing the consequences (Pearl, 1999). Directed functional connectivity in general, and G-causality analysis in particular, is therefore best understood as a statistical relationship among observed variables that reflects but may not be identical to the underlying physical mechanism.

The importance of identifying causal structure within data, especially during exploratory analysis phases, indicates a need for easy-to-apply, transparent, and extensible software methods. Such methods are provided by the GCCA toolbox described here. The toolbox includes several different types of function. The core functions implement G-causality analysis given multivariate time series data. Other functions test whether the provided data satisfies necessary assumptions, assess the statistical significance and validity of inferred interactions, generate network-level descriptions of patterns of causal interactions, and graphically display analysis results. Functions are also included to apply various preprocessing techniques and to demonstrate the toolbox capabilities. The toolbox is intentionally small when compared to several other brain signal analysis toolboxes (e.g., Delorme and Makeig, 2004; Cui et al., 2008), in order to facilitate assimilation of the methods into the repertoire of researchers. Also, the toolbox is not targeted specifically to any particular experimental technique. G-causality analysis has been successfully applied to a wide range of data types including spike trains (Cadotte et al., 2008; Nedungadi et al., 2009), local

\* Tel.: +44 1273 678549; fax: +44 1273 877873.

E-mail address: [a.k.seth@sussex.ac.uk](mailto:a.k.seth@sussex.ac.uk).

URL: <http://www.anilseth.com>.



**Fig. 1.** Demonstration of the GCCA toolbox. The data-generating process is shown in the top-left panel. Matrix and network representations of the corresponding G-causalities are shown in the remaining top panels (green lines depict unidirectional connections, red lines depict bidirectional connections). The bottom panels show (left-to-right) causal flow, unit causal density, and the raw time series generated by the process. Causal flow and causal density are network-level descriptions of causal patterns which are described further in Section 4; both ‘unweighted’ (bars) and ‘weighted’ (lines) versions are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

field potentials (Kaminski et al., 2001; Brovelli et al., 2004), EEG signals (Babiloni et al., 2005; Keil et al., 2009), fMRI BOLD signals (Roebroeck et al., 2005; Sato et al., 2006), among others. Several such studies have utilized an earlier version of the present toolbox (Sridharan et al., 2008; Shanahan, 2008; Gow et al., 2008; Gow and Segawa, 2009; Gaillard et al., 2009; Chang et al., 2008; Stevens et al., 2009) or have followed the procedures it advocated (Keil et al., 2009). By November 2009, over 600 researchers from many different countries had downloaded the toolbox. The present toolbox extends the earlier version in several important ways, including (i) frequency decomposition of causal interactions, (ii) partial Granger causality, (iii) ‘Granger autonomy’, (iv) analysis of event-related data, (v) additional preprocessing and validation techniques, for example for removal of single frequency line-noise from data and for checking model consistency, and (vi) built-in functions for bootstrap and permutation statistical testing. The toolbox is also enhanced with respect to the efficiency of its core algorithms.

A screenshot of the toolbox demonstration function is provided in Fig. 1. In this figure, data are generated according to the model described in Baccalá and Sameshima (2001) (Eq. (1.1), see top-left panel).

$$\begin{aligned}
 x_1(t) &= 0.95\sqrt{2}x_1(t-1) - 0.9025x_1(t-2) + w_1(t) \\
 x_2(t) &= 0.5x_1(t-2) + w_2(t) \\
 x_3(t) &= -0.4x_1(t-3) + w_3(t) \\
 x_4(t) &= -0.5x_1(t-2) + 0.25\sqrt{2}x_4(t-1) + 0.25\sqrt{2}x_5(t-1) + w_4(t) \\
 x_5(t) &= -0.25\sqrt{2}x_4(t-1) + 0.25\sqrt{2}x_5(t-1) + w_5(t)
 \end{aligned}
 \quad (1.1)$$

The model generates five time-series which are shown in the bottom-right panel of the figure.  $w_1 - w_5$  are independent normally-distributed processes with zero mean and unit standard deviation. The remaining panels show the results of GCCA of the time-series. The top panels show G-causalities in both matrix and network form. In this case the inferred causalities are identical to the underlying physical network structure (as noted this need not be true in general). The bottom panels show network-level

summary descriptions of the causal patterns, described further in Section 4. Briefly, causal flow reflects the extent to which a variable is influenced by or influences the remainder of the system; causal density expresses the overall degree of causal interactivity, either averaged across an entire network or assessed on a per-variable basis.

The remainder of this paper is organized as follows. Section 2 describes the underlying theory of G-causality and of the various extensions represented in the toolbox. Section 3 discusses the statistical assumptions involved in G-causality analysis and describes various preprocessing and validation techniques. Section 4 introduces some network-level descriptors of G-causality interaction patterns which can be useful for inferring macroscopic dynamical properties of neural systems. Section 5 discusses the important issue of filtering. Section 6 then outlines modality-specific issues involved in application of G-causality to data acquired from functional MRI (fMRI), neuroelectric and neuromagnetic signals, and spike train data. Limitations of the present instantiation of the toolbox are noted in Section 7, and Section 8 sets G-causality in the context of other frameworks for identifying directed functional connectivity. For detailed descriptions of the individual functions in the toolbox the reader is referred to the accompanying manual, available online at [www.anilseth.com](http://www.anilseth.com).

## 2. Granger causality

### 2.1. Bivariate and conditional G-causality

In 1969 Granger introduced the idea of G-causality as a formalization, in terms of linear regression modelling, of Wiener’s (and Akaike’s) intuition that  $X_2$  ‘causes’  $X_1$  if knowing  $X_2$  helps predict the future of  $X_1$  (Granger, 1969; Seth, 2007). According to G-causality,  $X_2$  causes  $X_1$  if the inclusion of past observations of  $X_2$  reduces the prediction error of  $X_1$  in a linear regression model of  $X_1$  and  $X_2$ , as compared to a model which includes only previous observations of  $X_1$ . To illustrate G-causality, suppose that the tem-

poral dynamics of two time series  $X_1(t)$  and  $X_2(t)$  (both of length  $T$ ) can be described by a bivariate autoregressive model:

$$\begin{aligned} X_1(t) &= \sum_{j=1}^p A_{11,j} X_1(t-j) + \sum_{j=1}^p A_{12,j} X_2(t-j) + \xi_1(t) \\ X_2(t) &= \sum_{j=1}^p A_{21,j} X_1(t-j) + \sum_{j=1}^p A_{22,j} X_2(t-j) + \xi_2(t) \end{aligned} \quad (2.1)$$

where  $p$  is the maximum number of lagged observations included in the model (the model order,  $p < T$ ),  $A$  contains the coefficients of the model, and  $\xi_1, \xi_2$  are the residuals (prediction errors) for each time series. If the variance of  $\xi_1$  (or  $\xi_2$ ) is reduced by the inclusion of the  $X_2$  (or  $X_1$ ) terms in the first (or second) equation, then it is said that  $X_2$  (or  $X_1$ ) *G-causes*  $X_1$  (or  $X_2$ ). Assuming that  $X_1$  and  $X_2$  are covariance stationarity (i.e., unchanging mean and variance), the magnitude of this interaction can be measured by the log ratio of the prediction error variances for the restricted (R) and unrestricted (U) models:

$$\mathcal{F}_{2 \rightarrow 1} = \ln \frac{\text{var}(\xi_{1R(12)})}{\text{var}(\xi_{1U})}, \quad (2.2)$$

where  $\xi_{1R(12)}$  is derived from the model omitting the  $A_{12,j}$  (for all  $j$ ) coefficients in the first equation and  $\xi_{1U}$  is derived from the full model.

Importantly, G-causality is easy to generalize to the multivariate (conditional) case in which the G-causality of  $X_2$  on  $X_1$  is tested in the context of multiple additional variables  $X_3 \dots X_n$  (Geweke, 1982). In this case,  $X_2$  G-causes  $X_1$  if knowing  $X_2$  reduces the variance in  $X_1$ 's prediction error when all other variables  $X_3 \dots X_n$  are also included in the regression model. To illustrate, for a system of three variables (1,2,3), we represent the noise covariance matrix of the unrestricted model as

$$\Sigma = \begin{bmatrix} \text{var}(\xi_{1U}) & \text{cov}(\xi_{1U} \xi_{2U}) & \text{cov}(\xi_{1U} \xi_{3U}) \\ \text{cov}(\xi_{2U} \xi_{1U}) & \text{var}(\xi_{2U}) & \text{cov}(\xi_{2U} \xi_{3U}) \\ \text{cov}(\xi_{3U} \xi_{1U}) & \text{cov}(\xi_{3U} \xi_{2U}) & \text{var}(\xi_{3U}) \end{bmatrix} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \Sigma_{13} \\ \Sigma_{21} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{31} & \Sigma_{32} & \Sigma_{33} \end{bmatrix}$$

where all  $\xi_{iU}$  are estimated from the (vector) autoregressive model including all variables. A useful partition of the matrix is given by the second equality. For  $n$  variables there are  $n$  restricted models, with each restricted model omitting a different predictor variable. For example, the noise covariance matrix of the restricted model omitting variable 2, with its partition, is

$$\rho = \begin{bmatrix} \text{var}(\xi_{1R}) & \text{cov}(\xi_{1R} \xi_{3R}) \\ \text{cov}(\xi_{3R} \xi_{1R}) & \text{var}(\xi_{3R}) \end{bmatrix} = \begin{bmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{bmatrix}$$

where all  $\xi_{iR}$  are estimated from the autoregressive model omitting variable 2. The G-causality from variable 2 to variable 1, conditioned on variable 3, is given by

$$\mathcal{F}_{2 \rightarrow 1|3} = \ln \frac{\rho_{11}}{\Sigma_{11}}. \quad (2.3)$$

The above development assumes that the observed data can be well represented by multivariate autoregressive (MVAR) models. The GCCA toolbox uses two different core algorithms to estimate these models. For data consisting of single long trials, the method of ordinary-least-squares is used to compute the regression coefficients ( $A$  in Eq. (2.1)). Alternatively, if the data are in the form of multiple repetitions of relatively short trials (e.g., event-related data), the method of Ding et al. (2000) is used. In this method, each trial is considered to be an independent realization of a single statistically stationary process, such that a single MVAR model can be estimated based on the entire data set. The algorithm enabling this

estimation is due to Morf et al. (1978), its computer implementation is available as part of the BSMART network analysis software (Cui et al., 2008).<sup>1</sup>

### 2.1.1. Model order

The estimation of MVAR models requires as a parameter the number of time-lags ( $p$ ) to include, i.e., the model order. Too few lags can lead to a poor representation of the data, whereas too many can lead to problems of model estimation. A principled means to specify the model order is to minimize a criterion that balances the variance accounted for by the model, against the number of coefficients to be estimated. Two criteria are implemented in the toolbox: the Akaike information criterion (AIC, Akaike, 1974) and the Bayesian information criterion (BIC, Schwartz, 1978). For  $n$  variables:

$$\text{AIC}(p) = \ln(\det(\Sigma)) + \frac{2pn^2}{T} \quad (2.4)$$

$$\text{BIC}(p) = \ln(\det(\Sigma)) + \frac{\ln(T)pn^2}{T} \quad (2.5)$$

The BIC is more often used for application to neural systems because it compensates for the large number of data points commonly found in neural data sets (fMRI excepted). In cases where the model order specified by the minimal BIC/AIC is too large to permit feasible computation, or in cases where the BIC/AIC does not reach a clear minimum over the range tested, a smaller model order can be chosen on condition that the BIC/AIC shows no further substantial decreases at higher orders (Brovelli et al., 2004).

### 2.1.2. Statistical significance

Having computed G-causality magnitudes, it is important to assess their statistical significance. A time-domain G-causality

interaction is significant if the coefficients in the corresponding  $A_{ij}$  are jointly significantly different from zero. This can be established via an  $F$ -test on the null hypothesis that  $A_{ij}$  are zero (Granger, 1969).<sup>2</sup> These tests should be corrected for multiple comparisons. The most conservative correction is the Bonferroni correction, in which the applied threshold is  $P_{\text{nom}}/n(n-1)$ , where  $P_{\text{nom}}$  is the nominal threshold (typically 0.01 or 0.05). A less conservative alternative is the *false discovery rate* (FDR) which controls the expected *proportion* of incorrectly rejected null hypotheses (type I errors) (Benjamini and Hochberg, 1995). (The Bonferroni, by contrast, controls the expected *number* of type I errors.) The GCCA toolbox is able to apply both of these corrections.

The core functions implementing time-domain G-causality analysis in the GCCA toolbox are `cca_granger_regress` for single trial data and `cca_granger_regress_mtrial` for multiple trial data. These functions return G-causality magnitudes and  $F$ -values. The function `cca_findsignificance` filters these values through the multiple comparison corrections just described.

<sup>1</sup> <http://www.brain-smart.org>.

<sup>2</sup> It is also known that the maximum likelihood estimator  $\hat{\mathcal{F}}$  will have (asymptotically for large samples) a  $\chi^2$ -distribution under the null hypothesis  $\mathcal{F}=0$  (Granger, 1963; Whittle, 1953) and a non-central  $\chi^2$ -distribution under the alternative hypothesis  $\mathcal{F} > 0$  (Geweke, 1982; Wald, 1943).

## 2.2. Spectral G-causality

Neural dynamics are often usefully interpreted in the frequency domain. A frequency-domain (spectral) interpretation of G-causality can be derived by examining the Fourier components of an MVAR model estimated in the time-domain (Geweke, 1982). The following treatment is based on the development in Kaminski et al. (2001) and Brovelli et al. (2004) [see also Seth, 2007]. The Fourier transform of (2.1) gives:

$$\begin{pmatrix} A_{11}(f) & A_{12}(f) \\ A_{21}(f) & A_{22}(f) \end{pmatrix} \begin{pmatrix} X_1(f) \\ X_2(f) \end{pmatrix} = \begin{pmatrix} E_1(f) \\ E_2(f) \end{pmatrix} \quad (2.6)$$

in which the components of  $A$  are

$$A_{lm}(f) = \delta_{lm} - \sum_{j=1}^p A_{lm}(j) e^{-i2\pi f j},$$

$$\delta_{lm} = 0 \quad (l = m),$$

$$\delta_{lm} = 1 \quad (l \neq m).$$

Rewriting (2.6) as

$$\begin{pmatrix} X_1(f) \\ X_2(f) \end{pmatrix} = \begin{pmatrix} H_{11}(f) & H_{12}(f) \\ H_{21}(f) & H_{22}(f) \end{pmatrix} \begin{pmatrix} E_1(f) \\ E_2(f) \end{pmatrix}$$

we have

$$\begin{pmatrix} H_{11}(f) & H_{12}(f) \\ H_{21}(f) & H_{22}(f) \end{pmatrix} = \begin{pmatrix} A_{11}(f) & A_{12}(f) \\ A_{21}(f) & A_{22}(f) \end{pmatrix}^{-1}$$

where  $H$  is the transfer matrix. The spectral matrix  $S$  can now be derived as

$$S(f) = \langle X(f)X^*(f) \rangle = \langle H(f)\Sigma H^*(f) \rangle$$

in which the asterisk denotes matrix transposition and complex conjugation and  $\Sigma$  is the noise covariance matrix. The spectral G-causality from  $j$  to  $i$  is then

$$I_{j \rightarrow i}(f) = -\ln \left( 1 - \frac{(\Sigma_{jj} - (\Sigma_{ij}^2 / \Sigma_{ii})) |H_{ij}(f)|^2}{S_{ii}(f)} \right) \quad (2.7)$$

in which  $S_{ii}(f)$  is the power spectrum of variable  $i$  at frequency  $f$  and  $\Sigma$ ,  $H$ , and  $S$  are defined for bivariate models only.

The GCCA toolbox contains the function `cca_pwcausal` which returns pairwise spectral G-causalities and coherence values for a matrix of time series. This function is a wrapper function for `pwcausal.m`, which is part of the BSMART toolbox.<sup>3</sup> It uses the Morf algorithm to estimate the underlying time-domain model such that it can be used both for single long time-series and for multi-trial data. Importantly and in contrast to the time-domain formulation, the statistical distribution of spectral G-causality does not follow any standard distribution entailing that statistical significance values and confidence intervals must be estimated using surrogate data methods (see Section 2.5).

Two alternative measures closely related to spectral G-causality are partial directed coherence (Baccalá and Sameshima, 2001) and the directed transfer function (Kaminski et al., 2001). For comparative results among these methods see Baccalá and Sameshima (2001), Pereda et al. (2005) and Gourévitch et al. (2006).

## 2.3. Partial G-causality

Time-series inference methods such as G-causality can be undermined by the confounding influence of environmental

(exogenous) and unmeasured (latent) inputs. An adaptation of G-causality that addresses this problem has been called *partial G-causality* (Guo et al., 2008). The intuition, based on the concept of partial coherence, is that the influence of exogenous and/or latent variables on a measured system will be reflected by correlations among the residuals of an MVAR model of the measured variables. The partial G-causality is defined as (compare with Eq. (2.3)):

$$\mathcal{F}_{2 \rightarrow 1|3}^p = \ln \left( \frac{\rho_{11} - \rho_{12}\rho_{22}^{-1}\rho_{21}}{\Sigma_{11} - \Sigma_{13}\Sigma_{33}^{-1}\Sigma_{31}} \right) \quad (2.8)$$

Note that partial G-causality cannot be computed for <3 variables because even the restricted model requires residual covariance terms in order to provide an estimate of the influence of exogenous/latent variables. It should also be noted that the analogy with partial coherence is not exact. Whereas partial coherence directly removes the influence of known external variables from a fundamentally bivariate measure (coherence), partial G-causality controls for the influence of unknown variables on a multivariate measure (G-causality) indirectly via their influence on residual covariances. For this reason, partial G-causality will only remove all traces of exogenous/latent variables in the unlikely case that these variables have equivalent effects on all components of the measured system. However, numerical investigations show that even when this condition is not met, partial G-causality nonetheless delivers substantially improved results as compared to standard conditional G-causality in many situations (Guo et al., 2008).

Partial G-causality is implemented in the GCCA toolbox by the function `cca_partialgc` for single trial data and by `cca_partialgc_mtrial` for multi-trial data. As with spectral G-causality, the statistical significance of partial G-causalities must be estimated by surrogate data methods (see Section 2.5).

## 2.4. G-autonomy

The framework of G-causality provides a novel approach to measuring the 'autonomy' of a variable, where autonomy is understood as the degree of self-determination or 'self-causation' exhibited by a variable (Seth, 2009; Bertschinger et al., 2008). Instead of asking whether the prediction error of  $X_1$  is reduced by including past observations of  $X_2$ , the G-autonomy measure asks whether the prediction error of  $X_1$  is reduced by inclusion of its own past, given a set of external variables  $X_2 \dots X_n$ . That is, a variable  $X_1$  is 'G-autonomous' to the extent that its own past states help predict its future states over and above predictions based on past states of a set of external variables  $X_2 \dots X_n$ . Put simply, a variable is G-autonomous to the extent that it is dependent on its own history and that these dependencies are not accounted for by external factors.

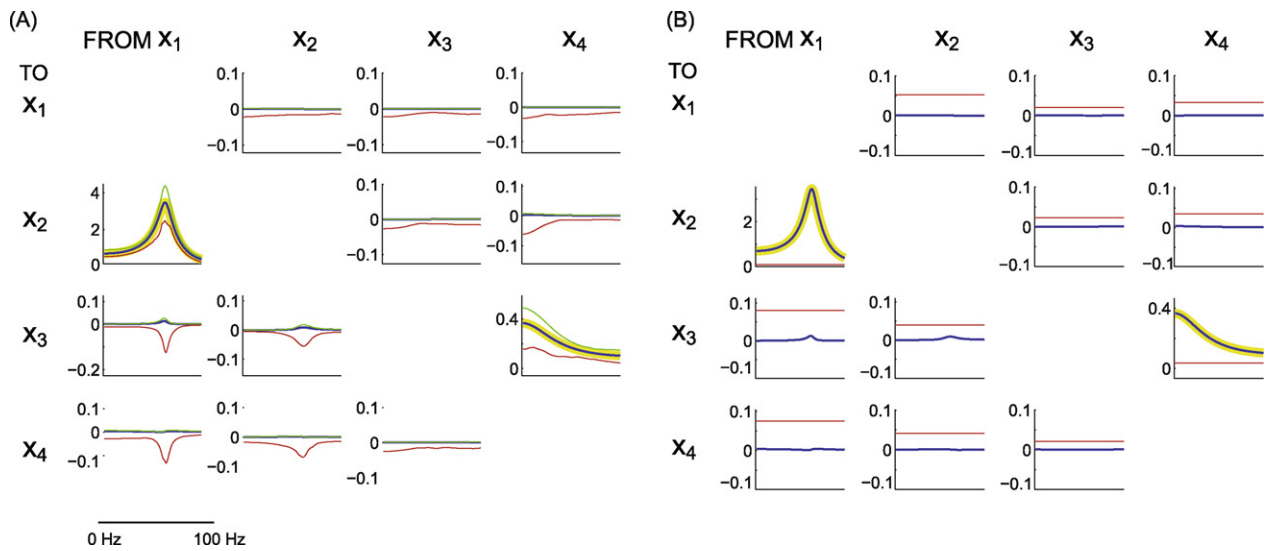
Recalling Eq. (2.2), if the variance of  $\xi_1$  (or  $\xi_2$ ) is reduced by the inclusion of the  $X_1$  (or  $X_2$ ) terms in the first (or second) equation, then it is said that  $X_1$  (or  $X_2$ ) is G-autonomous with respect to  $X_2$  (or  $X_1$ ). In other words,  $X_1$  is G-autonomous if the coefficients in  $A_{11}$  are jointly significantly different from zero. As with G-causality, this can be tested by performing an F-test of the null hypothesis that  $A_{11} = 0$ , given assumptions of covariance stationarity on  $X_1$  and  $X_2$ . By analogy with G-causality, the G-autonomy of  $X_1$  with respect to  $X_2$  is given by:

$$A_{1|2} = \ln \frac{\text{var}(\xi_{1R(11)})}{\text{var}(\xi_{1U})}, \quad (2.9)$$

where  $\xi_{1R(11)}$  is derived from the model omitting the  $A_{11,j}$  (for all  $j$ ) coefficients in the Granger equations.

In the toolbox, G-autonomy can be computed using the function `cca_autonomy_regress` and statistical significances identified using `cca_findsignificance_autonomy`. At present, G-autonomy

<sup>3</sup> `pwcausal.m` is available from <http://www.brain-smart.org>.

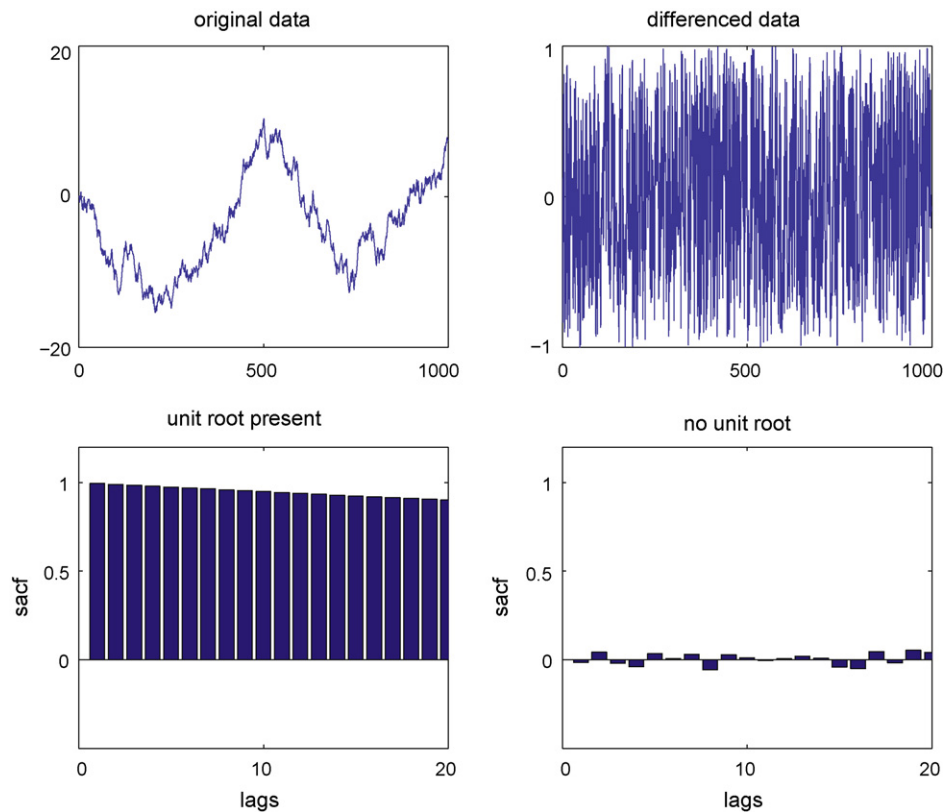


**Fig. 2.** Illustration of bootstrap (A) and permutation (B) resampling for setting significance thresholds for spectral G-causality (shown for 1–100 Hz); 500 resamples in each case. The generative model is given by Eq. (2.10), sampled at 500 Hz. (A) Thin lines show bootstrap confidence intervals around sample G-causality (thick line). (B) Thin red line shows permutation threshold; blue line shows sample G-causality. In both plots, yellow shading indicates frequencies at which G-causality significantly exceeds zero. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

has not been applied to neural data. It is however possible to envisage many interesting applications, for example in examining the conditions under which the activity of a particular brain area is G-autonomous with respect to a set of other areas. Illustrative applications of G-autonomy to computational models of interacting agents are described in Seth (2009).

### 2.5. Surrogate statistics

Several G-causality measures lack known statistical distributions. These measures include spectral G-causality, partial G-causality, and the ‘difference of influence’ metric that is useful for analyzing fMRI data (see Section 6). Establishing statistical



**Fig. 3.** The autocorrelation function can reveal non-CS variables. Left panels show a non-CS variable (top) with high and slowly declining autocorrelation (bottom). Right panels show the same variable after first-order differencing. In this example, the autocorrelation is greatly reduced at all lags.

significance in these cases requires surrogate methods such as bootstrapping and random permutation (Efron and Tibshirani, 1994). Permutation testing is useful for assessing whether an observed value is different from zero; bootstrapping places confidence intervals around a value, which can be useful either for distinguishing a value from zero or to compare apparently different values.

### 2.5.1. Bootstrap resampling

The premise of bootstrapping is that a single observation can stand in for a distribution if it is resampled with replacement (Efron and Tibshirani, 1994). In the context of MVAR models a data matrix is subdivided into many windows which are repeatedly sampled *with replacement* to generate surrogate data matrices. Importantly, causal relationships within each window are not disturbed. Confidence intervals on the sample value of the test statistic (e.g.,  $\mathcal{F}$ ) are then generated by examining the empirical quantiles of the bootstrap distribution. Typically, thousands of resamples are recommended for reliable testing. Bootstrap resampling is implemented in the toolbox via the functions `cca_partialgc_doi_bstrap` and `cca_pwcausal_bstrap` for time-domain and spectral G-causality analysis respectively. Both Bonferroni and FDR multiple comparison corrections are implemented within these functions.

### 2.5.2. Permutation resampling

In a permutation resampling test, a reference distribution is obtained by calculating many values of a test statistic following rearrangements of the labels on the data points. In other words, permutation tests involve resampling consistent with the assumption that the null hypothesis is true. Application to MVAR models involves subdividing a data matrix into many windows and constructing surrogate data matrices by rearranging the windows *for each variable separately*. The distribution of the test statistic (e.g.,  $\mathcal{F}$ ) over these resampled matrices then gives the expected distributions assuming a null hypothesis ( $\mathcal{F} = 0$ ). Significance thresholds can be derived by examining the empirical quantiles of this distribution. Permutation resampling is implemented in the toolbox via the functions `cca_partialgc_doi_permute` and `cca_pwcausal_permute`. In the spectral case, permutation significance thresholds for each potential causal interaction are determined by the method of Blair and Karniski (1993). According to this method, the G-causality spectrum is calculated for each permutation and the maximum G-causality value identified. A distribution of maximum G-causalities can then be derived following repeated permutations, and a per-interaction threshold is set against this distribution (Ding et al., 2006). After application of this procedure, Bonferroni and FDR corrections can be implemented to control for multiple comparisons among different potential interactions.

### 2.5.3. Example

To illustrate bootstrap and permutation resampling as applied to spectral G-causality, consider the simple system:

$$\begin{aligned} x_1(t) &= 0.95\sqrt{2}x_1(t-1) - 0.9025x_1(t-2) + w_1(t) \\ x_2(t) &= 0.5x_1(t-2) + w_2(t) \\ x_3(t) &= -0.4x_4(t-3) + w_3(t) \\ x_4(t) &= 0.35x_4(t-2) + w_4(t) \end{aligned} \quad (2.10)$$

In this system  $x_1$  drives  $x_2$  and  $x_4$  drives  $x_3$ ;  $w_1 - w_4$  are independent processes as in Eq. (1.1). There are no conditional relationships. Fig. 2 (right) shows the results of permutation significance testing using 500 resamples, a nominal  $P$  value of 0.05, Bonferroni correction for multiple comparisons, and the Blair and Karniski method for permutation resampling. As expected, only the correct causal

interactions exceed the permutation threshold. Similarly, only the correct causal interactions have bootstrap confidence intervals that do not intersect with zero (Fig. 2 left).

## 3. Preprocessing and validation

Meaningful application of the GCCA toolbox requires, minimally, that (i) the data satisfy certain preconditions and (ii) the MVAR models describe the data in a statistically satisfactory manner.

### 3.1. Preconditions and preprocessing

The primary precondition for G-causality analysis is that the variables must be *covariance stationary* (CS; also known as weak or wide-sense stationarity). Covariance stationarity requires that the first and second statistical moments (mean and variance) of each variable do not vary with time. If this condition is not satisfied, MVAR models will be invalid and may contain so-called ‘spurious regression’ results, i.e., correlations that arise from non-stationarities rather than from relations among variables (Granger and Newbold, 1974). One way to detect violations of CS is to examine the autocorrelation function of a variable (use `cca_sacf`). Non-CS variables typically will have an autocorrelation that declines slowly with increasing lags (Fig. 3 left). Most CS variables will have a sharply declining autocorrelation function (Fig. 3 right), though this is not true in all cases.

Formally, deviations from CS can be examined by testing for ‘unit roots’ within the data. The ‘augmented Dickey-Fuller’ (ADF) test assesses whether a unit root is present in a variable. The intuition behind this test is that if a variable is CS it will exhibit a tendency to return to a constant (or deterministically trending) mean. Therefore large values will tend to be followed by smaller values, and small values by larger values. The ADF test identifies the absence of this condition. The test is implemented in the toolbox by the function `cca_check_cov_stat`.

For non-CS variables, several preprocessing steps can be applied. First, deterministic linear trends can be removed via the function `cca_dettrend`. Second, unit roots can be removed via *differencing*, i.e.  $x'(t) = x(t) - x(t-1)$ , using function `cca_diff`. Often, one application of differencing is sufficient (see Fig. 3, right). Occasionally, repeated applications are required. It should be noted that differencing can change the interpretation of causal network analysis inasmuch as causal interactions are now among changes in variables rather than among the variables *per se*. Moreover, differencing can induce changes in the spectral profile of a time-series, complicating the interpretation of spectral G-causality. Repeated differencing obviously induces further challenges in interpretation.

If the data are in the form of multiple realizations (trials), non-stationarities due to variation of the mean during each trial can be removed by subtracting the ensemble mean, which is determined by averaging the values for each variable at each time point across trials (Ding et al., 2000). Similarly, intra-trial variation in the standard deviation can be controlled for by dividing through by the ensemble standard deviation. These capabilities are provided within the toolbox via the function `cca_rm_ensemblemean`. These steps can be extremely useful for event-related analyses, enabling causal interactions to be inferred based on the induced response. Note, in addition, that the MVAR algorithms in the toolbox require the time series to be zero mean (there is no constant term fitted in (2.1)). Therefore, the temporal mean should be removed either from each trial (for multitrial data) or from the entire time series. This step, as well as division through

by the temporal standard deviation, can be carried out using `cca_rm_temporalmean`.<sup>4</sup>

A different perspective on non-CS data is that causal relations among variables may be changing over time. In this case it may make sense to apply methods that are sensitive to time-varying causality. One approach is simply to divide the corresponding time series into short highly overlapping time windows each of which may be locally CS (Hesse et al., 2003; Ding et al., 2000). A constraint on this approach is to ensure that each window has sufficient data points (i.e., at least  $np$ ). Alternative approaches, not yet included in the toolbox but more generally applicable, include spectral factorization of wavelet transformations (Dhamala et al., 2008; see Section 7) and adaptive recursive least-squares modelling (Hesse et al., 2003).

In summary, it is good practice to apply the following pre-processing steps in the following order (Ding et al., 2000): (i) linear detrend, (ii) removal of temporal mean and division by temporal standard deviation, (iii) for multi-trial data, removal of ensemble mean and division by ensemble standard deviation, and (iv) differencing and/or windowing as necessary to achieve CS. If the objective is to examine time-varying causal connectivity then the windowing step could be introduced earlier, after the linear detrend.

### 3.2. Model validation

G-causality inferences are only valid if an MVAR model adequately captures the correlation structure in the data. There are several ways to check this. The simplest is to note the amount of variance accounted for by the model, in terms of the adjusted sum-square-error. This value is returned by the functions `cca_granger_regress` and `cca_granger_regress_mtrial`. Typically, a value of less than 0.3 signifies that the model may not have captured the data adequately. An alternative way to test the same intuition is by checking the model ‘consistency’, as defined by Ding et al. (2000). In the GCCA toolbox consistency is calculated as

$$c = \left(1 - \frac{|R_s - R_r|}{|R_r|}\right) \times 100 \quad (3.1)$$

where  $R_r$  is the correlation vector of the real data and  $R_s$  is the correlation vector of simulated data generated via the MVAR model. Each vector is a row vector (length  $n^2$ ) derived by reshaping the corresponding covariance matrix. As a rule of thumb, consistency values below 80% may give cause for concern.

A second important check on model validity makes use of the Durbin–Watson statistic which tests whether the residuals of a MVAR model are serially uncorrelated, as they should be if the model effectively captures the data (Durbin and Watson, 1950). The Durbin–Watson statistic  $d$  is calculated as

$$d = \frac{\sum_{t=2}^T (\xi_t - \xi_{t-1})^2}{\sum_{t=2}^T \xi_t^2} \quad (3.2)$$

As a rule of thumb, if  $d < 1.0$  there may be cause for concern. The toolbox returns  $d$  as well as a significance value based on the procedure described in Durbin and Watson (1950) (significance indicates serially correlated residuals). As implemented in the tool-

<sup>4</sup> Dividing by the standard deviation is useful for avoiding artifacts in causal connectivity due to differences in signal power (Ding et al., 2000). Of course, power differences can be physiologically significant and such differences should be analyzed separately, using data untreated by this step.

box, this test does *not* examine cross-correlations among residuals. It could be argued that such cross-correlations may signify model mis-specification in just the same way as does residual autocorrelation. However, as discussed in the context of partial G-causality (Section 2.3), residual cross-correlations could also signify the influence of latent and/or exogenous variables. For this reason, only the autocorrelation structure of residuals is examined in the GCCA toolbox.

All functions within the toolbox that perform a regression have the option of returning both the consistency value and the significance of the Durbin–Watson statistic.

## 4. Causal networks and their visualization

Methods for analyzing causal connectivity are especially powerful when applied in combination with graph-theoretic and network-theoretic techniques which allow their quantitative characterization (Seth, 2005). In causal networks, nodes represent variables or system elements and directed edges represent causal interactions. There is wide range of graph theory concepts that may find useful application to causal networks (Eichler, 2005; Bullmore and Sporns, 2009). Here, we describe those that are implemented in the GCCA toolbox.

### 4.1. Causal density

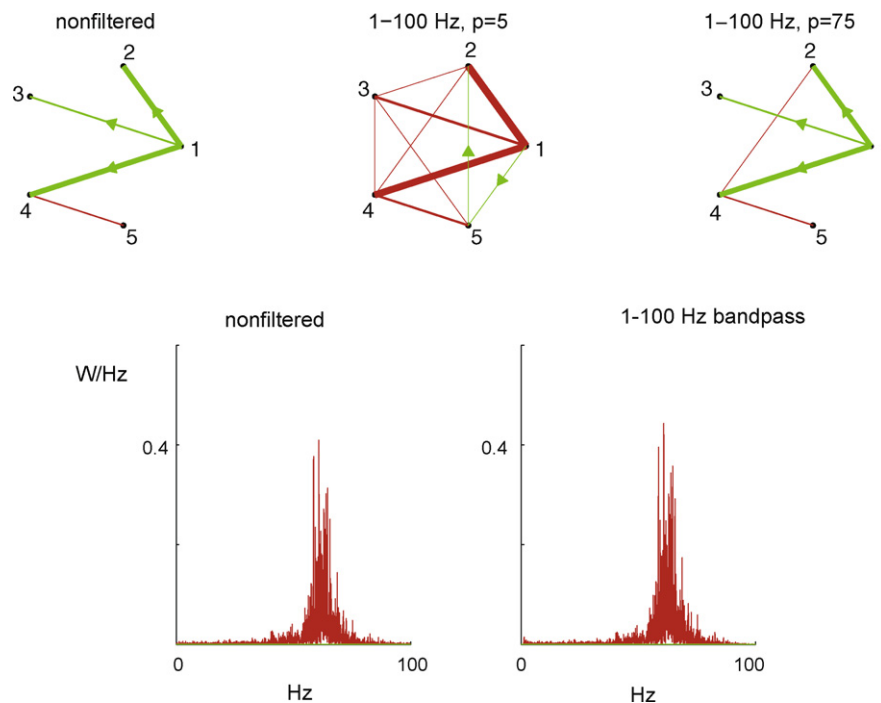
The *causal density* of the dynamics of a system  $\mathbf{X}$  is a global measure of causal interactivity (Seth, 2005, 2008). Causal density is defined as the mean of all pairwise G-causalities between system elements, conditioned on the remainder of the system; interactions which do not reach statistical significance are set to zero.

$$cd(\mathbf{X}) \equiv \frac{1}{n(n-1)} \sum_{i \neq j} \mathcal{F}_{X_i \rightarrow X_j | \mathbf{X}_{[ij]}} \quad (4.1)$$

where  $\mathbf{X}_{[ij]}$  denotes the subsystem of  $\mathbf{X}$  with variables  $X_i$  and  $X_j$  omitted. It is also possible to calculate a [0,1] bounded version of  $cd$  in which all statistically significant interactions are set to 1 (otherwise 0). Causal density provides a useful measure of *dynamical complexity* in the sense that complexity reflects the coexistence of integration and segregation in dynamics (Sporns, 2007; Shanahan, 2008). High values of  $cd$  indicate that system elements are globally coordinated in their activity (in order to be useful for predicting each other’s activity) and at the same time dynamically distinct (so that different elements contribute in different ways to these predictions) (Seth, 2005). A related quantity is *unit causal density*,  $cd_u(i)$ , which is the summed causal interactions involving node  $i$ , normalized by the number of nodes (see Fig. 1 bottom). A system with  $n$  elements will have  $n$   $cd_u$  values. (Again an unweighted version can be calculated by setting all significant interactions to 1.) Nodes with high values of  $cd_u$  can be considered to be *causal hubs* within  $\mathbf{X}$ . The function `cca_causal_density` can be used to calculate  $cd$  and  $cd_u$ . The toolbox also contains functions to calculate causal density by frequency, given the output of a spectral G-causality analysis.

### 4.2. Causal flow

The *causal flow* of node  $i$  in a causal network is defined as the difference between its (weighted or unweighted) in-degree and out-degree (see Fig. 1 bottom). A node with highly positive causal flow exerts a strong causal influence on the system as a whole and can be called a *causal source*. A node with a highly negative causal flow can be called a *causal sink*. Causal flows can be calculated using the function `cca_causal_flow`; calculation of spectral causal flows is also supported.



**Fig. 4.** Artifacts induced by bandpass filtering using FIR filter at 1–100 Hz. Raw data is from the model described in Fig. 1 with an assumed sampling rate of 500 Hz. Top-left panel shows the correct causal network (unfiltered data), top-middle panel shows the causal network following filtering and model estimation with  $p = 5$ . Top-right panel shows that most (but not all) artifacts are removed given a sufficiently high model order (e.g.,  $p = 75$ ). Green lines show unidirectional connections, red lines show bidirectional connections. Lower panels show power spectral density (W/Hz) for unfiltered (left) and filtered (right) data. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

#### 4.3. Visualization

The ability to effectively visualize causal networks is key to their interpretation. The GCCA toolbox includes functions for generating simple graphical depictions of network causal connectivity (Fig. 1, top right). Also included are functions for generating datafiles that describe a network in a format suitable for importing into the Pajek network analysis software<sup>5</sup> which contains many useful tools for network visualization and analysis. Functions enabling causal network visualization are `cca_plotcausality`, `cca_plotcausality_spectral`, and `cca_pajek`.

#### 5. Filtering

It is critical to ensure that preprocessing steps do not introduce spurious correlation structure into the data which can result in artifactual causal connectivity. In general, procedures that preserve the fine-grained timing relations among variable are safe, whereas those that do not, are not. One common preprocessing step that can cause problems is that of bandpass and/or notch filtering. According to Geweke (1982), the estimation of MVAR models in asymptotic conditions should not be affected by linear operations, such as application of standard finite impulse response (FIR) filters. However, under more realistic conditions filtering can be problematic. For example, Fig. 4 shows the time-domain G-causality artifacts introduced by bandpass filtering at 1–100 Hz, of data generated by the model described in Eq. (1.1). A standard FIR filter was used in two stages, low-pass with a 100 Hz high cutoff and high-pass with a 1 Hz low cutoff, implemented using the EEGLAB software (Delorme and Makeig, 2004).<sup>6</sup> One solution to this problem may be to use

substantially higher model orders so that the filter window is covered by the MVAR model. Fig. 4 (right) shows that with a model order  $p = 75$  the correct causal network is mostly, but not entirely recovered.

Some data sources require targeted removal of artifacts at specific frequencies. For example, EEG data is often contaminated by 'line noise' at the power supply frequency. The presence of line noise can severely disrupt time-domain G-causality analysis, as shown in Fig. 5A. Typically, line noise is removed via 'band-stop' or 'notch' filters tuned to the appropriate frequency. As with bandpass filters, band-stop/notch filters can induce time-domain causal network artifacts though the effects do not seem to be as severe (Fig. 5C). Again, high model orders may help remove or reduce these artifacts (Fig. 5D). An alternative, supported by the toolbox, is *multitaper filtering*, a procedure in which a sinusoidal oscillation is fitted to the line noise artifact and subtracted away (Mitra and Bokil, 2008; Brovelli et al., 2004). Because this operation is purely subtractive it should not introduce any additional G-causality artifacts (Fig. 5B). However, multitaper filtering is sensitive to choice of filter parameters and in practice may not always succeed in removing the line noise.

#### 6. Applications

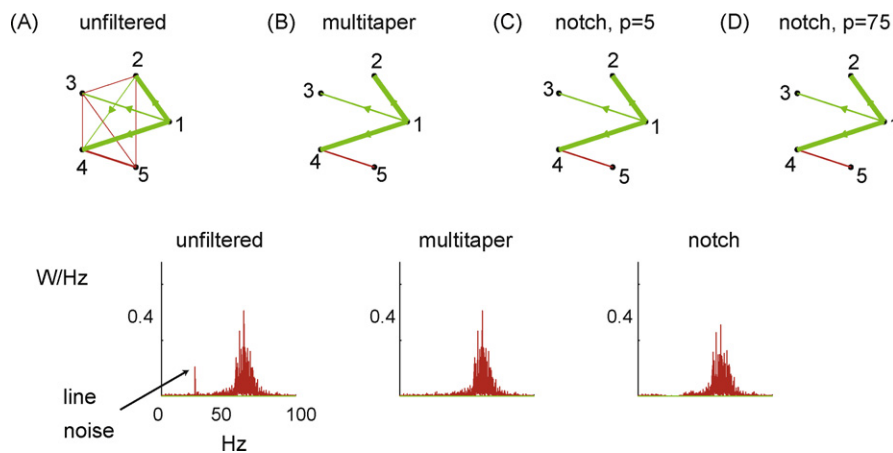
##### 6.1. Functional MRI

Application of GCCA to fMRI data is particularly attractive given the high spatial specificity of the fMRI BOLD signal. However, G-causality as applied to fMRI faces distinctive challenges and constraints. These challenges arise from the facts that the fMRI BOLD signal has relatively poor temporal resolution (on the order of seconds) as compared to other noninvasive neuroimaging techniques, and that it is an indirect measure of neural dynamics usually modelled by a convolution of underlying neural activity with a 'hemodynamic response function' (HRF), which reflects aspects of

<sup>5</sup> <http://vlado.fmf.uni-lj.si/pub/networks/pajek>.

<sup>6</sup> Each stage involves both forward and backward filtering in order to ensure precisely zero phase distortion.





**Fig. 5.** Artifacts induced by simulated line noise at 20 Hz. Raw data is from the model described in Fig. 1 with an assumed sampling rate of 500 Hz, contaminated with a 20 Hz oscillation applied separately to each variable. Top row shows causal networks, bottom row shows power spectral density (W/Hz). (A) Unfiltered data. (B) Following multitaper filtering (window size 1000, 5 tapers, half bandwidth of 3). (C) Following notch filtering (width 2 Hz) with low model order ( $p = 5$ ). (D) Following notch filtering with high model order ( $p = 75$ ). Notch filtering was applied using the `eeegfilt` function within the EEGLAB software. In this case, all treatments are effective but in general this outcome is not guaranteed.

neurovascular coupling. Of particular importance is that the delay embedded in the HRF is known to vary between subjects and among different brain regions within the same subject (Aguirre et al., 1998). Because G-causality is based on temporal precedence, such variability has the potential to introduce artifacts when assigning causality. The impact of HRF variation on G-causality analysis of the BOLD signal has recently been discussed extensively (Roebroeck et al., 2009 and commentaries). Below, we summarize parts of this discussion in the context of practical application.

The most useful response to challenge posed by neurovascular coupling would be to ‘undo’ the effects of the HRF via deconvolution, with GCCA then being applied to the recovered, uncontaminated, neural activity. Although reliable deconvolution requires information about the form of the HRF that is not normally available, some promising approaches have recently been suggested (Chang et al., 2008; Vokorin et al., 2007). In the absence of reliable deconvolution, two strategies can be followed to minimize the impact of HRF variability (Roebroeck et al., 2005). First, GCCA should focus on the modulation of causal connectivity by experimental condition (e.g., a higher level of causal density during waking than during sleep), as opposed to attempting to identify canonical causal structure for each condition independently. This is because, since HRFs are not expected to vary between conditions, it is unlikely that HRF variation across brain regions affects the comparison of conditions. Second, it may be useful to examine so-called ‘differences of influence’ (DOI) than causal connectivities *per se* (Roebroeck et al., 2005). The DOI term is determined by the asymmetry in causal connectivity (i.e.,  $\mathcal{F}_{A \rightarrow B} - \mathcal{F}_{B \rightarrow A}$ ). This term is relatively insensitive to artifactual bidirectional causal connectivities introduced by the slow and heterogeneous dynamics of the HRF, but of course at the expense of being constitutively unable to identify valid instances of reciprocal causality. DOI values are returned by the GCCA toolbox and their statistical significance can be assessed using the surrogate data methods provided. More generally, knowledge of the HRFs may not be necessary for G-causality analysis to disclose causal relations among event-related BOLD time series (Tang et al., 2009), and when combined with results obtained from ablation and/or stimulation studies, GCCA of BOLD signals may be particularly informative.

Prior to confronting issues of HRF shape and variability, fMRI analyses involve a preprocessing pipeline which must be undertaken with care when G-causality is applied. The follow-

ing procedure can be recommended: (i) slice timing correction; (ii) spatial realignment to control for movement artifacts; (iii) coregistration to structural MRI; (iv) normalization to a standard coordinate space (e.g., Talairach space); (v) spatial smoothing (if required), (vi) covarying out of so-called ‘nuisance covariates’ such as the global signal, white matter, cerebro-spinal fluid, and motion components), and finally (vii) time series extraction by averaging (or taking the first eigenvariate or principal component of) all voxels within a region-of-interest (ROI). Importantly, temporal smoothing is avoided. Standard preprocessing techniques can then be applied as described in Section 3. It should be noted that this pipeline is only suggestive and does not exclude other valid approaches.

Roebroeck and colleagues have proposed a strategy for applying G-causality to fMRI, which they call “Granger causality mapping” (GCM). GCM combines the steps described above with the procedure of selecting a single seed voxel (ROI) and then computing a large number of independent bivariate G-causality models including the seed voxel and – one-by-one – all other voxels in the fMRI dataset. GCM is therefore an exploratory approach which allows identification of voxels that are causally implicated in the activity of the seed voxel. It should be noted that this approach can be complemented with multivariate G-causality analyses of preselected networks of voxels or ROIs (Deshpande et al., 2009).

## 6.2. EEG, MEG, LFP

In contrast to fMRI, electromagnetic measures of neural activity [e.g., electroencephalography (EEG), magnetoencephalography (MEG) and intracranially recorded local field potentials (LFP)] have a time resolution on the order of action potential generation (i.e., milliseconds). Such measures are however lacking in spatial resolution and/or coverage, are susceptible to artifacts including electrical line-noise and, for MEG/EEG a range of other environmental and non-neural physiological influences, each of which has the potential to introduce confounds affecting G-causality analysis.

One common approach to artifact removal is to use bandpass filtering; however, we have already discussed issues involved in band-pass and band-stop filtering, suggesting that multi-taper filtering offers a useful alternative to band-stop (notch) filters for the selective removal of line-noise artifacts (see Section 5). For more general artifact removal, manual inspection and rejection of

contaminated data is time consuming but effective. Independent components analysis (ICA) may also remove many artifact sources on the logic that certain artifacts (e.g., stereotyped eye and muscle artifacts) will be temporally independent from neural activities (Jung et al., 2000; Delorme and Makeig, 2004). Because ICA is a subtractive method it should be compatible with G-causality analysis.

Following artifact removal or rejection, it can be useful to down-sample the data so that the corresponding MVAR model order lies within a reasonable range. For example, data sampled at 1 kHz can be downsampled to 250 Hz by retaining only every 4th data point in each time series (e.g., Brovelli et al., 2004). EEG data must also be re-referenced to a common source. Among several possibilities, re-referencing to the global average has been used successfully (Zhang and Ding, 2009) though care must be taken to ensure that artifactual cross-correlations are not introduced. Alternatively, conversion of sensor-space EEG/MEG data into estimates of surface current source density (e.g., using the BESA software package, MEGIS Software GmbH) removes the need to establish a reference. For LFP data, if required, bipolar montages can be constructed by subtracting signals recorded from adjacent sites belonging to each electrode (Gaillard et al., 2009). If the data are event-related, it is also important to subtract the pre-stimulus baseline from each trial. Event-related data also invite subtraction of the ensemble mean (see Section 3.1) so that G-causality analysis operates on the induced response, rather than on the event-related signal itself.

EEG and MEG measure the population effects of neural activity as they appear at the scalp surface, consequently localization to underlying neural sources must be indirectly inferred via inverse modelling. Reviewing the compatibility of source modelling methods with G-causality analysis is an active area of investigation that lies beyond the present scope (see, e.g., Babiloni et al., 2005). Note that surface current source density estimates may provide a useful 'middle ground' between sensor space and source localization (Cohen et al., 2009).

In summary, the following pipeline can be recommended for G-causality analysis of electrical/magnetic neural signals: (i) artifact removal/rejection; (ii) downsampling; (iii) re-referencing (EEG), current source density estimation (EEG/MEG), or optional bipolar montage generation (LFP); and finally for event-related data (iv) subtraction of pre-stimulus baseline. Standard preprocessing techniques (including ensemble mean removal) can then be applied as described in Section 3. Again, this pipeline is suggestive and does not exclude alternative approaches.

### 6.3. Spike trains

Recording of trains of individual action potentials (spikes), from multiple neurons in parallel, is increasingly prominent and holds great promise for elucidating the operating principles of neural systems. Such data can be acquired from implanted electrodes and from multi-electrode arrays on which brain slices or living neural cultures can be mounted (Cadotte et al., 2008; Harris et al., in press). Spike train data pose a challenge for G-causality analysis because they consist in point processes (i.e., sequences of 'on'/off' states, with 'on' states generally sparse in neural data), rather than in discrete samples of continuous processes. The simplest approach to dealing with spike train data is to convolve the spike train with a Gaussian (e.g., by low pass filtering), or half-Gaussian, of suitable width. Convolution effectively transforms the spike train into a continuous estimate of mean firing rate. A more complex but in principle more suitable method is to estimate a spectral representation of the point process data and then to apply specialized factorization techniques to derive causality metrics directly from this spectral representation (Nedungadi et al., 2009). Alternatively, Okatan et al. (2005) describe a maximum likelihood approach to estimating functional connectivity, designed specifically for point-

process data, that is not based on G-causality. These latter methods are not included in the GCCA toolbox.

## 7. Limitations

The current GCCA toolbox does not support several potentially useful analyses. One limitation is that *multivariate* spectral G-causality analysis cannot be performed. The best method for performing this analysis remains an active area of investigation. The standard approach is to analyze spectral features of autoregressive models following (Geweke, 1982), as implemented here for pairwise analysis by the function `cca_pwcausal`. However, in the multivariate case Geweke spectral causality can occasionally be negative, an outcome which eludes physical interpretation (Chen et al., 2006). Very recently, Ding and colleagues have introduced a nonparametric approach to computing spectral G-causality based on a factorization of wavelet transformations (Dhamala et al., 2008) [as already noted, this method is also suited to analysis of spike train point-process data (Nedungadi et al., 2009)]. This method always provides non-negative results, offers robustness to uncertainties in determining autoregressive model order and is better placed to capture complex spectral features of data. These advantages are especially important for multivariate (>3 variable) data and it is therefore intended that future versions of the GCCA toolbox will implement a version of this method.

A second limitation is that nonlinear G-causality analysis is not supported. Again, there are several competing approaches to this problem. A simple solution (Seth, 2009) is to fit autoregressive coefficients to Taylor expansions of the data, however this method requires estimating large numbers of parameters. Alternative approaches include locally nonlinear autoregressive models (Freiwald et al., 1999), nonlinear kernels such as radial basis functions (Ancona et al., 2004), and adapting methods from information theory such as transfer entropy (Schreiber, 2000). Interestingly, we have recently shown that, for Gaussian variables, G-causality and transfer entropy are entirely equivalent (Barnett et al., 2009). One implication of this result is that, under Gaussian assumptions, there is nothing extra to account for by nonlinear extensions to G-causality, since a stationary Gaussian autoregressive process is necessarily linear (Barnett et al., 2009). It should also be noted that although the brain is highly nonlinear at many levels of description, linear effects can be among the robust especially in the context of large-scale neurodynamic interactions (Freeman, 1972; McIntosh et al., 1994).

A general limitation of the toolbox, and one that attends most G-causality methods, is that scaling to massively multivariate data-sets is extremely challenging because of the large number parameters that would need to be estimated for a fully multivariate model. This problem is particularly challenging for fMRI BOLD data which typically consists of very large numbers of variables (voxels), each with only a few observations, as compared to other neuroscience data sources. The simplest approach to dealing with high-dimensional data is to perform repeated pairwise analysis. More advanced approaches could utilize sparse (regularized/penalized) regression techniques which are able to perform massively multivariate autoregression by imposing priors on the values of most coefficients (i.e., assume that most causal connections do not exist). For example, Valdés-Sosa et al. (2005) have developed a method that combines penalized regression with pruning of unlikely connections using a local false discovery rate. They validate their approach on simulated idealized cortical networks, and show that it can identify neural circuitry related to emotional processing as measured by fMRI.

Scaling could also be helped by increasing the computational performance of the toolbox, for example by combining the flexibil-

ity of MATLAB with the computational velocity of other languages (e.g., C/C++), and also with possibilities for hardware acceleration offered by parallel computation, either via standard Beowulf clusters or via graphical processor unit (GPU) technologies. GPU technology in particular is developing extremely rapidly and seems able to dramatically enhance the parallelization of the large-scale matrix manipulations required by G-causality analysis. Future toolbox versions will explore these important issues.

## 8. Discussion

The GCCA toolbox provides a range of MATLAB functions enabling the application of G-causality analysis to a broad range of neuroscience data. Together with the theoretical and practical context provided in this paper, it is well placed to facilitate progress in a wide cross-section of the neuroscience community. However, it is critical to ensure that the toolbox is not treated as ‘a black box’ whereby its output is assumed to be a valid and informative transformation of its input. As with any complex procedure the maxim ‘garbage in, garbage out’ applies. For this reason the toolbox incorporates a comprehensive range of validation and preprocessing tools, appropriate combinations of which must be tailored to each new application. Also for this reason, and in contrast to most similar projects, the GCCA toolbox deliberately avoids the use of graphical user interfaces (GUIs), on the logic that such interfaces can render the underlying algorithms opaque and inaccessible. The fact that the toolbox employs only MATLAB functions throughout should also help ensure transparency and accessibility.

### 8.1. Exploratory and confirmatory statistics

G-causality can usefully be considered as an *exploratory* rather than *confirmatory* approach to discerning directed functional connectivity. This is because G-causality is permissive with respect to structural constraints, typically making few if any assumptions about the causal connectivity patterns embedded in the data. Confirmatory approaches, by contrast, first specify a candidate causal mechanism and then test whether the data fit the proposed model. This distinction is sometimes phrased in terms of ‘data driven’ versus ‘model driven’, but this phraseology is misleading because both exploratory and confirmatory approaches involve the estimation of mathematical models from data: MVAR models in the case of G-causality, and other more constrained models for confirmatory approaches.

It is instructive to briefly compare G-causality with a popular and *prima facie* confirmatory alternative, *dynamic causal modelling* (DCM, Friston et al., 2003). As applied to neuroscience data, DCM involves the estimation of parameters reflecting directed functional connectivity and also of parameters reflecting the generation of observable data from underlying neural mechanisms. For example, DCM for fMRI involves the estimation of a hemodynamic forward model connecting the fMRI BOLD signal to neuronal activity.<sup>7</sup> As compared to G-causality, the explicit incorporation of a forward model can mitigate against artifactual causal connections arising from (for example) inter-areal variance in hemodynamic lag, but such mitigation is critically dependent on the accuracy of these forward models, the estimation of which still relies on the same data. To put the point slightly differently, given a certain covariance structure in the data, G-causality will account for that structure in terms of connectivity whereas DCM will account for that structure via some combination of connectivity and neurovascular coupling properties. Other significant differences also apply (Roebroeck et

al., 2009). Standard DCM relies on a well characterized input, since there is typically no stochastic term (though see Friston et al., 2008), whereas G-causality is based explicitly on stochastic models and so can be applied equally to spontaneously generated dynamics. In practice, DCM is limited to a small number of variables (nodes, regions-of-interest) as compared to G-causality, a consequence of the increased number of parameters that must be estimated for each variable. It should also be noted that DCM can also involve an exploratory component in the form of selection among multiple candidate models based on their ability to account for the data via their log-evidence, a Bayesian procedure (Penny et al., 2004).

In short, although DCM and G-causality are often viewed as competing methods (David et al., 2008; Friston, 2009) they may be better understood as complementary, inhabiting different parts of the exploratory-confirmatory spectrum but not confined to the extremes. Formally, both approaches can be given a state-space interpretation in which relations between measured variables and input variables are mediated via unobservable state variables (Roebroeck et al., 2009). In practice, G-causality may be easier to apply because of (i) the lack of explicit forward models and (ii) increased latitude with respect to number of variables incorporated. In either case, the advice of Box and Draper (1987, p. 424) should not be forgotten: “All models are wrong, but some are useful”.

### 8.2. Documentation and licensing

The GCCA toolbox can be freely downloaded from [www.anilseth.com](http://www.anilseth.com) under the GNU general public license (version 3, see [www.gnu.org](http://www.gnu.org)). A manual is included in the download package which describes in detail each MATLAB function, provides a walk-through tutorial of the toolbox in action, and includes demonstrations of other more specific functions. All functions are carefully commented so that MATLAB literate researchers should be able to understand and modify the code with ease. A dedicated web-site has been created on which users are encouraged to post comments, describe bugs and bug-fixes, and suggest proposed enhancements. A list of known errors and ‘frequently asked questions’ is maintained at this site, which can be found at <http://ccatoolbox.pbwiki.com>.

## Acknowledgements

I am grateful to A. Barrett, L. Barnett, and R. Giacomini for guidance and comments, and to many others for useful feedback on the first version of the GCCA toolbox (see the manual for individual acknowledgements). Many thanks to H. Liang and colleagues for providing links to the functions `pwcausal.m` and `armorf.m` which are part of the BSMART toolbox (Cui et al., 2008). The author’s research is funded by EPSRC Leadership Fellowship EP/G007543/1 and by a donation from the Mortimer and Theresa Sackler Foundation.

## References

- Aguirre GK, Zarahn E, D’Esposito M. The variability of human, BOLD hemodynamic responses. *Neuroimage* 1998;8:360–9.
- Akaike H. A new look at the statistical model identification. *IEEE Trans Autom Control* 1974;19:716–23.
- Ancona N, Marinazzo D, Stramaglia S. Radial basis function approaches to nonlinear granger causality of time series. *Phys Rev E* 2004;70:056221.
- Babiloni F, Cincotti F, Babiloni C, Carducci F, Mattia D, Astolfi L, et al. Estimation of the cortical functional connectivity with the multimodal integration of high-resolution EEG and fMRI data by directed transfer function. *Neuroimage* 2005;24:118–31.
- Baccalá LA, Sameshima K. Partial directed coherence: a new concept in neural structure determination. *Biol Cybern* 2001;84:463–74.
- Barnett L, Barrett A, Seth A. Granger causality and transfer entropy are equivalent for Gaussian variables. *Phys Rev Lett* 2009;103:238701.

<sup>7</sup> Different DCMs are required for each different neuroimaging methodology.

- Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc: Ser B (Methodol)* 1995;57:289–300.
- Bertschinger N, Olbrich E, Ay N, Jost J. Autonomy: an information theoretic perspective. *Biosystems* 2008;91:331–45.
- Blair RC, Karniski W. An alternative method for significance testing of waveform difference potentials. *Psychophysiology* 1993;30:518–24.
- Box G, Draper N. *Empirical model building and response surfaces*. New York: John Wiley and Sons; 1987.
- Brovelli A, Ding M, Ledberg A, Chen Y, Nakamura R, Bressler S. Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. *Proc Natl Acad Sci USA* 2004;101(26):9849–54.
- Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 2009;10:186–98.
- Cadotte AJ, DeMarse TB, He P, Ding M. Causal measures of structure and plasticity in simulated and living neural networks. *PLoS One* 2008;3:e3355.
- Chang C, Thomason ME, Glover GH. Mapping and correction of vascular hemodynamic latency in the BOLD signal. *Neuroimage* 2008;43:90–102.
- Chen Y, Bressler SL, Ding M. Frequency decomposition of conditional Granger causality and application to multivariate neural field potential data. *J Neurosci Methods* 2006;150:228–37.
- Cohen MX, van Gaal S, Ridderinkhof KR, Lamme VA. Unconscious errors enhance prefrontal-occipital oscillatory synchrony. *Frontiers in Human Neuroscience* 2009;3:54.
- Cui J, Xu L, Bressler SL, Ding M, Liang H. BSMART: a Matlab/C toolbox for analysis of multichannel neural time series. *Neural Netw* 2008;21:1094–104.
- David O, Guillemain I, Saillet S, Reyt S, Deransart C, Segebarth C, Depaulis A. Identifying neural drivers with functional MRI: an electrophysiological validation. *PLoS Biol* 2008;6:2683–97.
- Delorme A, Makeig S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 2004;134:9–21.
- Deshpande G, LaConte S, James GA, Peltier S, Hu X. Multivariate Granger causality analysis of fMRI data. *Hum Brain Mapp* 2009;30:1361–73.
- Dhamala M, Rangarajan G, Ding M. Analyzing information flow in brain networks with nonparametric Granger causality. *Neuroimage* 2008;41:354–62.
- Ding M, Bressler S, Yang W, Liang H. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. *Biol Cybern* 2000;83:35–45.
- Ding M, Chen Y, Bressler S. Granger causality: basic theory and application to neuroscience. In: Schelter S, Winterhalder M, Timmer J, editors. *Handbook of time series analysis*. Weinheim: Wiley; 2006. p. 438–60.
- Doesburg SM, Green JJ, McDonald JJ, Ward LM. Rhythms of consciousness: binocular rivalry reveals large-scale oscillatory network dynamics mediating visual perception. *PLoS One* 2009;4, e6142–e6142.
- Durbin J, Watson GS. Testing for serial correlation in least squares regression. I. *Biometrika* 1950;37:409–28.
- Efron B, Tibshirani R. *Introduction to the bootstrap*. Monographs on statistics and applied probability. New York, NY: Chapman & Hall; 1994.
- Eichler M. A graphical approach for evaluating effective connectivity in neural systems. *Philos Trans R Soc B* 2005;360:953–67.
- Engel A, Singer W. Temporal binding and the neural correlates of sensory awareness. *Trends Cogn Sci* 2001;5:16–25.
- Freeman W. Linear analysis of the dynamics of neural masses. *Annu Rev Biophys Bioeng* 1972;1:222–56.
- Freiwald W, Valdes P, Bosch J, Biscay R, Jimenez J, Rodriguez L, et al. Testing non-linearity and directedness of interactions between neural groups in the macaque inferotemporal cortex. *J Neurosci Methods* 1999;94:105–19.
- Friston KJ, Harrison L, Penny W. Dynamic causal modelling. *Neuroimage* 2003;19:1273–302.
- Friston KJ, Trujillo-Barreto N, Daunizeau J. DEM: a variational treatment of dynamic systems. *Neuroimage* 2008;41:849–85.
- Friston K. Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS Biol* 2009;7:e33.
- Gaillard R, Dehaene S, Adam C, Clémenceau S, Hasboun D, Baulac M, et al. Converging intracranial markers of conscious access. *PLoS Biol* 2009;7:e61.
- Geweke J. Measurement of linear dependence and feedback between multiple time series. *J Am Stat Assoc* 1982;77(378):304–13.
- Gourévitch B, Bouquin-Jeannes RL, Faucon G. Linear and nonlinear causality between signals: methods, examples and neurophysiological applications. *Biol Cybern* 2006;95:349–69.
- Gow DW, Segawa JA. Articulatory mediation of speech perception: a causal analysis of multi-modal imaging data. *Cognition* 2009;110:222–36.
- Gow DW, Segawa JA, Ahlfors SP, Lin F-H. Lexical influences on speech perception: a Granger causality analysis of MEG and EEG source estimates. *Neuroimage* 2008;43:614–23.
- Granger CWJ. Economic processes involving feedback. *Inform Control* 1963;6:28–48.
- Granger C. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 1969;37:424–38.
- Granger C, Newbold P. Spurious regressions in econometrics. *J Econom* 1974;2:111–20.
- Guo S, Seth A, Kendrick K, Zhou C, Feng J. Partial granger causality: eliminating exogenous inputs and latent variables. *J Neurosci Methods* 2008;172:79–93.
- Harris C, Passaro P, Kemenes I, Kemenes G, O'Shea M. Sensory driven multi-neural activity and associative learning monitored in an intact CNS on a multielectrode array. *J Neurosci Methods*; in press.
- Hesse W, Möller E, Arnold M, Schack B. The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies. *J Neurosci Methods* 2003;124:27–44.
- Jung TP, Makeig S, Humphries C, Lee TW, McKeown MJ, Iragui V, et al. Removing electroencephalographic artifacts by blind source separation. *Psychophysiology* 2000;37:163–78.
- Kaminski M, Ding M, Truccolo WA, Bressler SL. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. *Biol Cybern* 2001;85:145–57.
- Keil A, Sabatinelli D, Ding M, Lang PJ, Ihssen N, Heim S. Re-entrant projections modulate visual cortex in affective perception: evidence from Granger causality analysis. *Hum Brain Mapp* 2009;30:532–40.
- McIntosh A, Grady C, Ungerleider L, Haxby J, Rapoport S, Horwitz B. Network analysis of cortical visual pathways mapped with PET. *J Neurosci* 1994;14:655–66.
- Mitra P, Bokil H. *Observed brain dynamics*. Oxford: Oxford University Press; 2008.
- Morf M, Viera A, Lee D, Kailath T. Recursive multichannel maximum entropy spectral estimation. *IEEE Trans Geosci Electron* 1978;16:85–94.
- Nedungadi AG, Rangarajan G, Jain N, Ding M. Analyzing multiple spike trains with nonparametric granger causality. *J Comput Neurosci* 2009;27:55–64.
- Nunez PL, Wingeier BM, Silberstein RB. Spatial-temporal structures of human alpha rhythms: theory, microcurrent sources, multiscale measurements, and global binding of local networks. *Hum Brain Mapp* 2001;13:125–64.
- Okatan M, Wilson M, Brown E. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput* 2005;17(9):1927–61.
- Pearl J. *Causality: models, reasoning, and inference*. Cambridge, UK: Cambridge University Press; 1999.
- Penny WD, Stephan KE, Mechelli A, Friston KJ. Comparing dynamic causal models. *Neuroimage* 2004;22:1157–72.
- Pereda E, Quiroga RQ, Bhattacharya J. Nonlinear multivariate analysis of neurophysiological signals. *Prog Neurobiol* 2005;77:1–37.
- Roebroeck A, Formisano E, Goebel R. Mapping directed influence over the brain using granger causality and fmri. *Neuroimage* 2005;25(1):230–42.
- Roebroeck A, Formisano E, Goebel R. The identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. *Neuroimage* 2009.
- Sato JR, Junior EA, Takahashi DY, de Maria Felix M, Brammer MJ, Morettn PA. A method to produce evolving functional connectivity maps during the course of an fMRI experiment using wavelet-based time-varying Granger causality. *Neuroimage* 2006;31:187–96.
- Schreiber T. Measuring information transfer. *Phys Rev Lett* 2000;85(2):461–4.
- Schwartz G. Estimating the dimension of a model. *Ann Stat* 1978;5(2):461–4.
- Seth A. Causal connectivity of evolved neural networks during behavior. *Network: Comput Neural Syst* 2005;16:35–54.
- Seth A. Granger causality. *Scholarpedia* 2007;2(7):1667.
- Seth A. Causal networks in simulated neural systems. *Cogn Neurodyn* 2008;2:49–64.
- Seth A. Measuring autonomy and emergence via granger causality. *Artif Life* 2009;16:2.
- Seth A, Edelman G. Distinguishing causal interactions in neural populations. *Neural Comput* 2007;19:910–33.
- Shanahan M. Dynamical complexity in small-world networks of spiking neurons. *Phys Rev E Stat Nonlin Soft Matter Phys* 2008;78:041924.
- Sporns O. *Complexity*. Scholarpedia 2007;2(10):1623.
- Sridharan D, Levitin DJ, Menon V. A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proc Natl Acad Sci USA* 2008;105:12569–74.
- Stevens MC, Pearson GD, Calhoun VD. Changes in the interaction of resting-state neural networks from adolescence to adulthood. *Hum Brain Mapp* 2009;30:2356–66.
- Tang W, Bressler SL, Sylvester CM, Shulman GL, Corbetta M. Fmri-based granger causality is an effective measure of effective connectivity. Abstract at the Society for Neuroscience (SfN) annual meeting, programme number 188.15, 2009.
- Vakorin VA, Krakovska OO, Borowsky R, Sarty GE. Inferring neural activity from BOLD signals through nonlinear optimization. *Neuroimage* 2007;38:248–60.
- Valdés-Sosa PA, Sánchez-Bornot JM, Lage-Castellanos A, Vega-Hernández M, Bosch-Bayard J, Melie-García L, et al. Estimating brain functional connectivity with sparse multivariate autoregression. *Philos Trans R Soc Lond B Biol Sci* 2005;360:969–81.
- Wald A. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Trans Am Math Soc* 1943;54(3):426–82.
- Whittle P. The analysis of multiple stationary time series. *J R Stat Soc B* 1953;15(1):125–39.
- Zhang Y, Ding M. Detection of a weak somatosensory stimulus: role of the prestimulus Mu rhythm and its top-down modulation. *J Cogn Neurosci* 2009.