Published in final edited form as:

Magn Reson Imaging. 2010 October ; 28(8): 1058–1065. doi:10.1016/j.mri.2010.03.002.

Functional MRI and Multivariate Autoregressive Models

Baxter P. Rogers, Ph.D.,

Vanderbilt University Institute of Imaging Science Department of Radiology and Radiological Sciences Department of Biomedical Engineering

Santosh B. Katwal,

Vanderbilt University Department of Electrical Engineering and Computer Science

Victoria L. Morgan,

Vanderbilt University Institute of Imaging Science Department of Radiology and Radiological Sciences

Christopher L. Asplund, and

Vanderbilt University Department of Psychology

John C. Gore

Vanderbilt University Institute of Imaging Science Department of Radiology and Radiological Sciences Department of Biomedical Engineering Department of Molecular Physiology and Biophysics Department of Physics

Abstract

Connectivity refers to the relationships that exist between different regions of the brain. In the context of functional magnetic resonance imaging (fMRI), it implies a quantifiable relationship between hemodynamic signals from different regions. One aspect of this relationship is the existence of small timing differences in the signals in different regions. Delays of 100 ms or less may be measured with fMRI, and these may reflect important aspects of the manner in which brain circuits respond as well as the overall functional organization of the brain. The multivariate autoregressive time series model has features to recommend it for measuring these delays, and is straightforward to apply to hemodynamic data. In this review, we describe the current usage of the multivariate autoregressive model for fMRI, discuss the issues that arise when it is applied to hemodynamic time series, and consider several extensions. Connectivity measures like Granger causality that are based on the autoregressive model do not always reflect true neuronal connectivity; however, we conclude that careful experimental design could make this methodology quite useful in extending the information obtainable using fMRI.

INTRODUCTION

Connectivity

Connectivity refers to the relationships that exist between different regions of the brain. In the context of functional magnetic resonance imaging (fMRI), it implies some quantifiable relationship between hemodynamic signals from different regions. Measuring connectivity has

© 2010 Elsevier Inc. All rights reserved.

(Corresponding author) baxter.rogers@vanderbilt.edu 615-343-4636 1161 21st Ave S AA 1105 MCN Nashville, TN 37232-2310.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

provided useful information about the brain's functional organization. For example, the observation of multi-region brain networks with common patterns of activity [1] even when no particular stimulus is presented [2,3] has been used to identify discrete brain circuits. A wide variety of connectivity measures have been used [4,5]. They include simple correlations based on different sources of variance — intra-subject and inter-subject, intrinsic or task-based — as well as more sophisticated measures that incorporate a biophysical model of the hemodynamics [6] or temporal information in the time series data [7].

Temporal information in fMRI data

Small temporal differences in BOLD responses may play a crucial role in understanding the connectivity of the brain [8]. Functional MRI can detect differences in the timing of neural activity at 100ms or less, even though the BOLD response itself typically takes 5-8 seconds to reach its peak and 15-30 seconds to return to baseline. An early work clearly observed latency differences down to 125 ms for left/right hemifield visual stimuli, by taking pains to exclude high amplitude signals and apparent veins from the ROIs [9]. Temporal differences in this range have been correlated with behavioral differences [10-12]. Another study has shown the ability to trace the temporal evolution of brain activity at 100ms time scales in a stimulus-response reading task [13].

The multivariate autoregressive model

Multivariate autoregressive modeling (MAR modeling) is a time series analysis procedure often used to characterize dynamic systems because of its simplicity. It provides information on the temporal properties of a stationary linear system. Measures can be derived from the model parameters that represent the relationships within the system in time and frequency domains. A specific example is the Granger causality, a summary measure of how one time series predicts another. In the context of brain imaging, such measures may reflect the temporal precedence of hemodynamic activity in different voxels or brain regions. Under some circumstances this could be informative about effective connectivity, or how activity in one region influences activity in another. In this review, we describe the current usage of the multivariate autoregressive model for fMRI, discuss the issues that arise when it is applied to hemodynamic time series, and consider several extensions.

METHODOLOGY

Basic Model and Granger causality

The simplest MAR model is the first order bivariate model, considering time series from two channels X and Y that might represent the MR signals from two voxels or regions of interest within the brain:

$$\begin{aligned}X_t &= aX_{t-1} + bY_{t-1} + w \\ Y_t &= cX_{t-1} + dY_{t-1} + z\end{aligned}$$

The parameters a and d represent the autocorrelation of each signal. The parameters b and c represent the cross-correlation, the ability of X to predict Y and vice versa. The terms w and z represent random innovations, or noise. This is a simple parametrization, but the model may be extended to additional variables and higher order (more lags in time). A summary of the interdependence of X and Y can be obtained from the Granger causality [14,15], which is calculated as follows.

First, fit the univariate model for X alone:

$$X_t = aX_{t-1} + u$$

Then fit the bivariate model:

$$X_t = aX_{t-1} + bY_{t-1} + w$$

Calculate the Granger causality from the variances of the residuals u and w :

$$F(Y \rightarrow X) = \log \frac{\text{var}(u)}{\text{var}(w)}$$

The variance ratio cannot be less than 1, since additional model parameters cannot explain less variance, so the Granger causality exists on the interval $[0, \infty)$. It represents the degree to which the signal Y predicts the signal X , as represented by the decrease in residual variance after adding Y to the model. The corresponding measure $F(X \rightarrow Y)$ may be calculated in similar fashion. A convenient summary of the directionality of the relationship between X and Y is the quantity $F(X \rightarrow Y) - F(Y \rightarrow X)$, which has a value of zero in the absence of directional preference [16].

Basic properties of Granger causality in the context of fMRI

One clear-cut scenario where effective connectivity has been measured by Granger causality is when neural activity occurs in one region prior to another, with a predominantly unidirectional flow of information. An example would be a simple visuomotor reaction time experiment, where V1 neural activity precedes M1 activity by approximately the measured reaction time. This sort of neural delay may be measured in hemodynamic data using an MAR model [17]. Here, the MAR model is used to capture the temporal dynamics of the stimulus-induced BOLD response. Furthermore, the neural delay may be manipulated experimentally with any approach that affects reaction time, like increasing the complexity of the stimulus [18]. This manipulation ensures that a change in Granger causality reflects a change in neural processing, avoiding confounds introduced by the nonlinearity and spatial variability of the BOLD response.

Figures 1 and 2 show results from a simulation of this scenario. Five-minute BOLD time series were simulated for two regions of interest X and Y , with 15 trials spaced 20 sec apart. The neural event for Y was delayed relative to that for X by a value between 0 ms and 5 seconds. Gaussian noise was added to make the ratio of signal standard deviation to noise standard deviation equal to 6. Figures 1 and 2 show the Granger causality difference measure $F(X \rightarrow Y) - F(Y \rightarrow X)$ that expresses the degree of temporal precedence of the signal in X as captured by the bivariate AR model. As expected, when the neural events in X and Y were simultaneous, the causality difference was zero. As Y 's onset delay increased, the causality difference increased as well, up to an onset delay around 2 seconds. Beyond that, the measure decreased, indicating a nonlinear relationship and a decreasing sensitivity of the Granger causality to temporal differences that were too large. Furthermore, a faster sampling rate gave better sensitivity to small onset differences, but even typical whole-brain fMRI sampling rates of 1-2 seconds allowed measurement of sub-TR temporal dynamics.

We have observed this effect in fMRI data by manipulating the timing of visual stimuli. By presenting a flashing checkerboard pattern to the left visual hemifield, then to the right visual hemifield a short time later, BOLD responses with known relative onset are produced in right and left visual cortex. We have calculated the Granger causality difference measure $F(X \rightarrow$

$Y) - F(Y \rightarrow X)$ between regions of interest in right and left V1 at two relative onsets (0 ms and 112 ms). This measure was zero when the responses are simultaneous and became positive when the right V1 response preceded the left V1 response. Figure 3 shows measurements and 95% bootstrap confidence intervals (see Statistical Inference, below) in two individuals.

One concern that bears mention here is the issue of slice timing. fMRI data are typically acquired using 2D EPI sequences, where each slice is scanned at a different time 50-100 ms apart. Since the MAR model is capable of measuring these timing differences, they must be accounted for. Retrospective temporal interpolation (slice timing correction) may be effective, as may some means of accounting for these delays in the model [19]. Alternatively, using a single EPI slice or a volumetric technique like PRESTO could provide high temporal resolution and obviate slice timing problems, possibly at the cost of being limited in coverage.

Statistical Inference

The most convenient means to estimate the variability of MAR-derived connectivity measures is often to use a non-parametric bootstrap, resampling, or permutation technique. These can be applied to the estimated model parameters like a , b , c , d above, or more practically to summary measures of connectivity such as the Granger causality.

The method of surrogate data may be used to determine if a specific quantity is significantly different from zero [20,21]. Empirical null distributions of a quantity of interest are generated by randomly shuffling each time series independently, which destroys all temporal relationships within and between the signals. The value estimated from the data is then compared with the null distribution.

Some resampling methods retain the statistical properties of the individual time series, including their autocorrelation, while eliminating relationships between variables. These also provide empirical null distributions for inference. One example is the block bootstrap [22], where multiple samples of null data are created from random sub-blocks of each time series. A special case of this that has been applied to Granger causal mapping in fMRI is to swap the first half and last half of the reference or seed region time series and then examine the resulting null distribution of Granger causal measures over all brain voxels [16]. Wavelet-based resampling of the time series data may also be appropriate. Stochastic properties of individual time series may be retained after wavelet resampling [23], possibly to a greater degree than with resampling in the time or frequency domains [24].

Some situations may be amenable to use of the bootstrap to estimate the sampling distribution of a quantity of interest. For example, if a slow event-related design is used so that individual trials can be separated, bootstrap time series of the same length as the original experiment may be created by drawing randomly with replacement from the set of trials. The connectivity measure may be calculated in each bootstrap sample, and the distribution of these values approximates what would be expected if the experiment were actually run many times. This was the method used to generate confidence intervals for the statistic $F(X \rightarrow Y) - F(Y \rightarrow X)$ graphed in Figure 3.

Applying MAR models to fMRI time series

MAR connectivity measures between fMRI time series may be applied in several ways. The Granger causal mapping approach [7] considers a specific seed region of interest, and computes connectivity between it and every brain voxel using a bivariate model in each case. This permits a whole-brain analysis, but is susceptible to concerns about common input and indirect influences because only two time series are included in each model. Alternatively, a number of regions of interest may be included in the same multivariate AR model, and all pairwise

connectivity measures may be calculated [25]. In this case, summary metrics such as the directed transfer function, partial directed coherence, or partial Granger causality (see Extensions, below) permit assigning a single measure of connectivity to each path in the model even when model order is higher than 1. A third possibility is the preliminary use of a source separation technique like independent component analysis, and calculation of Granger causality between the separated components, e.g. [26,27,Londei2006] — each component has an associated time series that reflects the temporal dynamics of an entire hypothetical cognitive network.

An alternative application that sidesteps concerns about the low temporal resolution of fMRI is the use of the technique to study much slower effects; a recent example is the use of Granger causality to measure connectivity in different brain circuits during a motor task, with motor fatigue affecting the signal changes and hence the estimates of connectivity [28]. Since the fatigue process occurs over tens of seconds to minutes, it is easily resolved with fMRI.

INTERPRETATION

MAR model parameters and the Granger causality estimated from hemodynamic data cannot always reflect true underlying neural connectivity. It is important to understand what drives the connectivity estimates — MAR models use variations in the BOLD signals, much as with any other technique for studying connectivity in time series [5]. These may be intrinsic variations such as those that arise during the resting state, e.g. [29], or stimulus-induced responses. MAR-derived connectivity measures are distinct from correlation-based measures because of their sensitivity to temporal delays. In any case, both neural and hemodynamic factors will contribute. If onset or shape of the hemodynamic response differs between two regions, the Granger causality will not reflect the true underlying pattern of neural connectivity. This has been demonstrated in fMRI time series from a rat epilepsy model where spike-and-wave discharges were also measured with intracranial EEG [30]. Nonlinearities in the relationship between neural activity and the BOLD response may also affect the timing and shape [31]. These variations in the hemodynamic response are at least partly of non-neural origin and can significantly confound measurements of connectivity made with MAR models [32]. For this reason, applying an MAR model to a single multivariate fMRI time series may lead to difficulties in interpretation when connectivity values reflect both neural and hemodynamic effects to unknown degrees.

If an MAR model is preferred over a biologically realistic but more complex model such as the dynamic causal model for fMRI [6], it could be applied to neural time series estimated via deconvolution of the hemodynamic data [30]. Explicit measurements of hemodynamic response latency may be made with hypercapnia [33] or hypocapnia [34] paradigms, and these have been used to adjust measures of connectivity [33]. Another alternative may be the use of spin echo pulse sequences at high field to reduce nonlinear and delayed signals from macrovasculature [35].

This concern must be balanced against the demonstrated ability of a basic Granger causal mapping technique to capture known dynamics such as the temporal precedence of auditory cortex responses over motor cortex responses in an auditory-cued reaction time task [17]. Also, recent simulations suggest that the confounding effects of hemodynamic delay can be small under realistic circumstances [36]. If differences in the timing of neural events are much larger than the spatial variation in hemodynamic properties, the latter contribution may become unimportant. Multimodal studies of the BOLD response may clarify this, e.g. [37].

EXTENSIONS AND REFINEMENTS

Frequency domain representation

Multivariate autoregressive modeling may also be performed in the frequency domain to decompose signal relationships into their frequency components. BOLD signals may contain features of little interest which may be distinguished by their frequency content, such as cardiac and respiratory signals or scanner-related artifacts. The frequency domain analysis reveals interactions between channels at specific frequency bands and may allow a more thorough understanding of signal relationships than summary metrics like the Granger causality described above. The MAR model may be expressed in general form as

$$X(t) = \sum_{k=1}^p A_k X(t-k) + E(t)$$

where X is a multivariate time series containing the signals of all channels, A_k contains the model coefficients relating channels at lag k , E is a random innovation (noise), and p is the model order. In the frequency domain,

$$X_f(f) = H(f) E_f(f)$$

where H is the transfer matrix of the model:

$$H(f) = A^{-1}(f)$$

$$A(f) = I - \sum_{k=1}^p A_k e^{-i2\pi f k}$$

Pairwise measures of the relationships between channels cannot keep proper account of the influences when more than two channels are present, unless all variables are incorporated [38]. This can be accomplished in a multivariate situation using the directed transfer function (DTF). DTF is an effective frequency domain measure for inherently multivariate neuronal networks [39]. The non-normalized DTF between channel i and channel j is

$$\theta_{ij}^2(f) = |H_{ij}(f)|^2$$

which is equivalent to the spectral Granger causality [20]. The normalized DTF

$$\gamma_{ij}^2(f) = \frac{|H_{ij}(f)|^2}{\sum_{m=1}^M |H_{im}(f)|^2}$$

where M is the number of channels, describes the fraction of inflow to channel i that is accounted for by the input j . A normalized DTF value of 0 signifies the absence of direct influence, and a value of 1 indicates that channel j is the only direct influence on channel i .

The direct causal relations between signals can also be represented by a metric called partial directed coherence (PDC) [40], which has been recently discussed in the context of fMRI [41]. A normalized version that is independent of the time series scaling, the generalized PDC

(GPDC), is defined in terms of the AR coefficients in the frequency domain and the variance σ_i^2 of the innovation of channel i :

$$\pi_{ij}(f) = \frac{A_{ij}(f) \frac{1}{\sigma_i}}{\sqrt{\sum_{m=1}^M \frac{1}{\sigma_m^2} |A_{mj}(f)|^2}}$$

This measures the outflow from channel j to channel i as a fraction of all outflows from channel j . A PDC value of 0 signifies the absence of direct influence, and a value of 1 indicates that the only direct output of channel j is to channel i [41].

DTF and coherence measures appear to perform similarly in some practical situations for detecting patterns of connectivity, e.g. [42-44]. However, in principle the PDC is capable of separating direct from indirect influences [41], while deeper examination of the AR model parameters is needed to accomplish this in the case of DTF [20].

Conditional and partial Granger causality

The Granger causality describes the relationships between only two signals X and Y (though X and Y may themselves be multivariate time series). For fMRI data, this typically means two voxels or regions of interest. However, most neural systems will naturally consist of more than two, meaning that these bivariate Granger causality measures will be affected by unmodeled common input or indirect influences. The MAR model itself can describe multiple signals. To take advantage of this, the conditional Granger causality or partial Granger causality may be used. Conditional Granger causality [45,46] summarizes the relationship between X and Y , accounting for a third time series Z which may include signals from additional brain regions. The partial Granger causality [47,48] is based on the concept of partial covariance and attempts to account more completely for exogenous inputs. In the presence of common-mode noise, more sophisticated strategies may be effective [49].

Large data sets and data reduction

Models with more than a few signals can have too many parameters to estimate reliably, or at all. The number of free parameters in an unconstrained p -order AR model with n variables is pn^2 , while the number of available observations is tn with t the number of time points. This is intractable in whole-brain fMRI data sets with tens of thousands of voxels but only hundreds of time points, for example. Even modeling signals from a few hundred voxels becomes computationally demanding.

However, a simple average over an entire region of interest can be problematic. A multi-voxel average may include signals that are meaningfully different; different voxels within a single ROI can reliably distinguish different stimuli [50]. An a priori region also is likely to include delayed macrovascular signals when gradient echo EPI pulse sequences are used.

A simple solution is to manually remove voxels with excessively large or delayed responses, or voxels in obvious veins [9]. Voxel pairs with low values of zero-lag correlation may reflect very delayed signals from veins as well [16]. Modeling techniques that include all ROI time series but avoid averaging may reduce this problem [51]. Another possible approach is a data-reduction step in cases when large numbers of voxels are of interest. Principal component decomposition has been proposed for this exact situation in fMRI [52,53]. Other data reduction techniques may be effective as well; self-organizing maps have been applied to spike train data for this purpose, for instance [54]. The use of ICA described previously is another form of this

approach, except that components are separated on the basis of spatial sparsity [55] instead of temporal patterns.

Granger causality as a distance metric for clustering

The Granger causality may also be useful as a distance metric for clustering time series. For instance, this might be applied to time series from a multi-voxel region of interest to separate the early microvascular signals that are co-located with neural responses from delayed macrovascular signals in more distant veins, thereby improving the spatial specificity of fMRI activation maps. While standard Euclidean distance, correlation coefficient [56], or spatio-temporal methods [57] have some utility when responses are distant in time, they are less effective when there is little difference in the shape of the hemodynamics and the signals can be distinguished based solely upon relative onset. The Granger causality difference measure allows small differences in the latency of the BOLD responses to be discerned [58].

In a simulation to demonstrate this, we created 10 time series sampled at 4 Hz, each containing 16 repetitions of a canonical hemodynamic response plus Gaussian noise. Then 10 more were created, with the onsets of the responses delayed 100 ms relative to the first group. Additional groups were created with 200 ms and 300 ms onset delays. Hierarchical clustering was used to group the 40 time series into 4 clusters. Hierarchical clustering uses pairwise distances to determine which time series are similar to each other. We repeated the clustering three times with different distance metrics: the Euclidean distance, the correlation coefficient, and the Granger causality difference measure. The Granger causality difference measure was more accurate at classifying the time series correctly by their onset latency, with a larger advantage at higher SNR (Figure 4).

Relaxing model assumptions

The model assumes the signals are covariance stationary; in other words, that second-order statistical properties of the time series are fixed. One practical implication of this is that different experimental conditions must be in different fMRI runs, or in non-overlapping segments as in a slow event-related design. It also implies that connectivity during the course of a trial does not change; this may not always be reasonable. An alternative would be to expand the model with terms allowing connectivity to change with context, as is done with dynamic causal modeling [6]. The assumption of stationarity may be relaxed, allowing the estimation of time-varying connectivity, using reasonable parameterizations [59], methods that require only local stationarity [60], adaptive MAR models [42] or other means to permit time variation of the model parameters [61].

Another characteristic of the MAR model as described is its linearity. While linear models often capture important features of signals, in some cases nonlinear features may be of interest. The transfer entropy is an information-based metric that captures linear and non-linear dependence between two signals [62, Vakorin2009]. Similar to the Granger causality, it describes information flow bidirectionally, and can be calculated conditional on known common inputs. A locally linear expansion of the Granger causality [63] or a kernel expansion [64] can accommodate some nonlinear situations. There is a class of nonlinear models from which Granger causality measures may be derived from the prediction errors, with a convenient formulation in terms of radial basis functions [65]. The partial Granger causality described previously may also be extended to nonlinear situations [48]. Gourevitch and colleagues [66] have reviewed a number of linear and nonlinear causality metrics with empirical comparisons.

CONCLUSIONS

Autoregressive models are easy to apply and the Granger causality is straightforward to calculate. They provide an efficient and reasonable means to study the temporal information in hemodynamic data. But since these connectivity measures do not always reflect the true neuronal connectivity, careful planning is needed to ensure that they will be meaningful. This could include experimental manipulations intended to affect the Granger causality, use of deconvolution approaches to approximate neural time series, or use of pulse sequences that minimize delayed macrovascular signals. In general, we expect this methodology can be quite useful to extend the information obtainable using fMRI if appropriate care is taken.

Acknowledgments

This work was supported in part by NIH 5R01EB000461 (Principal Investigator JCG), NIH 5R01NS055822 (Principal Investigator VLM), and 5R03EY014437 (Principal Investigator R. Randolph Blake).

REFERENCES

1. Vincent JL, et al. Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *J Neurophysiol* 2008;100(6):3328–42. [PubMed: 18799601]
2. Raichle ME, Snyder AZ. A default mode of brain function: a brief history of an evolving idea. *Neuroimage* 2007;37(4):1083–90. discussion 1097–9. [PubMed: 17719799]
3. Shehzad Z, et al. The resting brain: unconstrained yet reliable. *Cereb Cortex* 2009;19(10):2209–29. [PubMed: 19221144]
4. Horwitz B. The elusive concept of brain connectivity. *Neuroimage* 2003;19(2 Pt 1):466–70. [PubMed: 12814595]
5. Rogers BP, et al. Assessing functional connectivity in the human brain by fMRI. *Magn Reson Imaging* 2007;25(10):1347–57. [PubMed: 17499467]
6. Friston KJ, Harrison L, Penny W. Dynamic causal modelling. *Neuroimage* 2003;19(4):1273–302. [PubMed: 12948688]
7. Goebel R, et al. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. *Magn Reson Imaging* 2003;21(10):1251–61. [PubMed: 14725933]
8. Roebroeck A, Formisano E, Goebel R. The identification of interacting networks in the brain using fMRI: Model selection, causality and deconvolution. *Neuroimage*. 2009
9. Menon RS, Luknowsky DC, Gati JS. Mental chronometry using latency-resolved functional MRI. *Proc Natl Acad Sci U S A* 1998;95(18):10902–7. [PubMed: 9724802]
10. Bandettini P. Functional MRI today. *Int J Psychophysiol* 2007;63(2):138–45. [PubMed: 16842871]
11. Bellgowan PS, Saad ZS, Bandettini PA. Understanding neural system dynamics through task modulation and measurement of functional MRI amplitude, latency, and width. *Proc Natl Acad Sci U S A* 2003;100(3):1415–9. [PubMed: 12552093]
12. Henson RN, Price CJ, Rugg MD, Turner R, Friston KJ. Detecting latency differences in event-related BOLD responses: application to words versus nonwords and initial versus repeated face presentations. *Neuroimage* Jan;2002 15(1):83–97. PubMed PMID: 11771976. [PubMed: 11771976]
13. Sigman M, et al. Parsing a sequence of brain activations at psychological times using fMRI. *Neuroimage* 2007;35(2):655–68. [PubMed: 17275341]
14. Geweke JF. Measurement of linear dependence and feedback between multiple time series. *J Am Stat Assoc* 1982;77:304–313.
15. Granger CWJ. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 1969;37:424–438.
16. Roebroeck A, Formisano E, Goebel R. Mapping directed influence over the brain using Granger causality and fMRI. *Neuroimage* 2005;25(1):230–42. [PubMed: 15734358]
17. Abler B, et al. Investigating directed influences between activated brain areas in a motor-response task using fMRI. *Magn Reson Imaging* 2006;24(2):181–5. [PubMed: 16455407]

18. Schumacher EH, Elston PA, D'Esposito M. Neural evidence for representation-specific response selection. *J Cogn Neurosci* 2003;15(8):1111–21. [PubMed: 14709230]
19. Schelter B, et al. Inference of causal interactions in fMRI data: The challenge of slice timing. *Neuroimage* 2009;47(Suppl 1):S124.
20. Kaminski M, et al. Evaluating causal relations in neural systems: granger causality, directed transfer function and statistical assessment of significance. *Biol Cybern* 2001;85(2):145–57. [PubMed: 11508777]
21. Theiler J, et al. Testing for nonlinearity in time series: The method of surrogate data. *Physica D* 1992;58:77–94.
22. Davison, AC.; Hinckley, DV. *Bootstrap methods and their application*. Cambridge University Press; 1997.
23. Bullmore E, et al. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum Brain Mapp* 2001;12(2):61–78. [PubMed: 11169871]
24. Laird AR, Rogers BP, Meyerand ME. Comparison of Fourier and wavelet resampling methods. *Magn Reson Med* 2004;51(2):418–22. [PubMed: 14755671]
25. Harrison L, Penny WD, Friston K. Multivariate autoregressive modeling of fMRI time series. *Neuroimage* 2003;19(4):1477–91. [PubMed: 12948704]
26. Demirci O, et al. Investigation of relationships between fMRI brain networks in the spectral domain using ICA and Granger causality reveals distinct differences between schizophrenia patients and healthy controls. *Neuroimage* 2009;46(2):419–31. [PubMed: 19245841]
27. Liao W, et al. Evaluating the effective connectivity of resting state networks using conditional Granger causality. *Biol Cybern* 102(1):57–69. [PubMed: 19937337]
28. Londei A, D'Ausilio A, Basso D, Belardinelli MO. A new method for detecting causality in fMRI data of cognitive processing. *Cogn Process* 2006;7(1):42–52. PMID 16628465. [PubMed: 16628465]
28. Deshpande G, et al. Multivariate Granger causality analysis of fMRI data. *Hum Brain Mapp* 2009;30(4):1361–73. [PubMed: 18537116]
29. Stevens MC, Pearson GD, Calhoun VD. Changes in the interaction of resting-state neural networks from adolescence to adulthood. *Hum Brain Mapp* 2009;30(8):2356–66. [PubMed: 19172655]
30. David O, et al. Identifying neural drivers with functional MRI: an electrophysiological validation. *PLoS Biol* 2008;6(12):2683–97. [PubMed: 19108604]
31. de Zwart JA, et al. Hemodynamic nonlinearities affect BOLD fMRI response timing and amplitude. *Neuroimage* 2009;47(4):1649–58. [PubMed: 19520175]
32. Friston K. Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS Biol* 2009;7(2):e33. [PubMed: 19226186]
33. Chang C, Thomason ME, Glover GH. Mapping and correction of vascular hemodynamic latency in the BOLD signal. *Neuroimage* 2008;43(1):90–102. [PubMed: 18656545]
34. Bright MG, et al. Characterization of regional heterogeneity in cerebrovascular reactivity dynamics using novel hypocapnia task and BOLD fMRI. *Neuroimage* 2009;48(1):166–75. [PubMed: 19450694]
35. Zhang N, et al. Linearity of blood-oxygenation-level dependent signal at microvasculature. *Neuroimage* 2009;48(2):313–8. [PubMed: 19580875]
36. Deshpande G, Sathian K, Hu X. Effect of hemodynamic variability on Granger causality analysis of fMRI. *Neuroimage*. 2009
37. Liu Z, et al. Linear and nonlinear relationships between visual stimuli, EEG and BOLD fMRI signals. *Neuroimage*. 2010 In press.
38. Kus R, Kaminski M, Blinowska KJ. Determination of EEG activity propagation: pair-wise versus multichannel estimate. *IEEE Trans Biomed Eng* 2004;51(9):1501–10. [PubMed: 15376498]
39. Kaminski MJ, Blinowska KJ. A new method of the description of the information flow in the brain structures. *Biol Cybern* 1991;65(3):203–10. [PubMed: 1912013]
40. Baccala LA, Sameshima K. Partial directed coherence: a new concept in neural structure determination. *Biol Cybern* 2001;84(6):463–74. [PubMed: 11417058]

41. Sato JR, et al. Frequency domain connectivity identification: an application of partial directed coherence in fMRI. *Hum Brain Mapp* 2009;30(2):452–61. [PubMed: 18064582]
42. Astolfi L, et al. Tracking the time-varying cortical connectivity patterns by adaptive multivariate estimators. *IEEE Trans Biomed Eng* 2008;55(3):902–13. [PubMed: 18334381]
43. Astolfi L, et al. Estimation of effective and functional cortical connectivity from neuroelectric and hemodynamic recordings. *IEEE Trans Neural Syst Rehabil Eng* 2009;17(3):224–33. [PubMed: 19273037]
44. Kayser AS, Sun FT, D'Esposito M. A comparison of Granger causality and coherency in fMRI -based analysis of the motor system. *Hum Brain Mapp* 2009;30(11):3475–94. [PubMed: 19387980]
45. 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(2):228–37. [PubMed: 16099512]
46. Geweke JF. Measures of conditional linear dependence and feedback between time series. *J Am Stat Assoc* 1984;79:907–915.
47. Guo S, et al. Partial Granger causality--eliminating exogenous inputs and latent variables. *J Neurosci Methods* 2008;172(1):79–93. [PubMed: 18508128]
48. Guo S, et al. Uncovering interactions in the frequency domain. *PLoS Comput Biol* 2008;4(5):e1000087. [PubMed: 18516243]
49. Nalatore H, Ding M, Rangarajan G. Mitigating the effects of measurement noise on Granger causality. *Phys Rev E Stat Nonlin Soft Matter Phys* 2007;75(3 Pt 1):031123. [PubMed: 17500684]
50. Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 2005;8(5):679–85. [PubMed: 15852014]
51. Wang X, et al. Granger causality between multiple interdependent neurobiological time series: blockwise versus pairwise methods. *Int J Neural Syst* 2007;17(2):71–8. [PubMed: 17565503]
52. Zhou Z, et al. Analyzing brain networks with PCA and conditional Granger causality. *Hum Brain Mapp* 2009;30(7):2197–206. [PubMed: 18830956]
53. Zhou Z, et al. Detecting directional influence in fMRI connectivity analysis using PCA based Granger causality. *Brain Res* 2009;1289:22–9. [PubMed: 19595679]
54. Passaro, P., et al. Spike sorting and functional connectivity analysis using self-organizing maps and Granger causality. 1st INCF Congress of Neuroinformatics; Stockholm, Sweden. 2008;
55. Daubechies I, et al. Independent component analysis for brain fMRI does not select for independence. *Proc Natl Acad Sci U S A* 2009;106(26):10415–22. [PubMed: 19556548]
56. Cordes D, et al. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn Reson Imaging* 2002;20(4):305–17. [PubMed: 12165349]
57. Liao W, et al. Analysis of fMRI data using improved self-organizing mapping and spatio-temporal metric hierarchical clustering. *IEEE Trans Med Imaging* 2008;27(10):1472–83. [PubMed: 18815099]
58. Katwal, SB., et al. International Society for Magnetic Resonance in Medicine. Honolulu, HI: 2009. Minimum resolvable latency difference of BOLD responses at 7T using autoregressive modeling.
59. Thompson WK, Siegle G. A stimulus-locked vector autoregressive model for slow event-related fMRI designs. *Neuroimage* 2009;46(3):739–48. [PubMed: 19236927]
60. Sato JR, et al. 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(1):187–96. [PubMed: 16434214]
61. Hemmelmann D, et al. Modelling and analysis of time-variant directed interrelations between brain regions based on BOLD-signals. *Neuroimage* 2009;45(3):722–37. [PubMed: 19280694]
62. Schreiber T. Measuring information transfer. *Phys Rev Lett* 2000;85(2):461–4. [PubMed: 10991308]
64. Vakorin VA, Krakovska OA, McIntosh AR. Confounding effects of indirect connections on causality estimation. *J Neurosci Methods* 2009;184(1):152–60. PMID 19628006. [PubMed: 19628006]
63. Chen Y, et al. Analyzing multiple nonlinear time series with extended Granger causality. *Phys Lett A* 2004;324(1):26.35.
64. Liao W, et al. Kernel Granger causality mapping effective connectivity on FMRI data. *IEEE Trans Med Imaging* 2009;28(11):1825–35. [PubMed: 19709972]

65. Ancona N, Marinazzo D, Stramaglia S. Radial basis function approach to nonlinear Granger causality of time series. *Phys Rev E Stat Nonlin Soft Matter Phys* 2004;70(5 Pt 2):056221. [PubMed: 15600742]
66. Gourevitch B, Bouquin-Jeannes RL, Faucon G. Linear and nonlinear causality between signals: methods, examples and neurophysiological applications. *Biol Cybern* 2006;95(4):349–69. [PubMed: 16927098]

Figure 1.

For small differences in hemodynamic response onset, the Granger causality difference is an increasing function of the onset difference. The graph shows 100 realizations of simulated 5-minute time series (95% of values fell within the error bars). Rapid sampling (250 ms TR) permitted detection of onset differences down to 50 ms at this SNR, while longer sampling times reduced sensitivity.

Figure 2.

At longer onset delays, the relationship between onset delay and Granger causality difference between simulated time series is no longer monotonic. Within this range of parameters, sensitivity appeared to peak when the onset delay was approximately matched to the sampling time.

Figure 3.

The Granger causality difference can measure short onset delays in the visual system. In this example, stimuli were presented to left hemifield, then right hemifield, with either 0 ms or 112 ms delay, during fMRI scanning at 7T. The Granger causality difference between right and left primary visual cortex is shown for two volunteers. As expected, it was zero when the induced fMRI responses were simultaneous. It was positive when the left hemisphere response was delayed, indicating measureable temporal precedence of the right hemisphere response. Error bars indicate 95% confidence intervals based on bootstrap samples of individual trials.

Figure 4.

The Granger causality difference outperforms standard distance measures for clustering time series with signals that have small relative delays. The graph shows classification accuracy for a group of signals containing simulated fMRI responses at 0, 100, 200, and 300 ms relative onset. Chance accuracy was 25%, and both the Euclidean distance and the correlation performed poorly compared to the Granger causality difference in the context of hierarchical clustering.