Title:

A Deconvolution-based Approach to Identifying Large-Scale Effective Connectivity

Authors:

Keith Bush^a, Suijian Zhou^a, Josh Cisler^b, Jiang Bian^c, Onder Hazaroglu^a, Keenan Gillispie^a, Kenji Yoshigoe^a, Clint Kilts^b

Affiliations:

^aDepartment of Computer Science, University of Arkansas at Little Rock (UALR), 2801 S. University Ave., Little Rock, Arkansas, USA 72204

^bBrain Imaging Research Center, University of Arkansas for Medical Sciences (UAMS), 4301 W. Markham St., Little Rock, Arkansas, USA 72205

^cDepartment of Health Outcomes and Policy, University of Florida (UF), 1329 SW 16th St., Gainesville, Florida, USA 32608

E-mail Addresses:

kabush@ualr.edu, sxzhou@ualr.edu, jcisler@uams.edu, bianjiang@ufl.edu, oxhazaroglu@ualr.edu, kxgillispie@ualr.edu, kxyoshigoe@ualr.edu, cdkilts@uams.edu

Corresponding Author:

Keith Bush, Ph.D.
Department of Computer Science
University of Arkansas at Little Rock
2801 S. University Ave.
Little Rock, AR 72204

E-mail: kabush@ualr.edu
Phone: 501.569.8143
Fax: 501.569.8144

Abstract

Rapid, robust computation of effective connectivity between neural regions is an important next step in characterizing the brain's organization, particularly in the resting state. However, recent work has called into question the value of causal inference computed directly from BOLD, demonstrating that valid inferences require transformation of the BOLD signal into its underlying neural events as necessary for accurate causal inference. In this work we develop an approach for effective connectivity estimation directly from deconvolution-based features and estimates of inter-regional communication lag. We then test, both in simulation as well as whole-brain fMRI BOLD signal, the viability of this approach. Our results show that deconvolution precision and network size play outsized roles in effective connectivity estimation performance. Idealized simulation conditions allow for statistically significant effective connectivity estimation of networks of up to four hundred regions-of-interest (ROI). Under simulation of realistic recording conditions and deconvolution performance, however, our result indicate that effective connectivity is viable in networks containing up to approximately sixty ROIs. We then validated the ability for the proposed method to reliably detect effective connectivity in whole-brain fMRI signal parcellated into networks of viable size.

1. Introduction

An open question in neuroimaging is whether whole-brain resting state functional magnetic resonance imaging (fMRI) blood-oxygen-level dependent (BOLD) signal contains sufficient information to consistently and accurately identify causal relationships between brain regions. Currently, large scale functional brain organization analyses, such as independent component analyses and graph theory analyses [1,2] rely primarily on non-causal measures, i.e. functional connectivity. Given the common understanding that the human brain operates as a distributed network of information processing [3], it is essential that we develop methods that accurately characterize causal connectivity networks, i.e., effective connectivity [4], in order to foster an understanding of both normative and pathological brain processes. Moreover, addressing this open question is a crucial step toward unifying functional brain organization with theoretical [5] and physiological [6] understanding of neural systems, which is deeply rooted in networks of causal information processing elements.

fMRI captures neuronal activity dependent changes in local magnetic fields [7,8,9]. Neuronal activity induces an influx of oxygenated hemoglobin molecules into the region of activity which alters the ratio between oxygenated and deoxygenated hemoglobin molecules in the local blood supply [10]; due to oxygen's role in masking the magnetic field of hemoglobin, this changing ratio alters the local magnetic field surrounding the neural activity and is captured as BOLD. The causal relationship between neural activation and BOLD contrast is well-approximated by the hemodynamic response function (HRF) [11,12]. In the HRF model, a BOLD response is formed in proportion to the neural activation according to a nonlinear kernel function, well-approximated by a double gamma distribution, shown in Figure 1a. The HRF model also assumes that kernels are linearly additive; neural events occurring in close temporal proximity produce a BOLD response that combines the individual HRF kernels, depicted in

Figure 1b. Noise processes (both physiological and thermal) confound real-world fMRI BOLD signal acquisition [13,14], as depicted in Figure 1c.

[Figure 1 approximately here]

It is well-known that there exists significant temporal variability (± 2 s) in the HRF kernel's time-to-peak across different regions of the brain [15,16,17,18]. Thus, there exists conflicting conclusions in the most recent literature on the viability of causal modeling based on fMRI BOLD. Smith et al. [19] used simulations of small scale brain activation networks to conduct a comparative survey of algorithms that estimate functional connectivity and effective connectivity directly from BOLD. However, this work cited, but failed to address, earlier work by David et al. [20] which called into question the value of causal inference computed directly from BOLD [20,21]. By pairing high-temporal fidelity intracerebral EEG recordings with fMRI recordings of the same brain regions, David et al. convincingly demonstrated that region-wise variations in HRF confounded the discovery of the system's key causal relationship directly from BOLD. Rather, identification of the BOLD signal's underlying HRF (in this case, through temporal comparison of EEG and fMRI BOLD) followed by explicit deconvolution was found to be a necessary condition for identifying the true, directed brain organization.

Causal inference from whole-brain resting state fMRI BOLD also raises concerns about algorithmic scale-up. Progress has been made toward solving causal inference on small-scale neural systems via dynamic causal modeling (DCM), a form of Bayesian model comparison which estimates the causal structure of coupled dynamical systems [22], as well as structural equation modeling (SEM) [23,24,25], Granger causal modeling (GCM) [26,27], and autoregressive methods [26,28]. There are, however, computational limits to scaling DCM to whole-brain as it is not feasible to search model spaces of this magnitude [19], a weakness also of structural equation modeling (SEM). Also, external system inputs are unknown or undefined for causal models of resting state data. Similar computational scale-up limitations have been reported for autoregressive approaches [28].

The purpose of this work is to explore an effective connectivity estimation approach that scales to whole-brain fMRI BOLD signals via massively parallelized maximum likelihood estimation of structured neural activation models. A key innovation of this work is to avoid the computational complexity of existing approaches by separating the neural activity identification step (i.e. deconvolution) from the effective connectivity identification step. This two-step process greatly reduces the size of the search space, making whole-brain causal structure identification tractable within the constraints of existing hardware. Separation of these search steps also facilitates independent verification of the effective connectivity estimation via controlled simulation experiments in which neural activity is known. This validation technique avoids questions surrounding the validity of causal inferences under hemodynamic variability [21].

The remainder of this manuscript details the connectivity estimation algorithm as well as supporting algorithms and implementation techniques required for its scale to whole brain; a simulation of whole-brain neural activation; the collection and subsequent deconvolution of real-world whole-brain fMRI BOLD signal; the effective connectivity estimation algorithm; and, the analysis methods by which effective connectivity was validated both in simulation and on real-world, human data.

2. Materials and Methods

2.1 Effective Connectivity Estimation

Our approach to effective connectivity estimation incorporates anatomical constraints to the problem [29,30] via a spatio-temporally structured model of neural activations underlying fMRI BOLD, as depicted in Figure 2. The proposed estimation method requires two inputs: 1) a dataset of estimated whole-brain neural activations (i.e., deconvolved fMRI BOLD), and 2) estimates of intra-voxel communication lag. The quantity $L(v_i, v_j)$ describes the time required for a neural activation in voxel, v_i , to influence voxel, v_j . Using these inputs, the modeling approach estimates the probability that an activation in v_i causes activation in v_j , defined as $p(v_j|v_i)$, by minimizing the cost function, J , the sum of voxel-wise squared errors between the whole-brain activations described by the model and the observed whole-brain activations. We achieve this via backtracking line-search gradient descent along a path minimizing the cost function with respect to the causal probabilities $\partial J / \partial p(v_j | v_i)$ for all v_i and v_j . The computational complexity of this approach is $O(TV^2/P)$ calculations per iteration of gradient descent where T is the number of time points of the signal, V is the number of voxels, and P is the number of available processors. Details of the optimization are provided in Appendix A. Analysis of the computational complexity is provided in the Appendix B.

[Figure 2 approximately here]

2.2 Simulation of the Functional Network and Neural Event Generation

We employ a parametric model of the functional network and neural event generation process to test the efficacy of the connectivity estimation approach against a known ground truth. The mathematical model used here is identical to the functional network and neural event generation processes described in earlier work [31]. Parameters of this model are given in Appendix C.

2.3 Simulation of Deconvolution Error

Hemodynamic deconvolution of fMRI BOLD signal is a poorly determined problem. A recent survey of deconvolution algorithms [32] found that classification performance, using deconvolved features of realistically confounded fMRI BOLD signal, achieves on the order of 0.6–0.7 area-under-the-curve (AUC) of the receiver-operator characteristic (ROC) curve. Using additional machine learning techniques can improve performance into the range 0.7–0.8 AUC but at the expense of removing time points from the classification problem as unknown (a knows-what-it-knows approach) [33]. Thus, deconvolution yields temporal identification of neural events that is significantly better than random but also contains error due to numerous confounding factors, notably noise and low sample-rate (high TR).

As effective connectivity estimation is directly linked to the error embedded in the underlying deconvolved neural events on which it relies, deconvolution error must be incorporated into effective connectivity performance analysis. We propose that the error in the estimation of the probability of a neural event at a given time by a deconvolution algorithm is well-approximated by a Gaussian distribution having a mean value equal to the neural event's true time and a variance dictated by the quality of the algorithm.

To simulate deconvolution error in our analysis we convolve the neural events generated via simulation (see Section 2.2) with a deconvolution error kernel, which we construct via discrete approximation to the Gaussian distribution (parameterized by both neural generate rate, FG (Appendix C) in units Hz, and σ , in units s). The kernel is bounded (truncated) at a width of $\pm 3\sigma$.¹ Thus, by varying the kernel's 3σ parameter the resulting convolved time-series approximates the neural events that would be estimated via deconvolution of fMRI BOLD signal using an algorithm with 3σ -precision. For clarity, we demonstrate the consequences of 3σ -precision on deconvolution performance in Figure 3.

[Figure 3: approximately here]

2.4 Algorithm Performance Analysis for Simulated Neural Events

We evaluated the performance of the effective connectivity estimation algorithm effectiveness according to AUC of the ROC curve calculated from classification of estimated conditional probabilities of effective connectivity. Performance was measured as follows. For each trial, we fit the effective connectivity estimate, $\tilde{p}(v_j | v_i)$, for all voxels i and j . We then classified effective connectivity, $b(j, i)$, for all voxels $i, j : i \neq j$,

$$b(j, i) = \begin{cases} 1 : & \tilde{p}(v_j | v_i) > \gamma \\ 0 : & \text{otherwise} \end{cases} \quad [1]$$

Where γ is the *ROC-threshold* parameter. By comparing the detected connectivity matrix against the true connectivity matrix, $p(v_j | v_i)$, the specificity and sensitivity of the detected effective connectivities was calculated for each threshold value, $\gamma \in [0, 1]$ sampled at intervals of 0.01, generating a 101-point ROC curve. Note, one ROC curve was generated for each trial. Using the trapezoid rule to numerically integrate the AUC, we computed the distribution of AUCs achieved over S i.i.d. samples. This distribution constitutes the performance of the estimation algorithm for the respective experimental parameters: random classification performance achieves AUC=0.5 and ideal classification performance achieves AUC=1.0.

2.5 Human Data Collection

We tested performance on real fMRI BOLD data collected among 17 healthy adults (9 female; mean age = 31.7; SD = 9.5) undergoing resting-state fMRI. Participants were included based on the absence of current mental health diagnoses, major medical conditions, or MRI contraindications (e.g., internal ferrous objects). The resting-state task was a 7.5 min scan during which participants were presented a fixation cross and told to lie still and not think about anything specific.

We used a 3T Achieva X-series small bore magnet using an 8-channel head coil (Philips Healthcare, USA) to acquire imaging data. We acquired anatomic images with a MPRAGE sequence (matrix=192x192, 160 slices, TR/TE/FA=min/min/90°, final resolution=1x1x1mm³ resolution). The echo planar imaging sequences used to collect the functional images were:

¹ We construct the kernel using Matlab's fspecial filter generation function.

TR/TE/FA=2000ms/30ms/90⁰, FOV=192x192mm, matrix=64x64, 34 oblique slices (parallel to AC-PC plane to minimize OFC signal artifact), slice thickness=3mm, final resolution 3x3x3 mm³.

We preprocessed images using AFNI software [34] and followed standard steps. In order, images underwent despiking, slice timing correction, deobliquing, motion correction using rigid body alignment, alignment to participant's normalized anatomical images, spatial smoothing using a 6 mm FWHM Gaussian filter, temporal bandpass filtering (.01–.1 Hz), and scaling into percent signal change. We normalized images using the ICBM 452 template brain. Additionally, to correct for respiratory and cardiovascular artifacts, we regressed fluctuations in white matter voxels and CSF out of time courses from grey matter voxels following segmentation using FSL [35] and used restricted maximum likelihood to account for autocorrelation. This step was implemented directly after motion correction and normalization of the EPI images in the preprocessing stream. We inspected the resulting images for each participant for artifacts and accuracy of normalization [36].

2.6 Data parcellation method

We used a previously published spatially constrained spectral clustering method [2] to construct variable-granularity parcellations of whole-brain fMRI BOLD. We then extracted the principal component time series of the voxels within each ROI through singular value decomposition (AFNI's 3dmaskSVD) [34]. We focused specifically on the posterior cingulate cortex (PCC), medial prefrontal cortex (mPFC), and right dorsolateral prefrontal cortex (dlPFC). Prior research suggests causal relationships between the PCC and mPFC [37], but not between the PCC and dlPFC [38]; thus, focusing analyses on identified effective connectivity between these regions provides a means of comparing the observed results to that which would be expected theoretically (significant PCC to mPFC effective connectivity, null PCC to dlPFC effective connectivity).

2.7 BOLD Deconvolution

We used the deconvolution algorithm of Bush and Cisler [32], in all experiments. For application to whole-brain human images we utilized a massively parallelized implementation based upon C++ threads [39] and Armadillo C++ linear algebra library [40]. For all real-world estimation experiments the parallelized code was executed (using a thread pool size of 60) on a Hewlett Packard ProLiant DL980 G7 Server (80 processors and 4TB single-addressable memory).

2.8 Estimation of Communication Lag

The range of neural communication velocities were computed as follows. First, we computed pairwise distances between all parcellations (using the mean coordinates of the regions as the endpoints of the segment being measured). From this set we selected the maximum distance. Second, we estimated the range of inter-brain communication velocity by combining the maximum lag distance with an estimated range of maximum temporal delays [21] having the form [maximum distance/maximum lag, maximum distance/minimum lag]. The resulting velocity range was [8.38,16.76] m/s. We can then generate a random lag model by uniformly, randomly sampling from the velocity range and multiplying this value with the pairwise distances.

We acknowledge that there exist numerous sources of error in computing these values including voxel registration errors [41,42,43,44], complex inter-neuron and inter-regional neural

pathways [45,46], neural conduction velocity [47,48,49,50,51], and physical connectivity [52,53]. Uniform, random sampling over the estimated velocity range is the mechanism by which we marginalized the impact of these factors in effective connectivity estimation.

2.9 Analysis of Human Effective Connectivity Estimation

Estimates of effective connectivity between each node pair for each participant were calculated as described above. Given prior research suggesting positive effective connectivity between PCC and mPFC, $\tilde{p}(v_{\text{mPFC}}|v_{\text{PCC}})$ and null effective connectivity between PCC and dlPFC, $\tilde{p}(v_{\text{dlPFC}}|v_{\text{PCC}})$, we focused group-level analyses on these regions. We compared with paired samples t-tests the mean effective connectivity estimates for these two node pair relationships against the mean degree of neural activation given PCC activity for all nodes $\tilde{p}(\cdot|v_{\text{PCC}})$. This comparison provides a test of the degree to which the hypothesized relationship (e.g., PCC to mPFC) is significantly greater than the mean PCC effective connectivity estimates for all other nodes.

3. Simulation Experiments

Our simulation experiments explore the performance and robustness of the effective connectivity estimation approach to various confounding factors, specifically, network size, deconvolution performance, and uncertainty of inter-voxel communication lag. Then, using the simulation results as parametric guidance, we applied the approach to estimate effective connectivity in human resting-state fMRI BOLD signal.

3.1 Effective Connectivity Estimation of Simulated Neural Events

We simulated networks of size $V \in \{5, 10, 20, 30, 40, 50, 75, 100, 200, 250, 500, 1000\}$. The total simulation time was 200 s. Neural events were generated at 20 Hz (i.e., each simulation step is 50 ms). For all ROIs, the probability of a neural event caused by an external (unmodeled) factor is $\rho = .05$. Time-lags between all nodes i and j , $i \neq j$, were uniformly, randomly sampled on the range [10, 50] ms. For each V -sized network, V non-zero conditional probabilities $p(v_j | v_i) = .982$ were chosen such that $i \neq j$.² Deconvolution performance was varied over 3σ -precision $\in \{75, 1525, 3375, 7525\}$ ms. Each combination of network size and deconvolution performance was sampled, independently identically $S=30$ times (thus, neural events, lag, and network connectivity vary for each trial).

4. Simulation Results

A summary of the performance and computational complexity of the effective connectivity approach (applied in simulation) is presented in Figure 4. As would be expected, estimation performance is negatively impacted by network size and deconvolution precision. The upper bound of statistically significant estimation is $V=500$, which is achieved only in the case of near

² For numerical purposes concerning simulation optimization, our simulation builds conditional probabilities via the sigmoid function, i.e., logistic function. Thus, causal link between nodes have the value $\text{sigmoid}(4) = .982$ and all other links is the connectivity model contain the value $\text{sigmoid}(-20) = 2.061\text{E-}9$

perfect deconvolution performance. Significant estimation performance is achievable for realistic deconvolution precisions for network sizes on the range $V \in [5, 60]$, approximately four times larger than has been detected in the prior literature [19].

Computational complexity grows quadratically with network size, and, as expected, increases inversely proportional to deconvolution precision (lack of precision introduces error in the optimization problem, which slow the convergence rate). The nonlinear structure of the plot occurring at $V \leq 250$ ROIs is likely an artifact of overhead produced by the specific hardware used in this experiment (see Section 2.7). Moreover, as the approach scales strictly linearly with the number of processors, we predict real-world convergence run-times of approximately $12/V$ s for existing deconvolution algorithms on network sizes V in $[5, 60]$ where simulation results have been shown to exhibit significant effective connectivity estimation performance.

[Figure 4: approximately here]

5. Whole-brain Resting State BOLD Experiments

Guided by the simulation results, each of the 17 subjects' whole-brain fMRI BOLD images were parcellated into ROI networks of $V \in [10, 60]$ sampled at intervals of 10 ROIs. In accordance with the procedures described in Sections 2.6–2.7, BOLD signals were extracted from the parcellations and deconvolved. Thirty ($S=30$) unique lag models were sampled by uniformly, randomly sampling neural communication velocities from the range calculated in Section 2. We fit effective connectivity maps, $\tilde{p}(v_j | v_i)$, for all voxels i and j , and computed the average connectivity for each subject.

[Figure 5: approximately here]

The parcellation indexes (i.e., ROI number) associated with the PCC, mPFC, and dIPFC were identified for each subject for each parcellation size, V (footnote: conducted by manually scanning parcellation indexes in AFNI). Two such parcellations (sizes 10- and 60-ROIs) with associated PCC, mPFC, and dIPFC labels are presented in Figure 5. From these indices we identified the mean $\tilde{p}(v_{\text{mPFC}} | v_{\text{PCC}})$ and mean $\tilde{p}(v_{\text{dIPFC}} | v_{\text{PCC}})$ values as well as the mean $\tilde{p}(v_{\text{PCC}} | v_{\text{PCC}})$. The comparison between these inter-subject effective connectivity estimates is shown in Figure 6. Consistent with hypotheses, our algorithm detects significantly stronger effective connectivity from the PCC to mPFC and significantly weaker effective connectivity from the PCC to dIPFC.

6. Discussion

We have described an effective connectivity approach, based on deconvolution and anatomical lag constraints, that is designed for scale to whole-brain fMRI BOLD. A key insight of this approach is that it decouples the identification of neural events occurring within an ROI from the estimation of the conditional probabilities by which ROIs communicate. This decoupling serves two roles. First, it dramatically reduces the search space of possible solutions. Second, it

provides insight into a potential source of algorithm failure, the quality of the neural event timings by which causality is based. Indeed, we know of no previous work in which the issue of neural event timing error has been discussed as a variable in effective connectivity estimation. Rather, the debate has centered on the fundamental utility of deconvolution.

The results here are decidedly split. Precise neural event identification leads to statistically significant effective connectivity estimation for extremely large networks (up to 400 ROIs). Such results, however, require neural event identification at a level of precision that is currently unavailable. Existing deconvolution algorithms, even when assuming idealized imaging conditions, yield neural event estimation precision that largely obviates the role of deconvolution (see Figure 4: 3σ -precisions induce effective connectivity estimations that are, largely, invariant to the role of deconvolution). Excellent effective connectivity performance is possible up to approximately 20 ROIs and statistically significant effective connectivity can be identified up to 60 ROIs. These scales are between 2 and 5 times larger than network sizes previously reported [19], however, they are far below a desired fine-grained scale of network connectivity (the 1000 ROI atlas [2]). Indeed, simulation results suggest that achievement of effective connectivity at this scale is not feasible without making additional assumptions about network structure, a fruitful area future exploration. It is possible, even likely, that there exists exploitable network structure at the ROI level. There are, however, enormous challenges in validating such methods. Ground truth of network level structure at the scale desired is unavailable; and, imaging technology that could conclusively identify such structure using existing methods does not yet scale to whole-brain.

Within the range of network sizes for which we expect effective connectivity estimation to work, results from human experiments are encouraging. At each ROI parcellation size, we observed significant effective connectivity between PCC and mPFC and null or negative effective connectivity between PCC and dlPFC, which is consistent with what would be expected based on prior default mode research [38,39]. Accordingly, these results provide evidence for the validity of this proposed effective connectivity methodology on real-world data. One consideration here is the high TR (2s) of the data in our real-world analysis. We could expect performance improvements in our estimations as TR decreases (shorter TRs, e.g., 1s are commonly used in practice) as this 1) correlates with significantly improved deconvolution performance [32] and 2) increases the number of datapoints on which the conditional probability is estimated for a fixed signal length. Nonetheless, in the absence of ground truth on real-world data, these results remain speculative. Future research along these lines may benefit from obtaining resting-state fMRI and DTI within the same participants, allowing mapping of structural connectivity pathways with which to compare estimates of the effective connectivity.

Acknowledgements

This work was supported in part by National Institutes of Health grants R21MH097784-01 and R01DA036360-01 as well as by the National Science Foundation grants CRI CNS-0855248 and MRI ACI-1429160.

References

1. McKeown M, Jung T, Makeig S, Brown G, Kindermann S, Lee T, Sejnowski T. Spatially independent activity patterns in functional MRI data during the stroop color-naming task. *Proc Natl Acad Sci*. 1998;95(3):803–810.
2. Craddock R, James G, Holtzheimer P, Xiaoping P, Mayberg H. A Whole brain fMRI Atlas generated via spatially constrained clustering. *Human Brain Mapping* 2012;33(8):1914–28.
3. Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews* 2009;10:186–98.
4. Friston K. Functional and effective connectivity in neuroimaging: A Synthesis. *Human Brain Mapping* 1994;2:56–78.
5. Dayan P, Abbott LF. Theoretical neuroscience: computational and mathematical modeling of neural systems, 2001. Cambridge, MA. MIT Press, 460 pp.
6. Kandel E, Schwartz J, Jessell T, Siegelbaum S, Hudspeth A. Principles of neural science (5th Edition), 2012. New York, NY. McGraw-Hill, 1760 pp.
7. Logothetis N, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 2001;412(6843):150–7.
8. Logothetis N. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci* 2003;23(10):3963–71.
9. Logothetis N. What we can do and what we cannot do with fMRI. *Nature* 2008;453(7197):869–78.
10. Ogawa S, Lee T, Kay A, Tank D. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci USA* 1990;87(24):9868–72.
11. Buxton R, Wong E, Frank L. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magn Reson Med* 1998; 39(6):855–64.
12. Friston K, Ashburner J, Kiebel, Nichols T, Penny W. Statistical Parametric Mapping: The Analysis of Functional Brain Images 2006. Elsevier, London. Academic Press.
13. Purdon P, Weisskoff R. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Human Brain Mapping* 1998; 6:239–49.
14. Lindquist M. The statistical analysis of fMRI data. *Statistical Science* 2008; 23(4):434–64.
15. Saad Z, Ropella K, Cox R, DeYoe E. Analysis and use fMRI response delays. *Human Brain Mapping* 2001;13(2): 74–93.
16. Huettel S, McCarthy G. Regional differences in the refractory period of the hemodynamic response: an event-related fMRI study. *NeuroImage* 2001;14(5):967–76.
17. Miezin F, Maccotta L, Ollinger J, Petersen S, Buckner R. Characterizing the Hemodynamic Response: Effect of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage* 2000;11:735–59.
18. Lee A, Glover G, Meyer C. Discrimination of Large Venous Vessels in Time-Course Spiral Blood-Oxygen-Level-Dependent Magnetic Resonance Functional Neuroimaging. *Magn Reson Med* 1995;33:745–54.
19. Smith S, Miller K, Salimi-Khorshidi G, Webster M, Beckmann C, Nichols E, Ramsey J, Woolrich M. Network modelling methods for fMRI. *Neuroimage* 2011; 54(2):875–91.

20. David O, Guillemain I, Sallet S, Reyt S, Deransart C, Segebarth C, Depaulis A. Identifying neural drivers with functional MRI: An electrophysiological validation. *PLOS Biology* 2008;6(12):2683–97.
21. Stephan K, Roebroeck A. A short history of causal modeling of fMRI data. *NeuroImage* 2012;62:856–63.
22. Friston K. Dynamic causal modeling. *NeuroImage* 2003;19:1273–302.
23. Buchel C, Friston K. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling an fMRI. *Cereb Cortex* 1997;8:768–78.
24. Bullmore E, Horwitz B, Honey G, Brammer M, Williams S, Sharma T. How good is good enough in path analysis of fMRI data? *NeuroImage* 2000;11:289–301.
25. Horwitz B, Tagamets M, McIntosh A. Neural modeling, functional brain imaging, and cognition. *Trends Cogn Sci* 1999;3:91–8.
26. Goebel R, Roebroeck A, Kim D, Formisano E. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. *Magn Reson Imaging* 2003;21:1251–61.
27. Roebroeck A, Formisano E, Goebel R. Mapping directed influence over the brain using Granger causality and fMRI. *Neuroimage* 2005;25(1):230–42.
28. Valdes-Sosa P. Spatio-temporal autoregressive models defined over brain manifolds. *Neuroinformatics* 2004;2:239–50.
29. Fonteijn H, Norris D, Verstraten F. Exploring the anatomical basis of effective connectivity models with DTI-based fiber tractography. *Int J Biomed Imaging* 2008;423192.
30. Sanchez-Bornot J, Martinez-Montes E, Lage-Castellanos A, Vega-Hernandez M, Valdes-Sosa P. Uncovering Sparse Brain Effective Connectivity: A Voxel-based Approach Using Penalized Regression. *Statistica Sinica* 2008;18:1501–18.
31. Bush K, Cisler J. Deconvolution Filtering: Temporal Smoothing Revisited. *Mag Reson Imaging* 2014;32(6):721–35.
32. Bush K, Cisler J. Decoding neural events from fMRI BOLD signal: A comparison of existing approaches and development of a new algorithm. *Magn Reson Imaging* 2013;29(3):353–64.
33. Bush K, Cisler J, Hazaroglu G, Hazaroglu O, Kilts C. Improving the Precision of fMRI BOLD Signal Deconvolution with Implications for Connectivity Analysis (in press), 2015.
34. Cox R. AFNI: Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages. *Computers and Biomedical Research* 1996;29(3): 162–73.
35. Smith S, Jenkinson M, Woolrich M, Beckmann C, Behrens T, Johansen-Berg H, Bannister P, De Luca M, Drobnjak I, Flitney D, Niazy R, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady J, Matthews P. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* 2004; 23(S1):S208–19.
36. Kelly RE, Alexopoulos G, Wang Z, Gunning F, Murphy C, Morimoto S. Visual inspection of independent components: defining a procedure for artifact removal from fMRI data. *J Neurosci Methods* 2010; 189(2): 233–45.
37. Greicius M, Supekar K, Menon V, Dougherty R. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cereb.Cortex* 2009;19(1):72–8.

38. Fox M, Snyder A, Vincent J, Corbetta M, Van Essen D, Raichle M. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *P Natl Acad Sci USA* 2005;102:9673–8.
39. Stallman R et al. Using the GNU Compiler Collection. GNU Press, Boston, MA, 2013.
40. Sanderson C. Armadillo: An Open Source C++ Linear Algebra Library for Fast Prototyping and Computationally Intensive Experiments. National ICT Australia (NICTA). St. Lucia, Queensland, Australia. Technical Report, 2010.
41. Ashburner J, Friston K. Multimodal image coregistration and partitioning—a unified framework. *NeuroImage* 1997;6:209–17.
42. Woods R, Grafton S, Holmes C, Cherry S, Mazziotta J. Automated image registration. I. General methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography* 1998;22:139–52.
43. Nestares O, Heeger D. Robust multiresolution alignment of MRI brain volumes. *Magnetic Resonance Medicine* 2000;43:705–15.
44. Reuter M, Rosas H, Fischl B. Highly accurate inverse consistent registration: a robust approach. *NeuroImage* 2010;53(4):1181–96.
45. Van Essen D, Smith S, Barch D, Behrens T, Yacoub E, Ugurbil K, for the WU-Minn HCP Consortium. The WU-Minn Human Connectome Project: An overview. *NeuroImage* 2013;80(2013):62-79.
46. Sporns O. *Networks of the Brain*, 2010. Cambridge, MA, MIT Press, 424 pp.
47. Weiss T. *Cellular Biophysics*, Vol 2, 1996. Cambridge, MA, MIT Press, 450 pp.
48. Waxman S. Determinants of conduction velocity in myelinated nerve fibers. *Muscle Nerve* 1980;3:141–50.
49. Waxman S. Action potential propagation and conduction velocity: new perspectives and questions. *Trends in Neuroscience* 1983;6:157–61.
50. Eccles J, Llinas R, Sasaki K. Parallel fiber stimulation and the responses induced thereby in the Purkinje cells of the cerebellum. *Experimental Brain Research* 1966;1:17–39.
51. Hodgkin A. *The conduction of the nervous impulse*, 1964. Cambridge, UK. Cambridge Univ Press, 108 pp.
52. Iturria-Medina Y, Canales-Rodriguez E, Melie-Garcia L, Valdes-Hernandes P, Martinez-Mondes E, Aleman-Gomez Y, Sanchez-Bornot J. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *NeuroImage* 2007;36(3):645–60.
53. Iturria-Medina Y, Sotero R, Canales-Rodriguez E, Aleman-Gomez Y, Melie-Garcia. Studying the human brain anatomical network via diffusion-weighted MRI and graphy theory. *NeuroImage* 2008;40(3):1064–76.

Appendix A

Via deconvolution, an observed BOLD signal for a voxel j at time t is transformed into the probability of a neural event generated in voxel j at time t denoted, $p(v_{j,t})$. According to a probabilistic causal model of this system, if we ignore communication time constraints then:

$$p(v_j) = \sum_{v_i} p(v_j, v_i) \quad [\text{A.1}]$$

where the joint probability of a neural event at voxel i coinciding with neural event at voxel j is marginalized over all voxels, i . We decompose this expression into an explicitly causal representation via the product rule,

$$p(v_j) = \sum_{v_i} p(v_j|v_i)p(v_i). \quad [\text{A.2}]$$

Physics, however, requires acknowledgement that causality is not instantaneous; there exists a lag in the influence between voxel i and voxel j , which we denote as the function $L(i, j)$ for which the values are provided in anatomically plausible ranges. Therefore, we may rewrite the causal model using a more rigorous temporal notation,

$$p(v_{j,t}) = \sum_{v_{i,t-L(i,j)}} p(v_{j,t}|v_{i,t-L(i,j)})p(v_{i,t-L(i,j)}) \quad [\text{A.3}]$$

The values $p(v_{j,t})$ and $p(v_{i,t})$ are known via deconvolution of fMRI BOLD signal. We propose that the true causal parameters, $p(v_{j,t}|v_{i,t-L(i,j)})$, i.e., the true conditional probabilities of neural communication, may be estimated as $\tilde{p}(v_{j,t}|v_{i,t-L(i,j)})$ through gradient descent optimization by defining the cost function,

$$J_j = \sum_t \frac{1}{2} [\tilde{p}(v_{j,t}) - p(v_{j,t})]^2. \quad [\text{A.4}]$$

We model the conditional probability as driven by some underlying activation parameter, $\gamma_{i,j}$, via the logistic transfer function,

$$\tilde{p}(v_{j,t}|v_{i,t-L(i,j)}) = [1 + \exp(-\gamma_{i,j})]^{-1}. \quad [\text{A.5}]$$

We differentiate J_j with respect to activation parameters using the chain rule,

$$\frac{\partial J}{\partial \gamma_{i,j}} = \frac{\partial J}{\partial \tilde{p}(v_{j,t})} \frac{\partial \tilde{p}(v_{j,t})}{\partial p(v_{j,t}|v_{i,t-L(i,j)})} \frac{\partial p(v_{j,t}|v_{i,t-L(i,j)})}{\partial \gamma_{i,j}}, \quad [\text{A.6}]$$

and apply a gradient descent to solve for parameters, $\gamma_{i,j}$,

$$\gamma_{i,j} = \gamma_{i,j} - \alpha \frac{\partial J}{\partial \gamma_{i,j}}|_{\gamma}, \quad [\text{A.7}]$$

yielding a locally cost-optimal approximation of whole-brain causal connectivity.

Appendix B

Gradient descent optimization of whole-brain effective connectivity scales efficiently via parallelization. Assuming V is the number of voxels, gradient optimization requires $O(V^2)$ unique

derivative calculations per iteration of descent. Each derivative requires $O(T)$ calculations as $\partial J / \partial p(v_j | v_i)$ is calculated for each of T observations. Thus, the algorithm's iteration complexity is $O(TV^2)$. The computational structure of this problem is highly parallelizable; each set of parameters $p(\cdot | v_i)$ determining the causality of voxel, v_j , is independent. Thus, when the algorithm is parallelized over P processing nodes, each node calculates V/P sets of parameters yielding $O(TV^2/P)$ calculations per iteration. Each parallelized process stores $O(TV+V^2/P)$ values (the whole-brain deconvolved fMRI BOLD signal as well as the condition probability values being optimized). Initial broadcast requires $O(TVP)$ communication and final collection requires $O(V^2)$ communication.

Appendix C

The default fMRI BOLD signal simulation [31], used in this work, employed parameters described in Table 1.

[Table 1: approximately here]

Figure Captions

Figure 1: Theoretical BOLD signal: **(a)** canonical hemodynamic response to a single neural event; **(b)** simulated BOLD signal formed via the linearly additive composition of HRFs in response to a sequence of neural events; **(c)** simulated observed fMRI BOLD signal confounded by autocorrelated physiological noise, white thermal noise, normalization, and low-fidelity observations. Note, neural events are plotted as vertical bars for temporal reference; they are not to scale.

Figure 2: A graphical depiction of the proposed anatomically-constrained causal model, configured in this example for the conditional probabilities describing the neural activity of voxel 3, v_3 (depicted in red). Prior probabilities $p(\cdot)$ of neural activation are known for each time step (estimated via deconvolution and depicted in blue). Conditional probabilities, $p(v_3 | \cdot)$ are calculated via gradient descent (depicted as black arrows. Temporal lags, $L(\cdot, \cdot)$, indicated by horizontal dotted lines, are determined by combining both physiological and anatomical prior knowledge.

Figure 3: A graphical depiction of the impacts of a deconvolution algorithm's 3σ -precision on estimating the timing of neural events from fMRI BOLD: $3\sigma=75$ ms exhibits close to perfect binary accuracy of the true neural events whereas $3\sigma=18$ exhibits significant loss of temporal structure. Existing deconvolution algorithms operating on ideally recorded fMRI BOLD signal (minimal motion artifacts, negligible autocorellation, $TR=1$, and hemodynamics close to the canonical HRF) are estimated to achieve 3σ -precision on the range 1,525–3,375 ms.

Figure 4: Effective connectivity estimation and computational complexity on simulated neural systems: **(a)** Effective connectivity performance, measured as AUC, with respect to network size

and simulated 3σ -precision of the deconvolution algorithm; the dashed lines represent the 95% confidence interval of a random classifier (classifying the true connections with randomly assigned connections); **(b)** quadratic complexity growth exhibited by the effective connectivity estimation algorithm with respect to network size and 3σ -precision of the deconvolution algorithm, see (a) for legend, executed on a single process (i.e. sequential).

Figure 5: Visual depiction of posterior cingulate (PCC), medial prefrontal (mPFC) and dorsolateral prefrontal (dlPFC) regions across the 10- and 60- ROI parcellations. Regions are depicted as sagittal views (at MNI coordinates $x=0$ and $x=28$) overlaid atop the ICMB_452 MRI atlas.

Figure 6: Comparison of the estimated effective connectivity from posterior cingulate (PCC) to medial prefrontal (mPFC), $\tilde{p}(v_{\text{mPFC}} | v_{\text{PCC}})$, and from PCC to dorsolateral prefrontal (dlPFC), $\tilde{p}(v_{\text{dlPFC}} | v_{\text{PCC}})$ in comparison to background PCC connectivity across varying ROI atlas granularities. Probabilities are plotted with respect to each ROI parcellation on the range [10,60] sampled at intervals of 10. Conditional probabilities are reported as means and standard errors.

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

10 ROI parcellation 60 ROI parcellation

Figure 6

Table 1

Parameter	Value	Description
TS	200 s	Simulation length.
V	See Section 3.1	The number of simulated regions-of-interest (ROI).
C	Sparse V -connected: a set of V ROIs were uniformly, randomly assigned non-zero probabilities.	Connectivity model, a matrix (size $V \times V$) of the conditional probabilities. Each element, $C(i,j)$ determines the probability of a neural event in region j caused by a neural event in region i .
L	See Section 3.1	Lag model, a matrix (size $V \times V$) of communication times between simulated ROIs. Each element, $L(i,j)$ determines the temporal delay of a neural event in region j occurring due to a neural event in region i .
ρ	0.05	External activity model, a V -sized vector containing the latent probability for each ROI of the model to fire based on unmodeled, external influences.
FG	20 Hz	Simulated generation frequency. $1/FG$ represents the smallest unit of time modeled and is assumed to be the span of time in which a single neural event occurs.
FO	1 Hz	Simulated observation frequency; equivalent to $1/TR$.