CANONICAL GRANGER CAUSALITY APPLIED TO FUNCTIONAL BRAIN DATA

Syed Ashrafulla, Justin P Haldar, Anand A Joshi, Richard M Leahy

Signal and Image Processing Institute, University of Southern California, Los Angeles, CA 90089

ABSTRACT

Dynamic images of functional activity in the brain offer the potential to measure connectivity between regions of interest. We want to measure causal activity between regions of interest (ROIs) with signals recorded from multiple channels or voxels in each ROI. Previous methods, such as Granger causality, look for causality between individual time series; hence, they suffer from local interactions or interferers obscuring signals of interest between two ROIs. We propose a metric that reduces the effect of interference by taking weighted sums of sensors in each ROI, as is done with canonical correlation. Hence, we measure region-to-region, rather than channel-to-channel or point-to-point, Granger causality. We show in simulation that our "canonical Granger causality" accurately mimics the underlying structure with few samples, unlike current methods of multivariate Granger causality. We then use anatomically relevant regions of interest in a visuomotor task in a multichannel intracortical EEG study to infer the direction of transmission in visual processing.

Index Terms— Intracranial EEG, Connectivity, Causality

1. INTRODUCTION

Our goal is to convert functional images of the brain into descriptions of interactions between the activity at different regions in the brain. We must analyze time series of activity to understand the functional relationships between anatomical structures. We can obtain these time series from any of several modalities (E/MEG, fMRI and depth electrodes are common examples) each with a common property: every recording is an amalgamation of different underlying signals within a region of the brain. We want to infer causal signal interactions between regions of the brain from these time series.

Communications in the brain occur between spatial clusters of neurons; thus, it is prudent to model the data in terms of anatomical ROIs rather than individual readings. We infer transfer of activity between ROIs by fusing multiple time series – from voxels, channels or cortical tessellations – within each ROI to determine whether some signal in an ROI affects any signal in a second ROI. Thus, we focus on a specific network of interest (NOI) in the presence of an interfering background. In addition, we obtain a functional description of the brain directly related to pre-defined ROIs.

We want to find which regions of interest help infer future activity at other regions of interest. Granger causality [1] can measure the effect of the past of one signal on the future of another, but is hard to extend to signals recorded in regions of interest. On the other hand, canonical correlation [2] can measure the relationship between regions of interest but does not find the direction of inference between participating regions.

This work was supported by the NIH under grants R01EB009048, R01EB000473 and the T32EB00438 Training Grant as well as the NSF under grant BCS-1028389.

We propose a framework for estimating inter-regional connections in which we find the maximum bivariate causal connectivity between components of two ROIs, rather than finding multivariate Granger causality [3]. Like other methods of regional Granger causality [4] we find combinations of signals within each ROI that are maximally related. Rather than maximizing residual cross-correlation, we maximize Granger causality between these combinations. Our metric also finds regional connections rather than summing connections between univariate signals [3]. In simulations we compare our proposed metric against multivariate Granger causality. By imposing a model in which only one signal from each ROI participates in the network, we can extract the structure of the NOI with fewer samples than other estimates of Granger causality.

2. BACKGROUND

2.1. Granger causality

For determining information transfer between regions we ask the question, "How much does one region's past behavior affect the behavior of a second region?" We measure this strength of connectivity by Granger causality [1]. Granger causality is the measurement of how the past of one signal x_2 (the "source") is able to predict of the future of another signal x_1 (the "sink"). This influence is measured as the change in prediction error of the sink $x_1[n]$ given only the sink's past $\{x_1[n-k]\}_{k>0}$ versus the prediction error of the sink given its past as well as the source's past $\{x_2[n-k]\}_{k>0}$:

$$\mathcal{G}_{2\to 1} = \frac{\mathrm{Var}\big(x_1[n]\,\big|\{x_1[n-k]\}_{k>0}\big)}{\mathrm{Var}\big(x_1[n]\,\big|\{x_1[n-k],x_2[n-k]\}_{k>0}\big)} \tag{1}$$

If $\mathcal{G}_{2\to 1}$ is significantly greater than 1, then we conclude that the source x_2 causes at least a part of the sink x_1 [1]. We calculate the prediction errors using an autoregressive model.

The multivariate extension [3] to measure causality from one set of signals $\mathbf{y}_2[n] \in \mathbb{R}^{M_1}$ to another set of signals $\mathbf{y}_1[n] \in \mathbb{R}^{M_2}$ requires the use of conditional variance matrices. The relative "size" of these matrices can be calculated via determinants:

$$\mathcal{G}_{2\to 1} = \frac{\big|\mathbf{\Sigma}\big(\mathbf{y}_1[n]\,\big|\{\mathbf{y}_1[n-k]\}_{k>0}\big)\big|}{\big|\mathbf{\Sigma}\big(\mathbf{y}_1[n]\,\big|\{\mathbf{y}_1[n-k],\mathbf{y}_2[n-k]\}_{k>0}\big)\big|} \tag{2}$$

where $\Sigma(\cdot)$ measures the covariance and $|\cdot|$ is the determinant. Using the determinant provides stability and invariance to linear transformations of channels [3]. We find the prediction errors via a multivariate autoregressive model with $O\big((M_1+M_2)^2\big)$ parameters.

2.2. Canonical Correlation

Canonical correlation [2] is one way to determine relationships between sets of channels. Canonical correlation finds the most related

components of $\mathbf{y}_1$ and $\mathbf{y}_2$ by maximizing a correlation measure between projections of each vector:

$$\mathcal{P}_{2\to1}(\boldsymbol{\alpha},\boldsymbol{\beta}) = \frac{\mathrm{Cov}\left(\boldsymbol{\alpha}^T\mathbf{y}_1[n], \boldsymbol{\beta}^T\mathbf{y}_2[n]\right)}{\mathrm{Var}(\boldsymbol{\alpha}^T\mathbf{y}_1)\mathrm{Var}\left(\boldsymbol{\beta}^T\mathbf{y}_2[n]\right)}$$
$$\mathcal{P}_{2\to1} = \max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \mathcal{P}_{2\to1}(\boldsymbol{\alpha},\boldsymbol{\beta}) \quad (3)$$

This measure reflects the assumption that there is only one signal from each group of sensors that is participating in the network. However, it does not reflect the causality between $\mathbf{y}_1$ and $\mathbf{y}_2$.

3. METHODS

3.1. Proposed: Canonical Granger Causality

Granger causality methods suffer from local interferers that increase the within-ROI connectivity and thus drown out between-ROI connectivity. Current methods to aggregate ROI sensors, such as principal or independent components, find the strongest components in a signal, not necessarily the components used within a network. Canonical correlation can reduce the effect of these non-participants by finding components that are most correlated to each other. We want to draw conclusions about inter-regional causality, so we combine these two approaches. Thus, we find signals within one group of sensors whose past is most related to the current signals within a second group of sensors. This method is then more resistant to local interferers (as in canonical correlation) while determining the directed interaction (as in Granger causality) between ROIs.

To do this, we reduce the causality between the two ROIs to a bivariate process. A group of signals in a region of interest ($\mathbf{y}_1$ or $\mathbf{y}_2$) can be broken into a signal (x_1 or x_2) pertinent to the network between the ROIs and interferers that do not enter the network. We aggregate each set of channels, as in section 2.2, while maximizing causality as calculated in section 2.1:

$$\mathcal{C}_{2\to1}(\boldsymbol{\alpha},\boldsymbol{\beta}) = \frac{\mathrm{Var}\left(\boldsymbol{\alpha}^T\mathbf{y}_1[n]\,\middle|\,\boldsymbol{\alpha}^T\mathbf{y}_1[n-1],\dots\right)}{\mathrm{Var}\left(\boldsymbol{\alpha}^T\mathbf{y}_1[n]\,\middle|\,\boldsymbol{\alpha}^T\mathbf{y}_1[n-1],\boldsymbol{\beta}^T\mathbf{y}_2[n-1],\dots\right)}$$
$$\mathcal{C}_{2\to1} = \max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \mathcal{C}_{2\to1}(\boldsymbol{\alpha},\boldsymbol{\beta}) \quad (4)$$

where $\boldsymbol{\alpha} \in \mathbb{R}^{M_1}$ and $\boldsymbol{\beta} \in \mathbb{R}^{M_2}$. When each group of sensors is located in an anatomical ROI, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ can represent how much of the signal is being generated near a sensor in an ROI.

When $M_1 = M_2 = 1$, optimization reduces to bivariate Granger causality. Since the calculation of variances is a bivariate model independent of M_1 and M_2, the number of parameters is $O(M_1 + M_2)$, an order lower than multivariate causality (2). In comparison to very recent independent work on regional causality [4] we maximize causality rather than residual correlation. In the future, we will investigate the relationships between these two metrics.

3.2. Optimization Parameters

The optimization in (4) is non-convex so local maxima must be avoided. To obtain a unique global maximum, we constrain $||\boldsymbol{\alpha}||^2 = ||\boldsymbol{\beta}||^2 = 1$. We solve the maximization in (4) using an interior-point method, which chooses direction first by direct calculation of the Hessian. If the Hessian is ill-conditioned, then the algorithm falls back to linear conjugate gradient descent.

We found effective convergence to a maximum in under 100 iterations for all simulations regardless of initialization. A multistart investigation indicated that the solutions we found were global maxima.

4. RESULTS

4.1. Simulated functional networks

We simulated two regions of activity x_1 and x_2 recorded on two sets of sensors $\mathbf{y}_1$ and $\mathbf{y}_2$ with random gains $\mathbf{g}_1 \in \mathbb{R}^{M_1}$, $\mathbf{g}_2 \in \mathbb{R}^{M_2}$ related by a first-order autoregressive model:

$$\begin{bmatrix} x_1[n] \\ x_2[n] \end{bmatrix} = \begin{bmatrix} 0.7 & A \\ 0 & 0.7 \end{bmatrix} \begin{bmatrix} x_1[n-1] \\ x_2[n-1] \end{bmatrix} + \begin{bmatrix} \eta_1[n] \\ \eta_2[n] \end{bmatrix}$$
$$\mathbf{y}_1[n] = \mathbf{g}_1 x_1[n] + \boldsymbol{\delta}_1[n] \quad (5)$$
$$\mathbf{y}_2[n] = \mathbf{g}_2 x_2[n] + \boldsymbol{\delta}_2[n] \qquad n = 1,\dots,N$$

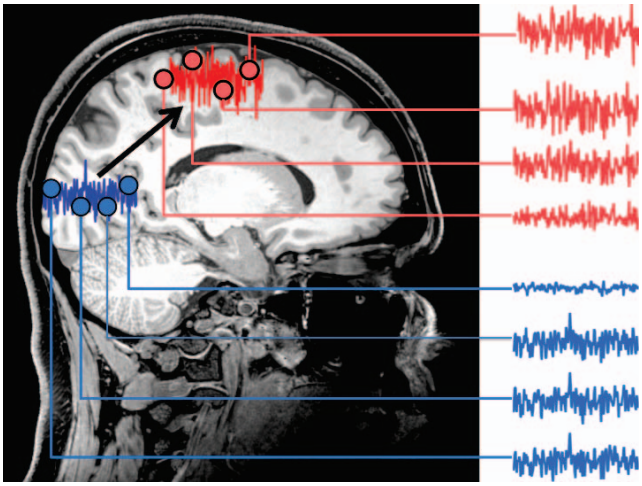


Fig. 1. Model of functionally connected regions. The past of one region – the source – affects the current signal at a second region – the sink. We infer this causal relationship from the sensor readings on the right.

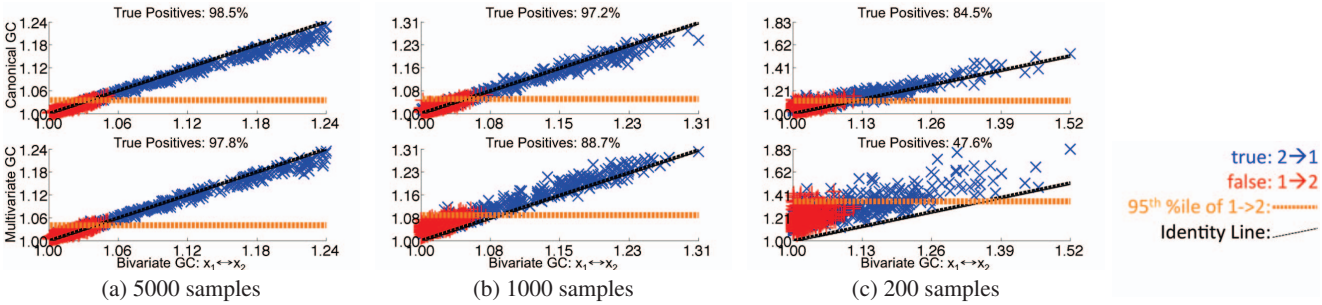
This system, illustrated in Figure 1, models bivariate connectivity between two regions detected by groups of sensors. We simulate 4 channels for each region. Our source innovations η_1, η_2 are i.i.d. normal. We generate simulations by setting $A = 0.2, 0.3, \dots, 0.8, 0.9$ for 50 simulations each. We build $\boldsymbol{\delta}_1$ and $\boldsymbol{\delta}_2$ from region-specific interferers and noise:

$$\boldsymbol{\delta}_1[n] = \sum_{q=1}^{Q_1} \boldsymbol{\delta}_1^q[n] + \boldsymbol{\gamma}_1[n] \qquad \boldsymbol{\delta}_2[n] = \sum_{q=1}^{Q_2} \boldsymbol{\delta}_2^q[n] + \boldsymbol{\gamma}_2[n]$$
$$\boldsymbol{\delta}_1^q[n] = 0.8\,\boldsymbol{\delta}_1^q[n-1] + \boldsymbol{\psi}_1^q[n] \quad \boldsymbol{\delta}_2^q[n] = 0.8\,\boldsymbol{\delta}_2^q[n-1] + \boldsymbol{\psi}_2^q[n]$$

Every innovation $\boldsymbol{\psi}_1^q[n], \boldsymbol{\psi}_2^q[n]$ is standard normal and i.i.d. In addition, $\boldsymbol{\gamma}_1[n]$ and $\boldsymbol{\gamma}_2[n]$ are colored Gaussian noise. We use $Q_1 = Q_2 = 3$ interferers per region and set the signal-to-interference-and-noise ratio (SINR) to 2.5.

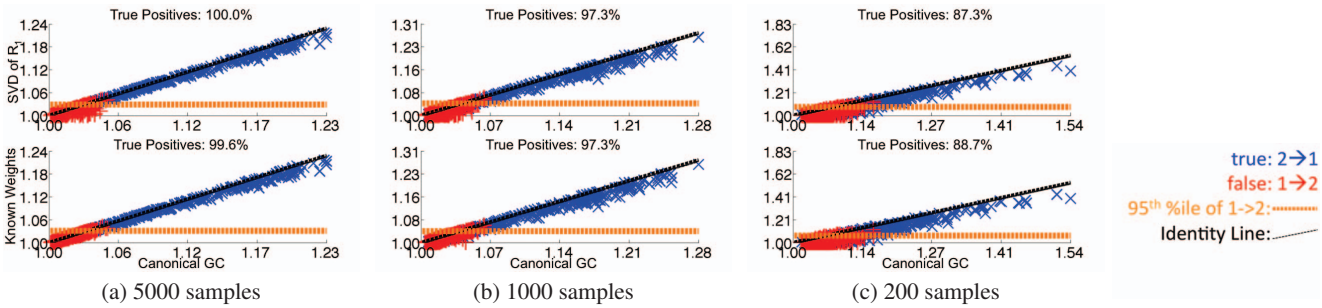
4.2. Comparison to multivariate causality

We expect the two measures of causality between $\mathbf{y}_1$ and $\mathbf{y}_2$ – canonical and multivariate – to each be closely related to the underlying causality between x_1 and x_2. Hence, we use the bivariate causality between x_1 and x_2 as our ground truth. In Figure 2 we compare



(a) 5000 samples (b) 1000 samples (c) 200 samples

Fig. 2. Canonical Granger causality and multivariate Granger causality between $\mathbf{y}_1$ and $\mathbf{y}_2$ plotted against the bivariate Granger causality of x_1 and x_2 in the simulation of section 4.1. We see from the legend that our proposed measure better follows the underlying bivariate causality. In addition, canonical causality has a much higher true positive rate than multivariate causality, whose true positive rate is at the level of chance.



(a) 5000 samples (b) 1000 samples (c) 200 samples

Fig. 3. Weighted bivariate Granger causality using fixed weights – most powerful singular vectors of $\mathbf{R}_{\mathbf{y}_\cdot\mathbf{y}_\cdot^T}[1]$ (*top*) and the known weights $\mathbf{g}_{1,2}$ (*bottom*) – against canonical Granger causality. Using the strongest singular vectors of $\mathbf{R}_{\mathbf{y}_\cdot\mathbf{y}_\cdot^T}[1]$ or the ground truth weights is nearly equivalent to maximizing Granger causality for this model.

canonical Granger causality (4) to multivariate Granger causality (2) in terms of their relationships to the underlying bivariate causality. Multivariate Granger causality tends to over-estimate causality for shorter time series. In contrast, our proposed metric more closely follows the dynamics of the hidden system, as seen by its tighter fit to the dashed black line representing zero bias.

We define a "true positive" as detection of significant causality in the regional metric given that the bivariate causality measure is significant. Then, from Figure 2(c) the true positive rate using $N = 200$ time points with multivariate Granger causality is 47.66% when controlling for a 5% false positive rate in the null case of no interaction. In contrast, canonical Granger causality under the same false positive rate has a true positive rate of 84.6%.

Again using the bivariate causality as the ground truth, the bias and variance of our causality estimator is $32.6 \cdot 10^{-3}$ and $89.6 \cdot 10^{-3}$, respectively. The bias and variance of multivariate causality estimation is $184.1 \cdot 10^{-3}$ and $123.2 \cdot 10^{-3}$, respectively. Thus, our measure has lower bias and variance, as shown in the general trend of Figure 2. Hence when there are fewer samples our measure is more robust and consistent than multivariate Granger causality.

4.3. Lagged cross-covariance in simulation

Based on some assumptions of our simulation, such as no cross-talk, we can solve for the weights $\mathbf{g}_1$ and $\mathbf{g}_2$ another way. We want our measure to be as good without those assumptions. The simulation forces $x_1[n]$ and $x_2[n]$ to be the only signals that co-vary at lag one:

$$\mathbf{R}_{\mathbf{y}_1\mathbf{y}_2^T}[1] = \mathbb{E}\{x_1[t]x_2[t-1]\}\mathbf{g}_1\mathbf{g}_2^T$$

so we can estimate the topography by singular value decomposition

$$\widehat{\mathbf{R}}_{\mathbf{y}_1\mathbf{y}_2^T}[1] = \sum_{n=1}^{N} \lambda_i u_i v_i^T, \lambda_1 > \lambda_2, \ldots, \lambda_N \Rightarrow \widehat{\boldsymbol{\alpha}}_0 = u_1 \quad \widehat{\boldsymbol{\beta}}_0 = v_1 \tag{6}$$

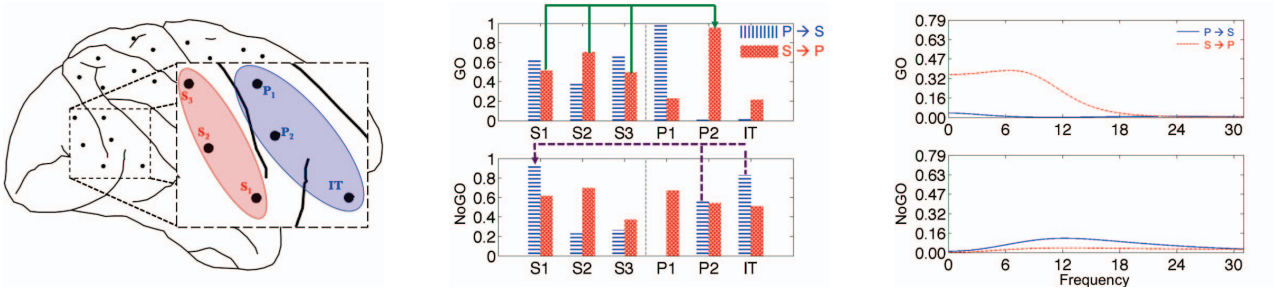
from which we measure causality by calculating $\mathcal{C}_{2\to1}(\widehat{\boldsymbol{\alpha}}_0, \widehat{\boldsymbol{\beta}}_0)$. We want our estimated weights to be nearly as good as the weights calculated from idealized assumptions.

In Figure 3, canonical GC is approximately equal to $\mathcal{C}_{2\to1}(\widehat{\boldsymbol{\alpha}}_0, \widehat{\boldsymbol{\beta}}_0)$ and bivariate causality using the known weights ($\mathcal{C}_{2\to1}(\mathbf{g}_1, \mathbf{g}_2)$). Hence, canonical GC is excellent at determining network topology for an ideal simulation. For fewer samples $\widehat{\mathbf{R}}_{\mathbf{y}_1\mathbf{y}_2^T}[1]$ leads to sub-optimal causality in Figure 3(c).

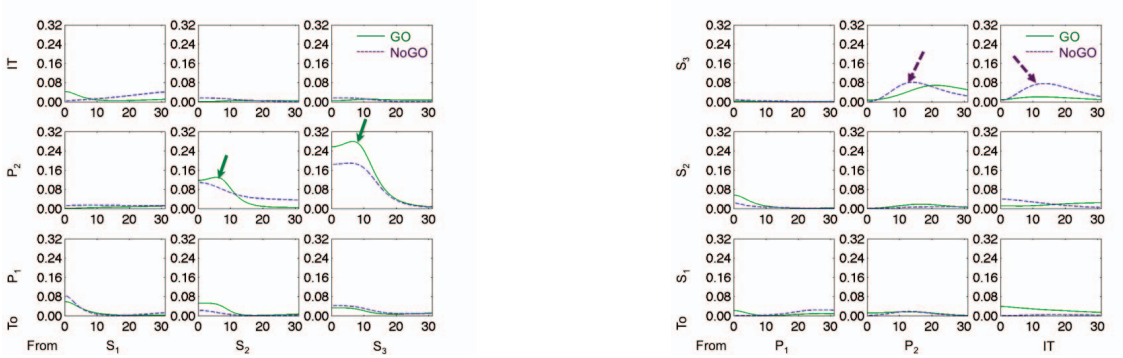
The true positive rate is best when using known weights at any number of samples, and almost as good when using the SVD of $\widehat{\mathbf{R}}_{\mathbf{y}_1\mathbf{y}_2^T}[1]$, a property specific to the model. Our canonical causality measure is almost as good, indicating that it is a consistent measure in the ideal case. We will explore variations of this model, including the effect of cross-talk and different noise models, in future work.

4.4. Application to monkey data

We use intracranial EEG signals recorded on the macaque cortex at six locations – three in striate (S_1, S_2, S_3) and three in prestriate (P_1, P_2, IT) shown in Figure 4(a) – during a visuomotor experiment. Intracranial EEG signals have a small footprint, so we delineate regions of interest anatomically. The visual cue tells the subject whether to perform a motor action ("GO") or refrain from



(a) Locations of sensors in striate and prestriate (b) Estimated topography (c) Frequency-domain canonical GC



(d) Bivariate GC: striate → prestriate (e) Bivariate GC: prestriate → striate

Fig. 4. (a) Locations of local field potential recordings in the striate and prestriate. (b) Weights optimizing canonical Granger causality during $[20, 120]$ms post-stimulus. The arrows show what weights are largest for the two important directions of causality. (c) Frequency-domain Granger causality between ROIs using projections from weights in (b), showing causality from striate to prestriate. (d-e) Bivariate Granger causality between striate and prestriate channels, agreeing with (b) (see arrows).

acting ("NoGO"). We use 137 trials restricted to 20-120ms poststimulus, a time period in which there is previously published evidence of causality from striate to prestriate in this data [5].

In Figure 4(c) our metric shows that striate appears to be driving prestriate during this critical window of visual processing. The topography shown in Figure 4(b) tells us that the entire striate region is sending to a small area in the prestriate around P_2. Figure 4(d) however only shows 2 of 3 striate channels sending information. Hence, S_1 affects the regional connection in a way undetected by bivariate causality. We see a small amount of information transfer from PST to STR during the NoGO condition in Figure 4(c), evidenced in the bivariate measurements of Figure 4(e). Overall, our metric agrees with previous results [5] supporting the idea that the striate region – and not just each sensor – sends visual information to the prestriate.

5. DISCUSSION AND CONCLUSION

Signals participating in a network of ROIs can be obscured by interferers within each ROI. We form a measure of Granger causality that is less affected by interference and time series length than other multivariate methods. With respect to multivariate Granger causality, canonical causality has less variance and more accurate estimation of causal inference. We attribute this to the reduction in the number of parameters as well as the focus on univariate causal relationships.

For the visuomotor experiment, we found a large difference in causality between the GO (act on stimulus) and NoGO (do not act on stimulus) conditions. With a smaller sample size we are still able to find causality from striate to prestriate cortex. The weights used to maximize canonical GC indicate which areas contribute to the network.

We built this measure as a way to calculate Granger causality between regions, similar to how canonical correlation calculates the correlation between regions. We selected anatomical ROIs and investigated possible functional relationships between those ROIs. In future work, we will use either anatomical or functional information to determine ROIs. To find the network, we maximized the Granger causality between projections of channels in each ROI. Canonical causality requires few samples to reliably estimate the topology of a network of ROIs, including which ROI members contribute to the network. We plan to derive optimization methods tailor-made to our metric and constraints, as well as expanding to multiple regions.

6. ACKNOWLEDGEMENTS

We thank Dr. Steven Bressler for providing the macaque data. We thank Dr. Brent Liu for his support of this work through the NIBIB T32EB00438 Bioinformatics Training Program.

7. REFERENCES

[1] C. Granger, "Causality, cointegration, and control," *Journal of Economic Dynamics and Control*, vol. 12, no. 2-3, pp. 551–559, Sep. 1988.

[2] J. Soto, D. Pantazis *et al.*, "Canonical correlation analysis applied to functional connectivity in MEG," in *2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, 2010, pp. 113–116.

[3] A. B. Barrett and A. K. Seth, "Multivariate Granger causality and generalized variance," *Physical Review E*, vol. 81, no. 4, pp. 1–14, Apr. 2010.

[4] G. R. Wu, F. Chen *et al.*, "Multiscale causal connectivity analysis by canonical correlation: theory and application to epileptic brain." *IEEE transactions on biomedical engineering*, vol. 58, no. 11, pp. 3088–96, Nov. 2011.

[5] H. Liang, M. Ding *et al.*, "Causal influences in primate cerebral cortex during visual pattern discrimination." *Neuroreport*, vol. 11, no. 13, pp. 2875–80, Sep. 2000.