Evaluation and calibration of functional network modeling methods based on known anatomical connections

Debra Ann Dawson

The Integrated Program in Neuroscience-MSc

Department of Neurology and Neurosurgery, faculty of Medicine

McGill University

Montreal, Quebec, Canada

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Science.

August 2012

Table of Contents

1- Abstract – English

Recent studies have identified large scale brain networks based on the spatio-temporal structure of spontaneous fluctuations in resting-state fMRI data. It is expected that functional connectivity based on resting-state data is reflective of- but not identical to the underlying anatomical connectivity. However, which functional connectivity analysis methods reliably predict the network structure remains unclear. Here we tested and compared network connectivity analysis methods by applying them to fMRI resting-state time-series obtained from the human visual cortex. The methods evaluated here are those previously tested against simulated data in Smith et al. (Neuroimage, 2011).

To this end, we defined regions within retinotopic visual areas V1, V2, and V3 according to their eccentricity in the visual field, delineating central, intermediate, and peripheral eccentricity regions of interest (ROIs). These ROIs served as nodes in the models we study. We based our evaluation on the 'ground-truth', thoroughly studied retinotopically organized anatomical connectivity in the monkey visual cortex. For each evaluated method, we computed the fractional rate of detecting connections known to exist ('c-sensitivity'), while using a threshold of the 95th percentile of the

distribution of interaction magnitudes of those connections not expected to exist.

Under optimal conditions, including session duration of 68 minutes, a relatively small network consisting of 9 nodes and artifact-free regression of the global effect, each of the top methods predicted the expected connections with 75%-83% c-sensitivity. Partial Correlation performed best (PCorr; 83%), followed by Regularized Inverse Covariance (ICOV; 79%), Bayesian Network methods (BayesNet; 77%), Correlation (75%), and General Synchronization measures (75%). With decreased session duration, these top methods saw decreases in c-sensitivities, achieving 66%-78% and 60%-70% for 34 and 17 minute sessions, respectively. With a short resting-state fMRI scan of 8.5 minutes (TR = 2s), none of the methods predicted the real network well, with ICOV (53%) and PCorr (51%) performing best. With increased complexity of the network from 9 to 36 nodes, multivariate methods including PCorr and BayesNet saw a decrease in performance. However, this decrease became small when using data from a long (68 minutes) session. Artifact-free regression of the global effect significantly increased the c-sensitivity of all top-performing methods. In an overall evaluation across all tests we performed, PCorr, ICOV and BayesNet set themselves somewhat above all other methods.

We propose that data-based calibration based on known anatomical connections be integrated into future network studies, in order to maximize sensitivity and reduce false positives.

2- Abstract – French

Des études récentes ont identifié des réseaux du cerveau à hautes-échelles basés sur la structure spatio-temporelle des fluctuations spontanées dans les donnés IRMf à l'état de repos. Il est prévu que la connectivité fonctionnelle basée sur des donnés de l'état de repos soit reflétée, sans être identique à la connectivité anatomique. Cependant, quelles méthodes analytiques de la connectivité fonctionnelle prédisent de manière fiable la structure du réseau n'est pas encore clair. Ici, nous avons testé et comparé des méthodes analytiques de la connectivité de réseaux en leur appliquant à des séries de temps IRMf de l'état de repos provenant du cortex visuel humain. Les méthodes évaluées ici sont celles qui ont été testées avec des données simulées dans l'article Smith et al. (Neuroimage, 2011).

À cette fin, nous avons défini des sous-régions contenues dans les régions visuelles V1, V2, et V3 selon leur excentricité dans le champ visuel, en créant une région d'intérêt (RI) centrale, intermédiaire, et périphérique. Ces RI ont servi comme nœuds dans les modèles que nous étudions. Nous avons basé notre évaluation sur la « vérité » de l'anatomie du cortex visuel des singes macaques qui a déjà été étudié en profondeur. Pour chaque méthode évaluée, nous avons calculé le taux fractionnel de

la détection des connections que nous savons existantes (‘c-sensitivité’), en utilisant un seuil étant défini comme le 95^{e} percentile de la distribution des magnitudes d’interactions des connections que nous savons inexistantes.

Sous conditions optimales, incluant une longueur de session de 68 minutes, un réseau de 9 nœuds et la régression sans artéfact de l’effet global, les méthodes les plus efficaces ont prédit les connections attendues avec une c-sensitivité de 75%-83%. La corrélation partielle était meilleure (PCorr; 83%), suivi par la Covariance Inverse Régularisé (ICOV; 79%), les méthodes de Réseau Bayesien (BayesNet; 77%), la corrélation (75%), et la Synchronisation Générale (GenSynch; 75%). Avec des sessions plus courtes, les meilleures méthodes ont vus des abaissements en c-sensitivité, réussissant seulement 66%-78% et 60%-70% pour des sessions de 34 et 17 minutes respectivement. Avec des sessions IRMf à l’état de repos courtes de 8.5 minutes (TR=2s), aucune méthode n’a prédit notre réseau avec succès; ICOV (53%) et PCorr (51%) ont été meilleurs. Quand nous avons augmenté la complexité du système en regardant 18 ou 36 nœuds en même temps, les méthodes multi-variantes incluant PCorr et BayesNet ont moins bien réussi qu’avec seulement 9 nœuds. La régression sans artéfact de l’effet global améliore significativement la c-

sensitivité de toutes les meilleures méthodes. Dans une évaluation générale à travers tous nos tests, PCorr, ICOV, et BayesNet se séparent des autres en étant les meilleurs.

Nous proposons qu'une calibration basée sur les donnés utilisant des connections anatomiques connues soit intégrée dans les futures études de réseaux pour maximiser la sensitivité des résultats et réduire la présence de faux positifs dans les réseaux prédits.

3- Preface

Contributions of Authors

Debra Dawson: Major contributions to ideas, analysis, figures, wrote the manuscript

Kuwook Cha: Contributions to ideas, acquired the resting-state data, preprocessing, wrote part of the methods section

Lindsay Lewis: Acquired the retinotopy data, retinotopy methods

Janine Mendola: Acquired the retinotopy data, retinotopy methods

Amir Shmuel: Initiated the study, major contributions to ideas, wrote the manuscript, oversaw the study

Acknowledgements:

We thank Felix Carbonell, Roberto Sotero, Jean Gotman and Bruce Pike for their comments on an earlier version of the manuscript. We thank Ilana Leppert and Bruce Pike for support with MRI physics. This work was supported by NIH grant R01 EY015219 awarded to JM, and grants from the Canadian Institute of Health Research (MOP-102599), the Human Frontier Science Program (RGY0080/2008), and Industry Canada (MNI, Center of Excellence in Commercialization and Research) awarded to AS.

4- Introduction and Literature Review

Advancements in the understanding of the human brain continue every day, and yet we still have a very limited understanding of the structures of the networks it contains. Since the development of nuclear magnetic resonance imaging (MRI) technology in the 1970's by Peter Mansfield and Paul Lauterbur, the potential to learn through imaging of the human brain has exploded. In 1990 Seiji Ogawa et al. originated functional MRI (fMRI) measurements, allowing the examination of a change in signal from the brain over time. It was quickly observed that fMRI activations from networks of regions of the brain are observed in relation to certain tasks or states. A stimulus presented to human subjects evokes increases in brain activity in the areas responsible for response to that stimulus. Modern MRI scanners, which detect changes in blood flow or blood oxygenation, are capable of detecting such increases in brain activity indirectly. Since increases in brain activity draw more blood to the active regions of the brain, these increases in blood flow can be detected and interpreted as increases in brain activity (i.e. fMRI measurements).

The development of this technology led to interest in examining the 'functional connectivity' (i.e. synchronous increases and decreases in

activity) between brain areas in the many task response conditions and brain states one can occupy. Some earlier such studies include a look at human visual cortex responses, namely in V1, with presentation of a patterned-flash photic stimulation (Kwong et al., 1992), imaging of the human primary motor cortex during task activation (Bandettini et al., 1992), and fMRI localization of Broca's area during internal speech word generation (Hinke et al., 1993). More recent studies have led to developments in our understanding of many things from group-ICA results of data collected during cinema viewing (Pamilo et al., 2012) to frontal-amygdala connectivity alterations during emotion down-regulation in Bipolar I Disorder (Townsend et al., 2012).

Studies have also demonstrated changes in fMRI-measured brain activity when the subjects are in a state of rest. Importantly, it has been shown that these spontaneous (i.e. without any external stimulation) fluctuations in fMRI signals demonstrate functional connectivity over large parts of the human brain (Xiong et al., 1999). These large parts of the brain correspond to regions that are active together in response to external stimuli (Cordes et al., 2000). It has been suggested that spontaneous brain activity is an important principle of brain function (Fox and Greicius, 2010).

Here we tested and compared functional connectivity analysis methods by applying them to fMRI resting-state time-series obtained from the human visual cortex. The text to follow will provide an introduction and background on the topics relevant to the study presented in this thesis. We begin with an introduction to functional connectivity and the resting state. Next we introduce the use of network models to analyze functional brain networks using fMRI data, followed by a look at some of the more popular models used to date, and comments on the interpretation of the results of such models. Finally we introduce the human visual cortex and the localization of lower retinotopic visual areas using fMRI as well as factors to consider during the preprocessing of fMRI data. These topics will lead to the presentation of the goals of our study and the specific hypotheses we make.

Functional connectivity

Functional connectivity is defined as the *correlations between spatially remote neurophysiological events* (Friston, 2011) and thus a study of functional connectivity can lend insight into functional brain networks. For example, if a positive correlation is found between the time series from two areas of the brain in data obtained when a task is being performed, we would conclude that these two areas are active simultaneously and belong

to a network which supports the processing of a certain group of tasks related to the task which was done. Likewise, for spontaneous activity, if two areas' time series are found to be correlated, this means that they show synchronized slow fMRI fluctuations, i.e. the blood oxygenation fMRI signal is synchronized. Furthermore, since the blood oxygenation fMRI signal is correlated with locally recorded neurophysiological activity (Shmuel and Leopold, Human Brain Mapping, 2008; Scholvinck et al., PNAS 2010), the spontaneous fluctuations of the neurophysiological activity of these areas are also correlated, indicating that the areas belong to a network that processes information jointly. The relationships can also reflect additional underlying information including anatomy (De Luca et al., 2006; Fox and Raichle, 2007; Koch et al., 2002; Quigley et al., 2003; van den Heuvel et al., 2009), or inherent co-fluctuation of areas (Fox et al., 2006; Broyd et al. 2009; Fornito et al., 2012). In order to extract the information which is interesting in your study from the functional imaging data, proper paradigms and methodologies must be used.

The resting state

The resting state of the brain is one particular type of activation which is very interesting and may contain information relevant to our understanding of a subject's response to a task, task performance, and

underlying anatomical structure (Fox and Raichle, 2007). The "resting-state" is a brain state in which the subject is awake and aware, however they attempt to keep their mind clear, not focusing on a task or idea; it is often done with eyes closed. Recently the neuroscience community has taken particular interest in resting state brain activations because of its potential to explain the brain on a fundamental level. The human brain consumes 20% of the body's energy even though it is only about 2% of a human's body mass, and most of this energy is utilized in continuing neuronal signaling. Task related processing shows increases in neuronal metabolism of less than 5% in most cases, a very small amount compared to the resting energy use (Ames, 2000; Attwell and Laughlin, 2001; Lennie, 2003; Shulman et al., 2004; Raichle and Mintun, 2006). This is particularly striking when one considers the likelihood that resting patterns of activity are maintained during task conditions, where the latter's activation patterns are simply superimposed, as seen in a recent study in the somatomotor system (Arfanakis et al., 2000). These ideas together propose that the major functioning of the brain is contained in resting activations. Koch et al. (2002), Quigley et al. (2003), and van den Heuvel et al. (2009) among others, have theorized that from functional connectivity in the resting state we can extract the anatomical connections between areas due to findings

linking functional and anatomical connectivity in diffusion tensor imaging studies and studies of patients with absent anatomical connections.

Most functional studies will consider brain activations following a task or stimulus which leads, in general, to stronger signal strengths in the task response areas relative to the rest of the brain unlike in unprovoked resting activations. There is evidence that supports the idea that upon presenting the brain with a task, the regions responsible for responding to the task increase their relative ensemble activity while those involved in the ongoing non-task activations, see a decrease in the collective quality of their activity, resulting in a greater contrast in network activations (Hampson et al., 2004; Lowe et al., 2000; Morgan and Price, 2004; Jiang et al., 2004; Bartels and Zeki, 2005; Sun et al., 2006). Considering that this is not the case in resting conditions, in order to reliably detect connections using resting state data, we consider the strongest connections we can expect to observe. Direct connections, mediated by an axon emerging in one area and terminating in another, are often associated with stronger functional connectivity than indirect connections since in indirect connections there is more energy lost in the propagation of the signal through additional brain areas (Reiss et al., 2011).

Network models

Many mathematical models aim to predict brain networks, however thus far their successes have been largely limited to networks on a large spatial scale. Adjustments in the approaches taken with these models will need to be made in order to optimize for use on smaller scale networks. Methods of determining the functional connectivity of a brain network are varied. A recent paper by Smith et al. (2011), summarizes a selection of the more popular and promising methodologies including correlation analyses, Granger causality, and mutual information. Each of the considered models has its shortcomings such as the ability to differentiate direct versus indirect connections, sensitivity, or practicality. The models also have their unique strengths and so the analysis type in which one would use each of the models should be chosen carefully. In the Smith et al. (2011) paper, as we will do here, an exploration of how a set of models can perform in the context of anatomical network modeling is in question. The most popular models for network analyses using fMRI data are correlation based analyses, coherence based analyses, lag based analyses, and Bayesian network analyses.

Correlation based models

The most popular of all the methods are the correlation models which are appealing because of their simplicity and intuitiveness. Correlation considers how the variation in one region co-varies with the variation in another. If the way the two regions vary is related, then they are said to be correlated. Correlation analyses of the brain's networks using fMRI data are plentiful and have looked at many things from the modular organization of functional networks in children with frontal lobe epilepsy (Vaessen et al., 2012), to language networks in subjects with anophthalmia (Watkins et al., 2012), to the reliability of the identification of the default mode network (Long et al., 2008). Partial correlation is also a widely used method which seeks to identify direct connections, rather than the direct and the indirect as identified with correlation (Smith et al., 2011), by regressing out information from all additional regions in the network in considering the relationship between two regions. Partial correlation has been used in studies looking at cortical networks mediating object motion (Calabro and Vaina, 2012), network abnormalities in patients with schizophrenia (He et al., 2012), and the effects of electro-acupuncture stimulation on the default mode network (Liu et al., 2009), to name only a few. Independent component analysis (ICA) is another very commonly

used correlation based approach to studying fMRI brain activations. In the reverse of usual correlation analyses, it seeks to identify which regions in a network are uncorrelated and by default, any items which are not uncorrelated, would be exhibiting functional connectivity. ICA analyses have been used in fMRI studies such as the study by Esposito et al. (2008) where the authors examined individual and subject level ICA of the default mode network and are often used to segregate seeds regions which may show unique behavior in full brain data sets (Long et al., 2008;Li et al., 2012; Arbabshirani and Calhoun, 2011).

Coherence based models

Coherence based analyses look at the spectral density of two regions, examining the relationship of their power at the contained frequency elements. These types of analyses are commonly done in order to see how the data relates in certain frequency bands. Some examples of coherence models include Wavelet Coherence (Grinsted et al., 2004) and a normalized cross-spectral density (Zhou et al. 2009). Some relevant past studies include Sun et al. (2004) where interregional functional connectivity was examined under different motor tasks and Salvador et al. (2005) where whole brain networks were estimated.

Lag based models

The lag-based methods, such as Granger causality, seek to identify the causal relationship between two regions given information from a set number of previous time points in the two time series. If the data from previous time points in the time series of one region improves the prediction of a second region's time series (improvement from the prediction using only the second region's own time series) then the first is said to granger-cause the second (Seth, 2007). This method was first used to study neural interactions in 1999 by Bernasconi and Konig. Today, there are many variations which exist including a sparse multivariate autoregressive model (Valdes-Sosa et al., 2005), nonlinear models (Freiwald et al., 1999), time-varying models for non-stationary data (Havlicek et al., 2010; Hesse et al., 2003), and models with the structure of non-parametric spectral factorization (Dhamala et al., 2008) (Stephan and Roebroeck, 2012). The use of these methods with functional MRI data has been criticized partly because of a mismatch in the time frames of the studied neuronal communications and the scanner image acquisition time (Smith et al., 2011; David et al., 2008). Most fMRI studies to date have been done with repetition time of 2 or more seconds, meaning a full 2 second will pass before the scanner can acquire the next image of a given

slice of the brain. For studies of small scale brain networks, such as in the visual cortex, the flow of information between areas is expected to be quite fast, on the scale of a few tens to hundreds of milliseconds (Schmolesky et al., 1998), thus in these networks with fast spreading information, it seems unlikely that a study solely based on the progression over time would be effective. However in studies of slower acting influence, such as verbal comprehension, a repetition time (TR) of around 2 seconds may be able to capture the influence it aims to (Yang and Shu, 2012). Some believe that with adjustments the lag based methods can still be used in network analyses with long TR relative to the neuronal spread of information, such as David et al. in their 2008 study (TR= 3s) where they deconvolve the hemodynamic response function (HRF).

Bayesian Network models

Bayesian network approaches are showing promise for their use with fMRI data, as seen in Smith et al. (2011) where they were among the top models in tests using simulated data and network structures, as well as in Sun et al. (2012) where they were able to produce meaningful, consistent and reproducible network predictions from fMRI data of the human brain during video watching. These approaches assume that each region in a network has a set of parent regions: some subset of regions in

the system which was activated before the one in question. The Parent subset is defined such that all earlier activated regions not in the subset are conditionally independent from the node in question (Scheines et al. 1996). This conditional independence test (often a g-square test) allows one to establish which regions in the network are connected and which are not. The Bayesian network models also attempt to extract causal information from the data by establishing a set of criteria for the edges between regions and testing for the likelihood of one type/ direction of connection versus other possible types. These criteria and the way in which one tests the likelihood of an edge's existence in the network can vary between models. Some examples of different models include the Peter and Clark (PC) model where acyclic networks are forced and no latent regions are accounted for (Meek, 1995), the Cyclic Causal Discovery (CCD) model where cyclic structure is allowed (Richardson and Spirtes, 2001), and the Greedy Equivalence Search (GES) model where a validity score is associated with the network given each added or removed edge and the optimal score is sought (Chickering, 2003; Ramsey et al., 2010). Some Bayesian network studies that have been conducted using fMRI data thus far include a dynamic Bayesian network modeling study for longitudinal brain morphometry (Chen et al., 2012), and a large-scale

connectivity study between several resting-state networks in the human brain (Li et al., 2011).

Statistical interpretations

While the correct statistical approach to extract information from data is of great importance, how the results is interpreted is equally so. Consider correlation: most often correlation analyses will result in nearly all brain areas considered being correlated to some extent (Schwarz and McGonigle, 2011; Tohka et al., 2012) so how does one decide which of these correlations are significant? Statistical tests exist to determine the statistical significance of interactions measures, however in order to achieve real statistical significance, while accounting for multiple comparisons, avoiding circularity, and satisfying assumptions, many of these tests end up being unusable or impractical. Hochberg and Benjamini (1990) and Turkheimer et al. (2001) found that traditional approaches to multiple comparisons corrections are impractical in many medical studies including neural network studies as they eliminate a large number of true connections. Often statistical tests require the independence of sources, a criterion which is not necessarily met in fMRI studies (Wang et al., 2003). Moving outside of strict statistical methods for significance, one often sees thresholding implemented as in Buckner et al.

(2009) and Tomasi and Volcow (2010). The problem with thresholding is that it is largely arbitrary and thus it is likely to be imprecise in terms of really sorting the true interactions from the false (Turkheimer et al., 2000). Oftentimes in thresholding a high number of false positives exist or with a stricter threshold, a large proportion of true positives are missed (Tomasi and Volcow, 2010). Viable options for this "post processing" of interaction measure from network analyses are currently being sought as in Smith et al. (2011) where they considered a thresholding based on the data and some background information/ assumptions of the truth and falseness of the measures. In this, working with simulated networks, they were able to establish a threshold based on the false positive distribution. Others who have proposed post processing approaches include the following: Cole et al. (2010) have created a weighted global brain connectivity measure which can quantify inter-subject consistency while avoiding arbitrary connection strength thresholding; Scheinost et al. (2012) present "the intrinsic connectivity distribution", a voxel-based contrast procedure for resting state fMRI that does not need a priori information to define an ROI and characterizes a connection without an arbitrary threshold.

The visual cortex and retinotopy

One area of the human brain which is relatively well understood due to its testability and its relatively straight forward organization is the visual cortex. The concept of retinotopy was established in 1900 by Tatsui Inouye during his experiments involving victims of war with gunshot wounds to the head (Inouye, 1900). Further studies have refined our knowledge of the human visual cortex such that we are now in a position to quite reliably define visual areas using fMRI data, as is well outlined in Larsson and Heeger (2006). Visual stimuli changing in the visual field according to polar angle and according to eccentricity allow for the delineation of visual areas by identifying where the gradient of the polar angle is reversed. With these delineations being defined orthogonally to eccentricity phase lines, early visual areas V1, V2 and V3, among others, can be reliably identified on a subject by subject basis.

Fluctuations of non-neuronal origin and fMRI data preprocessing

Non-neuronal physiological fluctuations in fMRI studies and scanner instabilities which may impact the signals are real concerns (Glover et al., 2000; Tanabe et al., 2002; Wise et al., 2004; Birn et al., 2006; Lund et al., 2006). Blood Oxygenated Level Dependent (BOLD) fMRI, as is most commonly used in recent years, in essence captures the change in

presence of deoxyhemoglobin brought on by activation in regions of the brain. Thus one can agree that the signal captured by the MRI would be influenced by things affecting blood flow and the oxygen content of the blood such as the heart beat and breathing. Also significant is the major effect motion has on the magnetic field within the scanner and thus the signal the sensors perceive. There are processing approaches to correct these unavoidable confounds such as motion correction and a field gradient which attempt to align scans and adjust for field changes (Ing and Shwarzbauer, 2012; O'Donnell and Edelstein, 1985), measuring the heartbeat and rate of breathing to extract those frequency components from the measured signal (Glover et al., 2000; Birn et al., 2006; Lund et al., 2006), and removing the signal which is common to all regions, called the global signal (Zarahn et al., 1997; Macey et al., 2004; Fox et al., 2005; Carbonell et al., 2011). While motion correction is now widely accepted as a necessary step in fMRI studies (Churchill et al., 2012), the consensus on the removal of the global signal has not yet been met. While it seems probable that such a signal would not contain information interesting in studies which contrast the differences in activations of brain areas (Macey et al., 2004; Fox et al., 2009; Carbonell et al., 2011; Smith et al., 2011), some believe that it may contain useful information in potentially contained

networks (Golland et al., 2007), or that its removal will alter the relationship structures across regions in the brain (Aguirre et al., 1998; Desjardins et al., 2001; Gavrilescu et al., 2002; Laurienti, 2004; Macey et al., 2004).

Goals of the study

Our goal in this work is to identify the most appropriate network models for use with real human resting state fMRI data, and improve the usage of all models by means of a data calibration. The methods we test here against human resting state data are those previously tested against simulated data in a paper by Smith et al. (2011).

We have chosen to focus on the human visual cortex in the resting state and use our knowledge of monkey anatomy in our ground truth network. Evidence showing the relationship between the human visual cortex and the monkey visual cortex was found in Engel et al. (1997), Tootell et al. (1997), and Hinds et al. (2008), to name a few. Strong resemblances have been found thus far and they are expected to be robustly maintained in the lower visual areas.

Our hypotheses include: 1- The ability of a mathematical model to predict network connections between brain regions may differ when implemented using simulated or real data. Real data has any number of

confounds including physiological and mechanical, whereas simulated data will have a finite number of confounds as established by the neural model. 2- Even the best mathematical network models are not able to accurately predict brain networks and improvements will need to be made in order to make them usable in predictive studies. The shortcomings to overcome will be a high number of false positive connection predictions, and/or a low number of true connections predictions. We hope to propose a network "calibration" to remedy some of the impact of these shortcomings. 3- The parameters of the analysis and steps in data preprocessing will impact the successfulness of the models at accurately predicting networks. We expect that a fine balance between the presence of noise in the signal and the inclusion of additional useful information will need to be met and hope to shed light on the important elements of that balance. 4- The human visual cortex behaves similarly to the macaque monkey cortex between visual areas V1, V2, and V3. 5- Resting state BOLD fMRI functional connectivity reflects the anatomical connectivity between brain areas, but is not identical to the anatomical connectivity due to 'network effects'.

5- Manuscript

Evaluation and calibration of functional network modeling methods based on known anatomical connections

Debra Ann Dawson[1,2,5], Kuwook Cha[1,2,5], Lindsay B. Lewis[3,5], Janine D. Mendola[3,5], Amir Shmuel[1,2,4,5]

[1]Montreal Neurological Institute; [2]Dept. of Neurol. and Neurosurg.; [3]McGill Vision Res.; [4]Dept. of Biomedical Eng.; [5]McGill Univ., Montreal, QC, Canada

Keywords: Resting-state, Resting-state networks, Spontaneous activity, Functional connectivity, Coherence.

Submitted to: Neuroimage, 2012

Introduction

One of the most widely used and insightful methods for measuring activity in the brain is functional Magnetic Resonance Imaging (fMRI; Kwong et al., 1992; Ogawa et al., 1992; Bandettini et al., 1992). Efforts are constantly being made to improve processing techniques, data interpretation, and modeling of such data. A widely used analysis method is that of computing functional connectivity measures associated with resting-state data, obtained in the absence of a task (Biswal et al., 1995). The term "functional connectivity" refers to the correlations between spatially remote neurophysiological events. While this is the working definition, measures other than correlation per se can quantify a relationship in the temporal domain between neurophysiological events. Here we evaluate the performance of network modeling methods applied to resting-state fMRI time-series obtained from the human brain.

Past considerations to modeling of networks using fMRI data have found that careful model design and selection are required for the results of the study to achieve a semblance of validity. The more robust the model is, the more accuracy we can expect from the model. Specifically, successful models include as many network confounds as possible, such as forward, backward, or cyclic connections (Friston, 2011) as well as any

number of foreseeable contributions from non-neuronal, physiological sources (Fox and Raichle, 2007). Flexibility is also an important quality in a model; it should be able to represent the parameters and conditions of a wide array of experiments (Roebroeck et al., 2009).

A major shortcoming of many network models is the inability to distinguish indirect from direct connections based on fMRI data (Fox and Raichle, 2007). Such a task is even more difficult when analyzing strong activations such as in fMRI experiments with stimulus presentation. Studying functional connectivity with measurements of task evoked responses have been useful in identifying larger brain regions which co-activate (Bullmore et al., 1996; Calhoun et al., 2001; Mizuhara et al., 2004), however due to the strength of the responses and the fact that they are time-locked to the onset of a stimulus or a task, the details achievable in these connectivity analyses are limited. It has been suggested that the much lower amplitude signal of resting state activity is a fraction of the task response amplitude (Jiang et al., 2004; Ng et al., 2011) and thus that connectivity analyses of the brain at rest may give insight into functional connectivity at a finer scale; in fact it is expected that functional connectivity based on resting-state data is reflective of anatomical connections (De Luca et al., 2006; Fox and Raichle, 2007; Koch et al.,

2002; Quigley et al., 2003; van den Heuvel et al., 2009), although this reflection is not trivially identical (Honey et al. 2010).

In a recent paper by Smith et al. (2011), the abilities of various network models to correctly predict networks of simulated fMRI data were evaluated. The simulations considered a range of scenarios, including variations in number of nodes, session durations, connection strengths, input strengths, and other potential confounds thought to be present in real fMRI data. The main result was that, with simulated data, Partial Correlation, Inverse Covariance, and Bayesian network models could often predict over 90% of correct network edges with greater connection strengths than the 95th percentile of the false positive distribution. This being said, even these strongest methods showed weakness when faced with some of the tested confounds, indicating a necessity to test the model performances with real data, where any number of confounds may be significant.

Here we explore the validity of various network modeling methods by evaluating their successes at predicting expected anatomical connections between different regions in the human visual cortex. We chose retinotopic visual areas V1, V2, and V3 for pursuing this evaluation because of the in depth understanding we have of the anatomical

connections within and between these areas based on studies of the macaque monkey brain. The methods we test here against human resting state data are those previously tested against simulated data in a paper by Smith et al. (2011). Our goal is to identify the most appropriate models for use with real human resting state fMRI data. To quantify a method's ability to accurately predict a set of connections, we use measures of "c-sensitivity", a term coined by Smith et al. representing the fractional rate of detecting true connections. We found that Partial Correlation, Regularized Inverse Covariance and Bayesian Network methods achieved the highest c-sensitivities across tests. We propose to improve data-acquisition parameters and analysis by integrating a data-driven calibration into functional connectivity studies.

Methods

Data acquisition and preprocessing

Subjects and scanning sessions

Seven normally sighted subjects participated in this study after giving written consent in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki). Each subject was scanned in two sessions: one to obtain high-resolution anatomical images and fMRI

for retinotopy and the other to obtain high-resolution anatomical images and fMRI in the resting-state. During the retinotopic mapping scans, the subjects were instructed to fixate their eyes on a fixation spot at the center of the screen through a projection mirror. During the resting-state functional scans, the subjects kept their eyes closed.

MRI data acquisition

Data were acquired on a 3T Magnetom TIM Trio scanner (Siemens, Erlangen, Germany). A phased array head coil was used, with 20 and 32 channels employed for retinotopy mapping and resting state scans, respectively. Echo-planar imaging was used to measure blood oxygenation level-dependent (BOLD) changes in image intensity. Resting-state functional MRI was pursued in 8 runs per each subject. Each run consisted of 256 contiguous EPI whole-brain functional volumes [repetition time (TR) = 2000 ms; echo time (TE) = 30 ms; flip angle = 90°; 30 slices; matrix = 112 ×112; field of view (FOV) = 224 mm; acquisition voxel size = 2 × 2× 2.2 mm]. Retinotopic mapping was pursued with axial-oblique slices [repetition time (TR) = 2000 ms; echo time (TE) = 30 ms; flip angle = 90°; 28 slices; matrix= 128; field of view (FOV) = 263 mm; acquisition voxel size = 2.05 × 2.05 × 3 mm]. In both sessions, T1-weighted (MPRAGE)

anatomical volumes were acquired with high resolution (1 × 1 × 1 mm) for between session co-registration of the fMRI data.

Retinotopic mapping and definition of visual areas

The visual cortex's functional organization corresponds spatially to the inverse of the visual field representation, a phenomenon called retinotopy (Inouye, 1900; Holmes, 1945). Using retinotopic analysis we can identify a hierarchy of regions which process different sub-fields of the visual field and at different levels of detail. Here we focus on visual areas V1, V2, and V3, further defining the ROIs according to their eccentricities in the visual field. Visual areas V1, V2, and V3 of each subject were defined according to conventional phase encoding retinotopy (Engel et al., 1994; Larsson and Heeger, 2006; Sereno et al., 1995). The visual stimuli were projected from a liquid crystal display (LCD) projector at 1024 × 768 resolution and 60 Hz refresh rate onto a translucent screen at the end of the scanner bore. The subjects viewed the screen at the total viewing distance of 138 cm through a mirror mounted to the coil, which yielded 32×24° (40° along the diagonal) of viewing angle. During the retinotopic mapping scans, the conventional retinotopy mapping stimulus of expanding/contracting rings or clockwise/counter-clockwise rotating wedges periodically traversed the visual field every 64 seconds. The width

and expansion/contraction speed of the ring stimulus were logarithmically changed as a function of eccentricity in consideration of the cortical magnification factor whereas the arc angle of wedge stimulus was constant at 10°. The ring and the wedge stimuli consisted of checkerboard patterns where brightness and colors changed at a rate of 8 Hz to maximize the neural responses of the visual areas of interest. In order to aid the subjects in maintaining their fixation, it was requested that they report the direction of a 0.5° arrowhead which randomly changed directions (up, down, left, right) using a fibre-optic button box.

In order to define visual areas, the temporal phases of the travelling waves of fMRI responses evoked by the ring and wedge stimuli were matched to the visual field locations in the polar coordinate system: the eccentricity and polar angle positions were taken from the ring and wedge stimuli respectively. To match the phase values of fMRI responses to visual field location, we shifted the time-series in accordance with hemodynamic delay (6 s, 3 TRs), reversed the time-series of contracting rings scans, averaged the time-series of the two types of ring scans, and took cosine coefficients of the fast Fourier transform of the resultant time-series (Larsson & Heeger, 2006). The phase-encoded retinotopic maps were visualized on the flattened surface of each subject's occipital cortex

using the FreeSurfer package (http://surfer.nmr.mgh.harvard.edu/; Sereno et al., 1995; Dale et al., 1999); Fischl et al., 1999)(Fig. 1A). Retinotopic visual areas were identified according to the criteria documented in Larsson and Heeger (2006).

Preprocessing of resting-state fMRI and construction of resting-state time-series

The 3 first volumes were discarded in the preprocessing of each run, in order to avoid non steady-state effects. The fMRI data were preprocessed using the NeuroImaging Analysis Kit (NIAK; http://code.google.com/p/niak/; Bellec et al., 2010) including slice acquisition time correction, motion correction, removal of ultra-slow time drifts (high-pass filter with a 0.01 Hz cut-off) and high temporal frequencies not contributing to resting state functional connectivity (low-pass filtering with a 0.1 Hz cut-off), and removal of modeled physiological noise (CORSICA; Perlbarg et al., 2007). Spatial smoothing and co-registration between subjects were not applied to avoid unnecessary signal spread over space.

To obtain resting-state time-series, we regressed out confounds from the time series of each voxel. These include the first principal component (PC), motion correction parameters, and the linear trend

(http://www.math.mcgill.ca/keith/fmristat/; Worsley et al., 2002). The purpose of removing the first PC was to eliminate the spatially-global resting state fluctuations that were common to all regions in the imaged volume, without introducing artifactual negative correlations (Carbonell et al., 2011). After removing all of these covariates, the residuals were normalized by subtracting their mean over time and by dividing them by their standard deviations.

Regions of interest and resting-state time-courses

The resting-state fMRI time-series were integrated over the voxels which were classified into three eccentricity bins for each visual area we defined through retinotopic mapping. We projected the visual area ROIs originally defined on the cortical surface into the native functional image space where the resting-state fMRI data were acquired. For this projection, we aligned the high-resolution anatomical volume images from the retinotopy and resting state sessions (used for cortical surface extraction) using an automated robust image registration algorithm (Nestares & Heeger, 2000) to obtain the transformation matrices.

The phase values obtained from retinotopic mapping analysis were assigned to the individual voxels in the ROIs. Voxels from each area were classified into three eccentricity bins, i.e., central, intermediate, and

peripheral, with the boundaries 0.40° - 2.50°, 2.50° - 5.97°, and 5.97° - 12.00° respectively. The most central boundary was arbitrarily set to 0.4° in order to avoid the region prone to reflect signals of response to the fixation point, yet still use as much information as possible. The steps were defined such that each region would span approximately the same length in the cortex based on the cortical magnification model from Duncan and Boynton (2003). Given 3 visual areas and 3 eccentricity regions defined in each hemisphere, and given that we looked at the ventral and dorsal regions of these areas separately within a hemisphere, we had 36 ROIs to work with (3 areas x 3 eccentricity bins x 4 quadrants). The time courses of the spontaneous activity within the voxels in each bin were averaged. To prevent trivial enhancement of interactions due to the dependence of signal strength on the number of voxels in the bins, the voxel numbers were equalized across bins for each subject. To achieve this, we found the maximum number of voxels which could be selected evenly from all bins across areas in a given subject. In the bins where more voxels were available than the identified maximum, the voxels which had corresponding phase values closest to the midpoint eccentricity between the boundaries were selected for inclusion. Thus, the number of voxels across areas and

eccentricity bins were equal within a subject. The number of voxels used per bin for each of the 7 subjects was in the range of between 8 and 30.

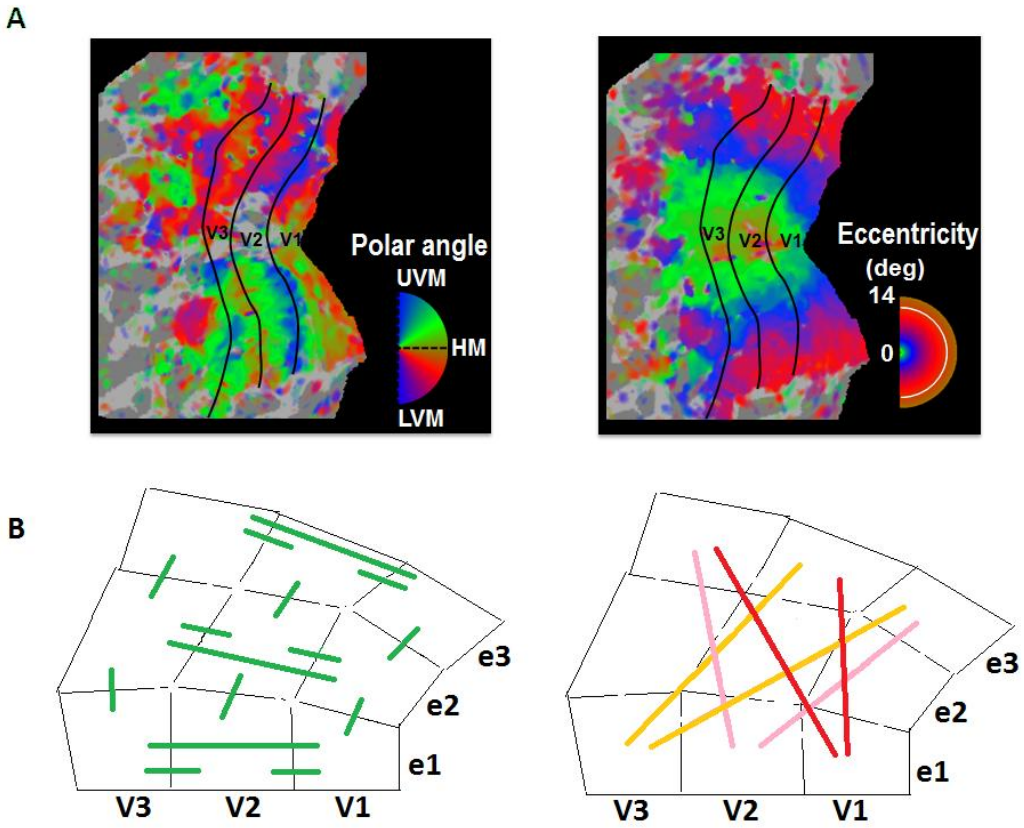

Figure 1. Retinotopy and connections tested for use in c-sensitivity calculations. A) Visual areas V1, V2 and V3 were defined by examining the reversal of direction of increasing polar angle phases perpendicular to the iso-eccentricity delineation. Voxels from each area were further classified into three eccentricity bins: central, intermediate, and peripheral, with the boundaries 0.40° - 2.50°, 2.50° - 5.97°, and 5.97° - 12.00°, respectively. On the left we see the functional response to a stimulus changing in terms of polar angle and on the right we see the response to a stimulus changing in terms of eccentricity. B) Connections tested between ROIs for c-sensitivity calculations. Each hemisphere was divided for analysis according to its dorsal and ventral parts, based on the mapping of the

horizontal meridian, resulting in brain regions representing the four quadrants of visual space. Within a quadrant, connections between adjacent ROIs within a visual area as well as those of the same eccentricity between visual areas were expected to be connected (left; represented in green lines). Non-adjacent eccentricity regions in different areas were expected to not be directly connected (right; represented in red, pink, and yellow lines). UVM = upper vertical meridian, HM = horizontal meridian, LVM = lower vertical meridian.

Models

The most promising variations of all models studied in Smith et al. (2011) are studied here. They are briefly presented in Table 1. We note that while many of these models attempt to predict the directionality of connections, here we prefer to focus on predicting the existence of the connections only. Thus we do not test the directionality aspect of their outputs.

Retinotopic visual areas as a network model

We based our analysis on the 'ground-truth', thoroughly studied, anatomical connectivity in the macaque monkey visual cortex. We assume that the principles governing visuotopic anatomical connections within and between retinotopic visual areas V1, V2, and V3 are similar in humans and macaque monkeys. Indeed, numerous human fMRI studies (e.g. Sereno et al., 1995; Engel et al., 1997; Dumoulin and Wandell, 2008) have shown

that functional retinotopy of these areas in the human brain closely follows the well studied principles demonstrated in the monkey brain. More recently, it has been shown that resting-state functional connectivity between these areas is topographically specific too (Heinzle et al., 2011). These findings strongly indicate that the anatomical connections between V1, V2, and V3 in the human brain are retinotopically specific, just like their homologous connections in the monkey brain. We therefore test the network predictions of each method against network expectations from earlier studies on the monkey cortex. ROIs within each of the considered areas that represent adjacent regions in the visual field have direct anatomical connections (Gilbert and Wiesel, 1979; Rockland and Lund, 1982; Amir et al., 1993; Lyon et al., 1998), while non-adjacent ROIs do not. Visuotopic feedforward connections (Salin and Bullier, 1995) and visuotopic feedback connections with broader extent relative to feedforward connections (Salin and Bullier, 1995; ; Angelucci et al., 2002; Shmuel et al., 2005) exist between V1, V2, and V3. Although areas V1 and V3 are not adjacent, tight links have been shown between their retinotopically matched regions (Girard et al., 1991).

Table 1: Connectivity models tested against human resting state FMRI data

Model name	Description	Variations	Implementation
Full Correlation	Pearson's linear correlation coefficients are calculated.	Bandpass filtering of the times series to retain only the bottom half of the frequency range (fullcorr1/2)	Matlab's "corr" function
Partial correlation	We find Pearson's linear correlation between the residuals of two nodes after linearly regressing out all other nodes' information from them.		Calculated using the inverse covariance matrix
Regularized inverse covariance (ICOV)	The model has parameter lambda, which regularizes the covariance matrix to below a certain threshold. The regularization used here is the Lasso method (Banerjee et al., 2006; Friedman et al., 2008) which results in entries that are far from zero maintaining their magnitudes, and those that are close to zero being reduced more.	Lambda=5, 100	L1precision implementation of ICOV available at www.cs.ubc.ca/~schmidtm/Software/L1precision.html
Mutual information (MI)	Mutual Information (Shannon, 1948) is a measure of how much information we can extract about a node given the time series data of another node (Quiroga et al., 2002). It is calculated using the entropies of each node and the joint entropy of both nodes given by the joint probability of the nodes within a certain radius around each timeseries datum.	Partial mutual information (regress the information out of the two nodes in question that comes from the other nodes in the network)	Functional Connectivity Toolbox (FCT) (Zhou et al., 2009) (groups.google.com/group/fc-toolbox)
Coherence	The cross wavelet power is computed for the time series from two nodes and then a coherence value is calculated in a manner similar to correlation coefficients using the cross wavelet powers.	Wavelet Coherence, or Coherence A (cohA); the normalized cross-spectral density, or Coherence B (cohB)	Crosswavelet and Wavelet Coherence toolbox (Grinsted et al., 2004) www.pol.ac.uk/home/research/waveletcoherence); FCT (see above)
Generalized Synchronization (GenSynch)	These are nonlinear interdependency measures. They all consider the mean over time points of the mean squared Euclidean distance to a given number (model order) of nearest neighbor signals. To this they compare the node-conditioned mean squared Euclidean distance which considers the signal at time points where a second node's nearest neighbor signals occur. Lags=1,2	S calculates the interdependency of node B on node A by defining the average over all time points of node A's mean squared distance (MSD)divided by the node B-conditioned MSD; H calculates the average of the log of this same ratio; N computes a normalized average by taking the difference between node A's MSD and the node B-conditioned distance, all divided by A's MSD (Quiroga et al,2002).	www.vis.caltech.edu/~rodri/programs/synchro.m

Table 1 (continued): Connectivity models tested against human resting state FMRI data

Model name	Description	Variations	Implementation
Granger causality and related lag-based methods	At the basis of the lag based models used here is the multivariate linear autoregressive model which aims to approximate an ROI time series by assigning weights, or interaction effects, to the previous time points up to a maximal lag/model order (specifically, the number of recent time points to include in the autoregressive model). The idea is that if two nodes are interacting, the past information in the time series from one node will contain information useful in the prediction of the other's. Model order=1 for all except GrangerA.	GrangerA is the standard described to the left. Model order= 1, 3, 20; of the following. GrangerB uses the Bayesian information criterion to determine the ideal model order, up to a set maximum; Direct Causality (DC) is the sum of the squares of the interaction magnitudes using the MVAR model; Geweke Granger Causality (GGC) determines by use of an F-statistic if the interaction magnitude between a pair of nodes is significantly different from zero; Partial Directed Coherence (PDC) is calculated in the frequency domain and seeks to remove instantaneous interaction effects; Directed Transfer Function (DTF) uses the MVAR model in the frequency domain.	"Causal Connectivity Analysis" toolbox created by Seth (2010) anilseth.com; (www.mathworks.co.kr/matlabcentral/fileexchange/ 25467-grangercausality-test); BioSig toolbox (biosig.sourceforge.net) for DC, GGC, PDC, DTF
Patel's conditional dependence measures	This measure is derived from a bivariate Bernouilli Bayesian model which aims to identify the joint activation of a pair of brain voxels using a multinomial likelihood with a Dirichlet prior distribution.	Kapa is calculated using the probabilities of activation of two nodes as well as the probabilities of each node given the other. Tau is calculated using the same probabilities but with the goal of establishing ascendency (directionality); We also binarize the data at a 0.75 threshold.	(Patel et al. 2006)
BayesNet methods	The Baysian Network methods assume that each node in a network has a set of parent nodes: some subset of nodes in the system which was activated before the node in question. The Parent subset is defined such that all earlier activated nodes not in the subset are conditionally independent from the node in question. The product of the activation probabilities of the node given information from the parent node subset constitutes the interdependency between the node and the parent nodes (Pearl, 1985).	The PC (Meek, 1995) method "is a pattern search which assumes that the underlying causal structure of the input data is acyclic, and that no two variables are caused by the same latent (unmeasured) variable" (Spirtes et al., 1993); The CPC (Ramsey et al.,2006) method differs from PC in its detection of direction of interactions; GES (Chickering, 2003; Ramsey et al., 2010) scores the validity of the network with each edge added or removed and seeks to maximize this score; FCI (Zhang, 2008) relaxes the condition that unmeasured variables can't contribute to the network; CCD(Richardson & Spirtes, 2001) allows cyclic causal structure.	Tetrad IV toolbox (www.phil.cmu.edu/projects/tetrad/tetrad4.html)
Linear, Non-Gaussian, Acyclic causal Models (LiNGAM)	An iterative algorithm which uses independent component analysis including whitening (applied to the full set of nodes) to extract the independent components from the timeseries data. The remaining interaction information is normalized		www.cs.helsinki.fi/group/neuroinf/lingam (Shimizu et al, 2006)

The 36 brain regions defined above served as “nodes” in the various models we studied. The regions of interest studied lie in visual areas V1, V2, and V3, each of these being divided into three regions according to eccentricity: central eccentricity (E1), intermediate eccentricity (E2), and peripheral eccentricity (E3). Each hemisphere was divided for analysis according to dorsal and ventral delineations resulting in brain regions representing the four quadrants of the visual space. Connections between adjacent ROIs within a visual area were expected to be directly connected. Similarly, ROIs of the same eccentricity between visual areas were expected to be directly connected. ROIs representing non-adjacent regions of visual space in two different areas were considered as non-directly connected. These considerations resulted in 15 expected connections and 6 pairs of ROIs expected to not be connected per quadrant (Fig. 1B) (or 60 expected and 24 unexpected connections per scan).

Statistics

Fig. 2 describes the statistical method we applied schematically. The 9 node (quadrant) network’s mean over voxels within each of the 9 ROIs’ time-series were input into the models which would each output a 9X9 matrix of coefficients of interaction. A null set of time-series was also

input into each model to obtain a null set of interaction coefficients usable in normalizing the raw connection strengths (utilizing the mean and standard deviation of the distribution), allowing better comparison of interaction strengths across models. The null data sets were created by testing for connections between time series from different subjects' datasets, which have no causal connections between them (i.e., the subject labels for each node in the network were randomly shuffled; Smith et al., 2011). Recall that we obtained data from 56 runs total (8 runs for each of 7 subjects) and attributed a mean (across voxels) time-series to each of 36 nodes. The representation from the 7 subjects was distributed as evenly as possible; each subject contributed a random node's time-series from each of 5 randomly selected runs, the time-series then being randomly assigned to null nodes. In order to create a 36 node null, a 6th run's data from a randomly-selected subject was used. Eight 36 node nulls were created. Upon creating a 36 node null distribution, the data set was broken down into quadrants, 4 data sets of 9 nodes, as done with the real data. We note that the nulls were only used to visualize the relative interaction magnitudes across methods (top panels in Figs. 3 and 6); quantitative evaluations were exclusively based on non-shuffled data, as described below.

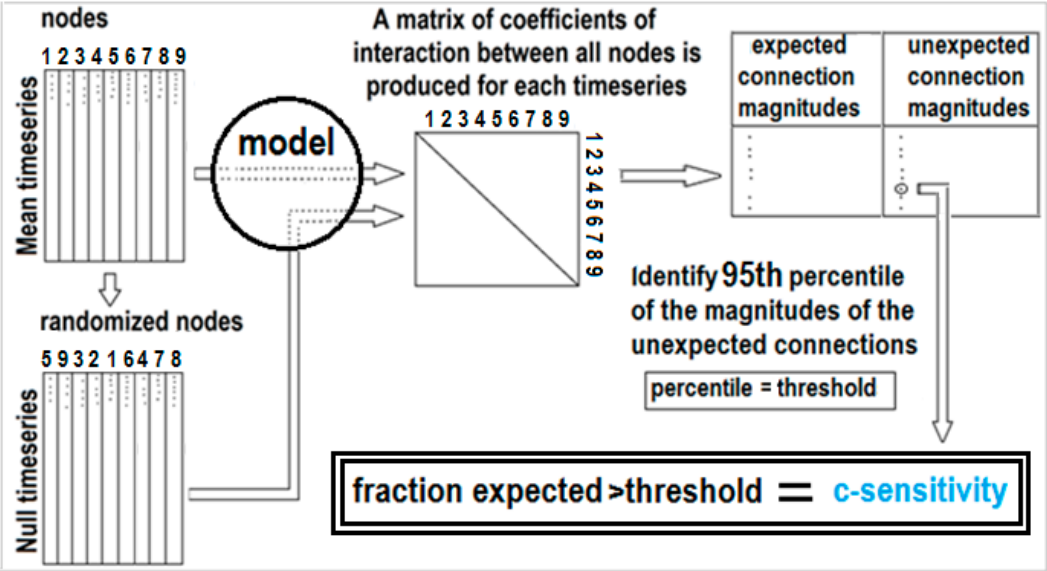

Figure 2 . Method of calibrating the data according to the distribution of unexpected connections. This method is modeled after that used by Smith et al. (2011).

For quantitative evaluations, for each network modeling method we identified the interaction magnitudes of the expected and unexpected connections outlined in Fig. 1B. Note that the BayesNet models as well as GrangerB provide binary outputs, 1 for a predicted connection, 0 for no connection. In order to attribute an interaction magnitude to a pair of nodes using these methods, the parameter controlling the method's sensitivity was varied. 100 logarithmically spaced sensitivity settings' networks were computed, and the (-log of the) least stringent setting which predicted a connection was attributed to that connection as its strength (Smith et al., 2011).

For each network modeling method and for each of the 8 runs per subject, the 95th percentile of the interaction measures of the 24 unexpected connections (6 per quadrant) was identified and used as a threshold above which interaction magnitude connections are likely to be direct. To obtain the 'c-sensitivity', the expected connection interaction magnitudes above this threshold were counted and their number was then divided by the total number of expected connections per run (60).

In order to consider how session length would impact the performance of different network models, single runs of 8.5 minutes obtained from a subject were concatenated to give longer session durations of 17 minutes (2 runs), 34 minutes (4 runs), and 68 minutes (all 8 runs.) Since all of these runs were obtained in the same fMRI session with only short breaks in between, any inconsistencies occurring between runs are expected to be negligible. To avoid discontinuities due to small scanner-related gain differences between scans, we normalized the time-courses of individual runs before concatenating them. To this end, individual runs were normalized to have mean of 0 and standard deviation of 1 only for the purposes of the concatenation, as this is not a requirement of the models. Given 253 time-points (TR = 2 s) in each 8.5 minute session, discontinuities due to concatenation are not expected to introduce

notable biases. Indeed, based on concatenation of a significantly larger number (24) of significantly shorter duration (26.67 s, TR = 2.5 s) of resting-state blocks, Fair et al. (2007) concluded that the functional connectivity for concatenated resting periods is qualitatively and quantitatively very similar to that of continuous resting-state data. In addition, concatenation is commonly used in current resting-state functional connectivity studies (e.g., Smith et al., 2012).

Given the success of the longer runs in the Smith et al. (2011) simulation study, we use the 68 minute sessions in our main analyses.

Results

Overall c-sensitivity results

Fig. 3 (top panel) presents the normalized interaction measures of the expected and unexpected interaction coefficient magnitudes predicted by each model. We hoped to see as little overlap in the distributions as possible, this being especially important when a ground truth set of connections isn't known beforehand. The number of data points plotted for a model is the number of expected (60; in blue) or unexpected (24; in red) connections times the number of subjects (7). To evaluate the performance of each method, we first computed the separation between

the interactions measures obtained for the expected and unexpected connections. To this end, for each method we found the difference between the means of the distributions of the normalized interaction measures for the expected connections and the unexpected connections relative to the overall range of both distributions. Pcorr, fullcorr, fullcorr1/2, ICOV 5 and 100, GenSynchH1 and H2, N1 and N2, GenSynchS2, Patel-Kapa, Patel-Kapa bin0.75, as well as MI and cohB all have separations of 20-28%. This is compared to poorly separated methods which have most often less than 5% separation between the respective means. Notably, the BayesNet methods have by far the best separations between their expected and unexpected interaction magnitudes, the means being offset by 38% to 64%.

We calculated the fraction of expected connections with connection strength greater than the 95th percentile of the strengths of the unexpected connections for each subject (68 minute session length). The mean fraction across subjects is considered as the c-sensitivity score of a model; this is plotted in Fig. 3 (bottom panel). As one would expect, the most c-sensitive models show greater separation between the expected and unexpected connection magnitudes. The top model was pcorr at 83% c-sensitivity, followed by ICOV5 (79%), BayesNet methods PC, CCD, CPC,

and FCI (77%), fullcorr (75%), fullcorr1/2 (75%), ICOV 100 (71%), Patel-Kapa (71%), Patel-Kapa bin0.75 (70%), GES (67%), and the GenSynch measures (67-75%). The methods with the lowest c-sensitivity scores, all scoring below 20%, were the lag-based methods including Granger A and B, pMI, Patel's Tau, and LiNGAM. MI, and cohA and B showed intermediate performance, scoring between 52% and 57% c-sensitivity.

It is important to note that while a method may have good separation of normalized interaction measure means, its c-sensitivity may not necessarily be high since we set our threshold by the 95th percentile of the interaction measures of the unexpected connections, not the mean. The consequence of this is that even if the majority of the interaction measures of the unexpected connections are low, if 5% of them are relatively high, the threshold may still be high enough to eliminate a fair portion of the expected connections. This is why the BayesNet methods have better separations than any other methods (Fig. 3 top panel), but do not achieve outstanding c-sensitivities (Fig. 3 bottom panel). We still discuss the separation of the means, however, because we find it a good representation of the trends in the interaction measures of the expected and unexpected connections.

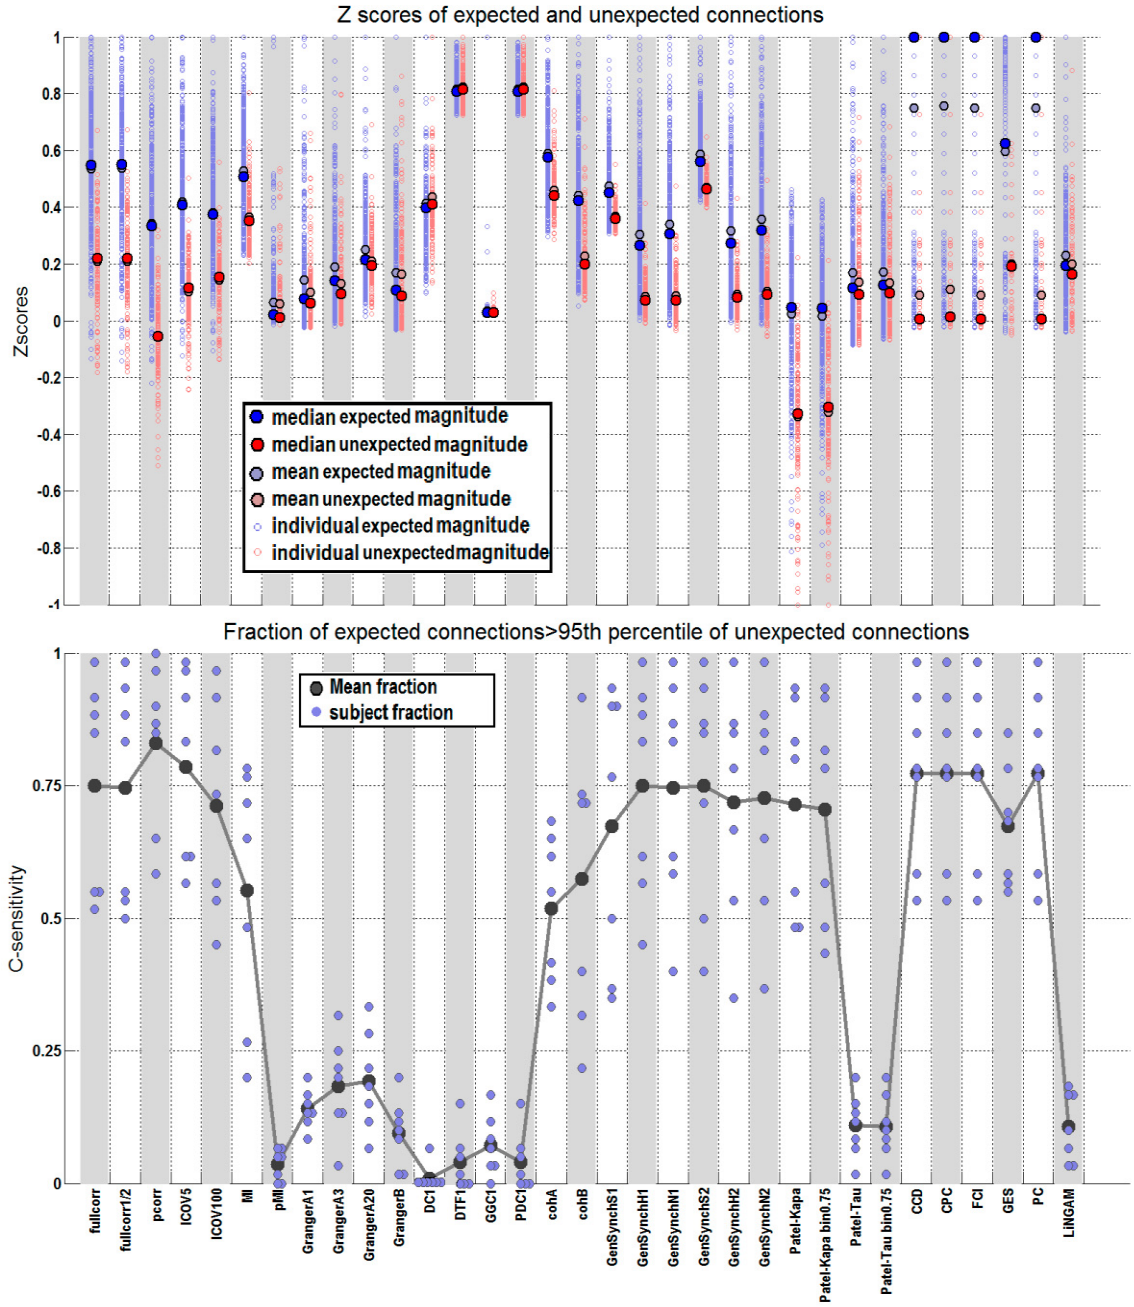

Figure 3. Evaluation of the network modeling methods using resting-state data. In the top panel, Z scores of the expected (blue) and unexpected (red) connection interaction magnitudes predicted by each model are plotted. The vertical scatter plots show the Z score of each connection for each run and subject. In the lower panel, the fraction of expected connections that are estimated with higher connection strengths than the 95th percentile of the unexpected connection distribution is plotted. The

vertical distributions shown here present the 7 subject fractions, where overlapping data points are horizontally adjacent. The 8 runs' time-series were concatenated prior to input into the models for optimal model performances. The mean values from the lower plot represent the c-sensitivity values of the models.

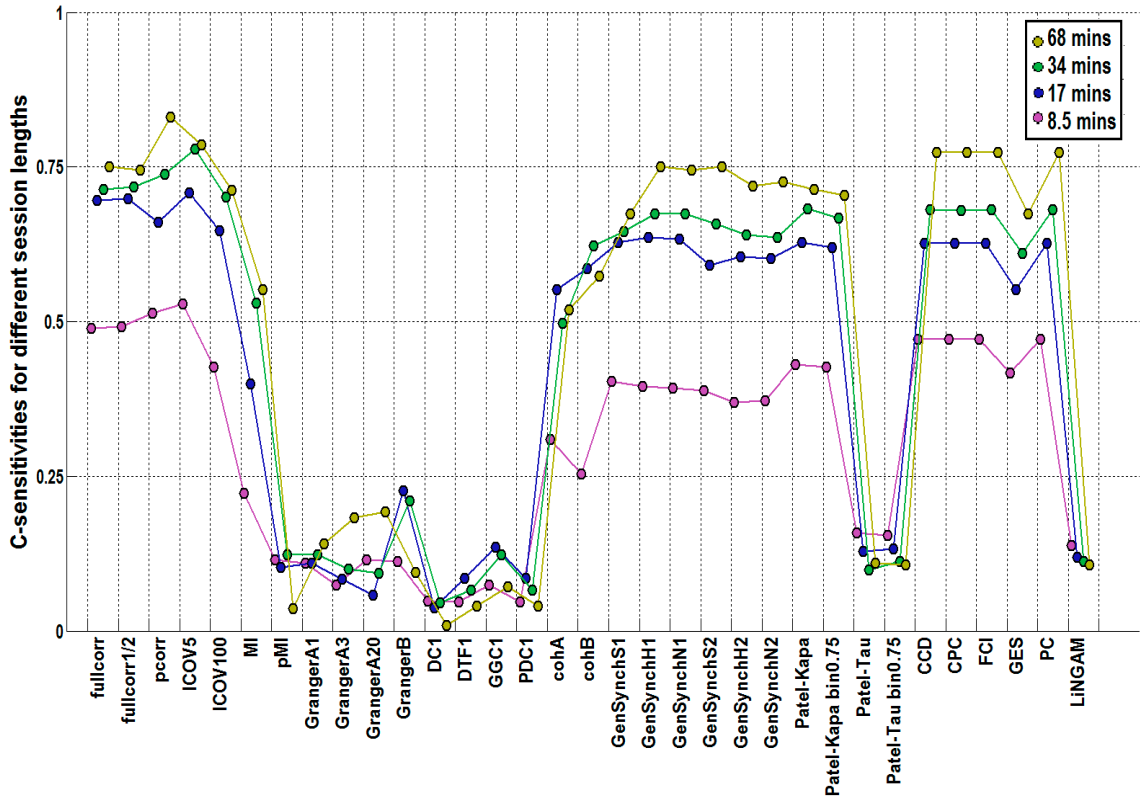

Figure 4. C-sensitivity values plotted for different MRI session lengths. All connections tested lie within a quadrant, as shown in figure 1B and the network considered is that of the quadrant, with 9 nodes. The session durations tested are multiples of the run length, which is approximately 8.5 minutes. We tested 1, 2, 4, and 8 runs, or 8.5, 17, 34, and 68 minutes by concatenating timeseries of the 8.5 minute runs.

Session duration

Recall that we obtained from each subject a total of 8 runs of 8.5 minutes each, resulting in resting-state fMRI sessions of over 1 hour. In most resting-state studies, less data are collected per session. To test whether the decreased data would affect the results significantly, we conducted the same analysis as above with 9 input nodes and all 8 runs concatenated, this time concatenating only 1, 2, or 4 runs of data.

As in the 8 run case, we concatenated the runs in the order they were obtained to assure as little variation between runs as possible, such as can occur because of a change of state of the subject over time. The eight 8.5 minute null runs were concatenated in the same way as the subject runs in these tests.

Fig. 4 shows the c-sensitivity results from each of these session durations. The 8.5 minute session length yielded low c-sensitivity for all methods. In this case, only pcorr and ICOV5 obtained c-sensitivities slightly higher than 50%, in the range of 51- 53%. As one would expect, with longer session duration, the results obtained from most of the networks were more reliable. With the doubling of the session length from 8.5 to 17 minutes, a sizeable increase in c-sensitivities occurs in the top models, now most achieving in the range of 60-70% c-sensitivity. Further

doubling of session length does result in increased c-sensitivities with the top models, however the increase is not as large as seen in the initial doubling. The same models which perform best in our main test using 68 minute sessions, also perform well with other session durations. The models which performed poorly in our main test (lag-based, pMI, Patel's tau, LiNGAM), also performed poorly with session duration shorter than 68 minutes. They often do not show a clear trend of improved results with increasing session duration.

Number of nodes

As many of the models we looked at were multivariate, the results could depend largely on the number of nodes input to the model. Upon dividing the visual cortex into quadrants, we saw a logical definition of networks based on 4 quadrants or 2 hemispheres or the entire visual cortex. Recall that we also tested the effect of session duration by concatenating scans. After concatenation, the number of extended scans will be referred to as the 'effective number of scans' in the text that follows. These network definitions equated to 9 nodes (one quadrant; resulting in 4 quadrants x 7 subjects [9 x effective number of scans] time-series inputs for the models), 18 nodes (one hemisphere; resulting in 2 hemispheres x 7 subjects [18 x effective number of scans] time-series inputs for the

models), and 36 nodes (2 hemispheres = the entire visual cortex; resulting in 7 subject [36 x effective number of scans] time-series inputs for the models) respectively. Fig. 5 shows the c-sensitivity plots for each of these networks in the case of 17 minute sessions. In general, for many of the models c-sensitivity remained almost unchanged with differing number of nodes, but this was expected as many of the models are bivariate. Partial Correlation, the BayesNet methods, and to a lesser extent ICOV showed decreased performance with increased number of nodes. Partial Correlation seemed to suffer most with increased number of nodes, which could be expected due to the increased number of time-series that it had to regress out. We were curious if increased data would compensate for the noise in the larger networks so we also pursued this analysis with 68 minute sessions. When the session duration is increased, the models which had seen variation in their c-sensitivities in the shorter (17 min) session duration test still showed decreases in c-sensitivities, however less dramatic, and sometimes no decrease at all. Given this, and that the connections we tested lay within a quadrant and not spanning quadrants, the 9 node network was sufficient and most practical for further evaluation in other parts of our study since it also required the least computational time.

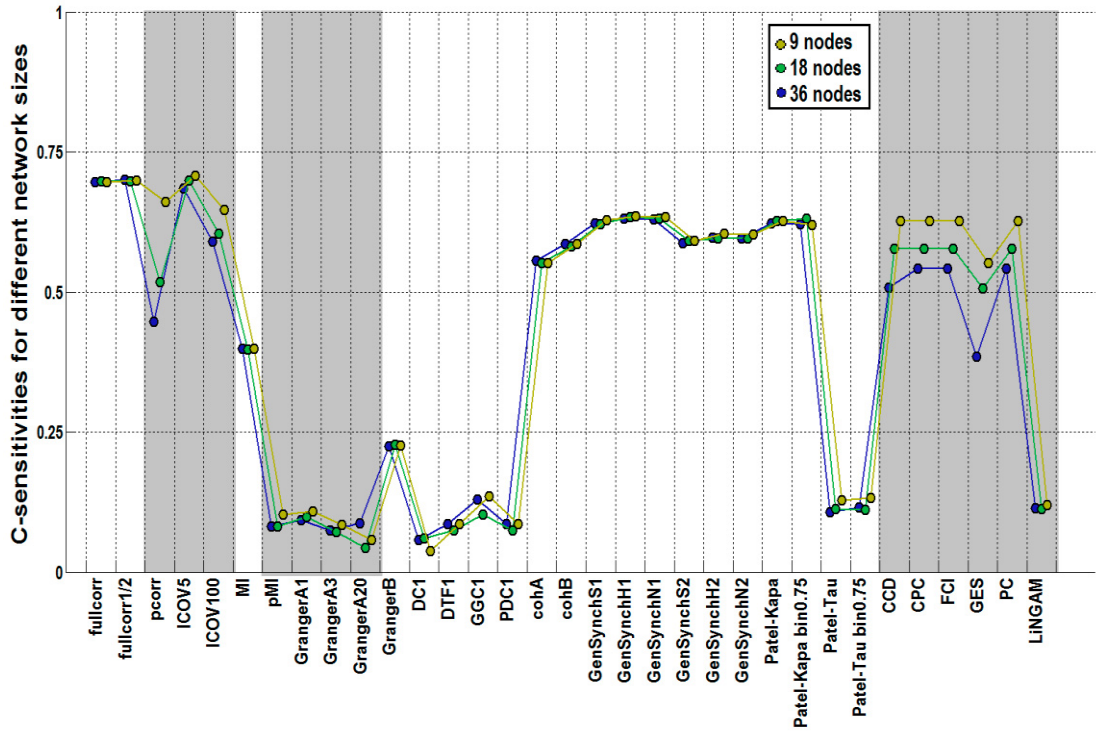

Figure 5. C-sensitivity values plotted for different network sizes. The visual cortex networks considered are those of the quadrant (9 nodes), hemisphere (18 nodes), and both hemispheres (36 nodes). All connections tested lie within a quadrant, as shown in figure 1B. The session duration tested is 17 minutes. Shaded methods are multivariate and expected to vary with different network sizes, other methods are bivariate and should show no more than small variation due to rounding.

Without removal of the global effect

We next tested the case where we do not remove the global signal in preprocessing. For the majority of the methods, the effect of including the global effect was to decrease the separations between normalized interaction measures of expected and unexpected connections.

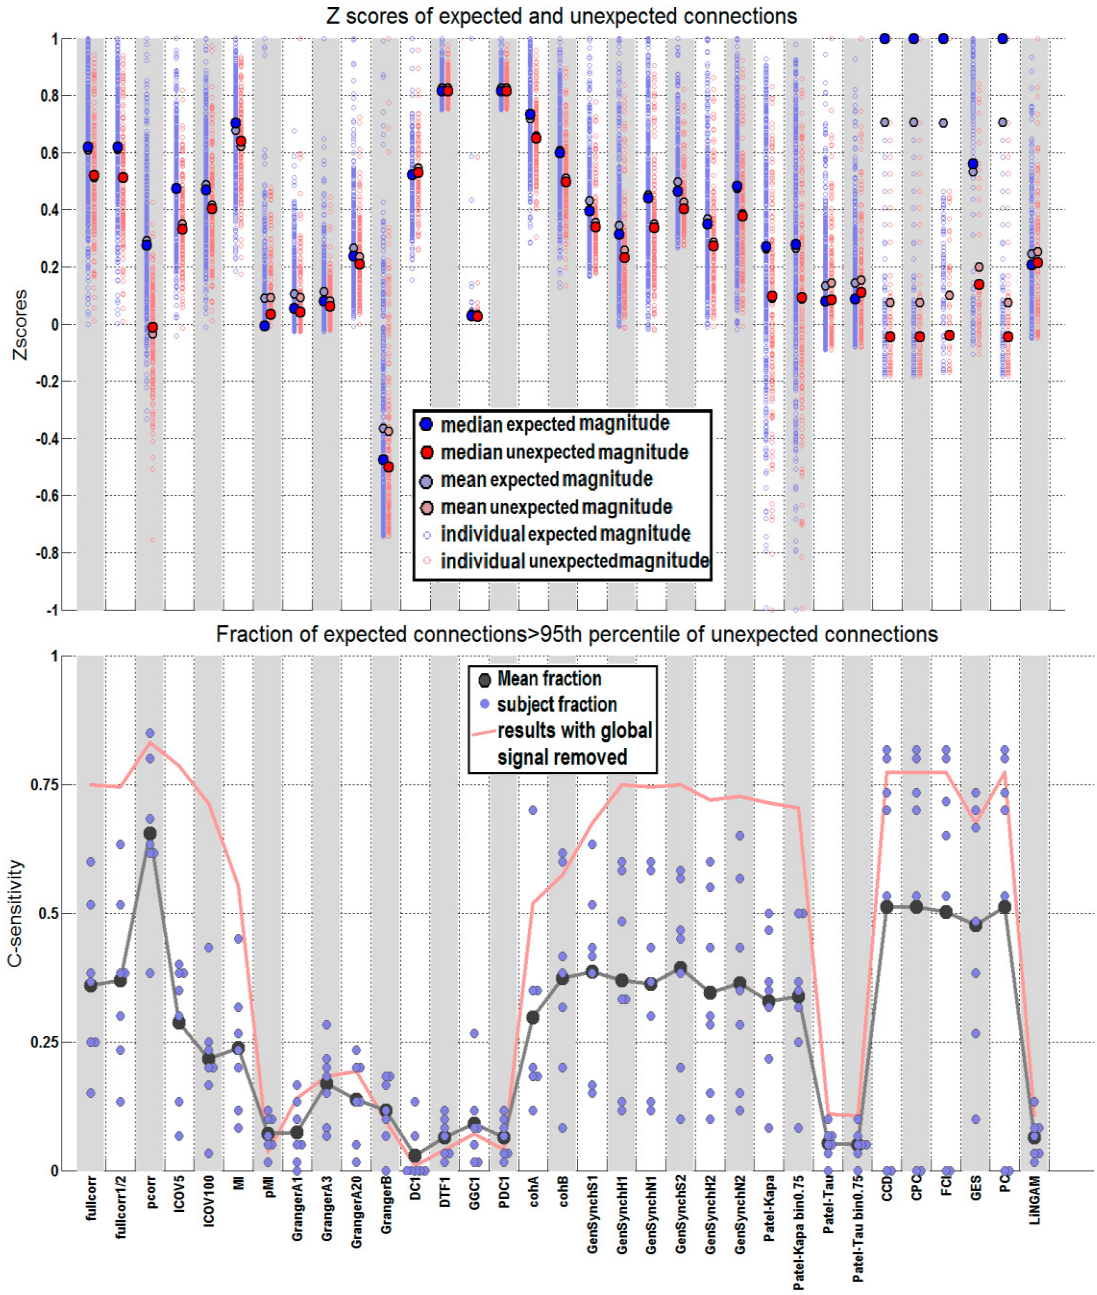

Figure 6. Evaluation of the network modeling methods using resting-state data without the removal of the global signal (here approximated as the first principle component). Results are for the 9 node network with 68 minute fMRI sessions. Panels are as in figure 3.

The separations of expected and unexpected connection magnitude means are now from 30% to 53% for the BayesNet methods, 19% for pcorr, and from 7% to 12% for fullcorr, fullcorr1/2, ICOV5 and 100, the GenSynch Measures, cohA and B, MI, Patel-Kapa, and Patel-Kapa bin0.75 (Fig. 6, top panel).

The distribution of c-sensitivities across models with the global signal not removed (Fig. 6, bottom panel) is significantly lower than the distribution obtained in our main test which does include removal of the first principal component (Fig. 3, bottom panel). By far, the top model in this case is still pcorr, achieving here only 65% c-sensitivity; the next best models are the BayesNet models with 48-51% c-sensitivity. The other top models see drops in c-sensitivities into the range of 22-39%.

Effect of adjusting the threshold: ROC analysis

The choice of threshold involves a trade-off between a strict threshold which eliminates nearly all false positives while potentially eliminating true positives, and a more lenient threshold which obtains more true positives but also increases the number of false positives. In order to best demonstrate this trade-off, we plotted receiver operating characteristic (ROC) curves (Fig. 7). Here we define a true positive (TP) as an expected connection which does indeed have interaction magnitude greater than the

set threshold. Likewise, a false positive (FP) is an unexpected connection with interaction magnitude greater than the threshold. Since we recognize that most often 68 minutes of resting data is not available, we plot an ROC curve for the cases of 17 minute sessions (Fig. 7A) as well as 68 minute sessions (Fig. 7B). In both of these plots, there is a clear separation across all threshold levels between the more c-sensitive models (correlation based methods, Patel's Kapa, GenSynch, coherence, and BayesNet methods) and all others. With the increased session duration (Fig. 7B), MI also separate itself from the poor models and begins to show similar trends as seen in the better models. When the sessions are long, most of these top models, such as pcorr, ICOV, fullcorr, and some GenSynch and BayesNet, have curves with small slopes even at low FP fractions (for example, pcorr has a slope of 0.34), indicating that not much is to be gained by decreasing the threshold, and that these methods have very good separation between the interaction magnitudes of the expected and unexpected connections. Others, like Patel's Kapa and some of the GenSynch methods, have steeper slopes (for example, Patel's Kapa has a slope of 1.2) indicating that there would be more gain in terms of TPs with increased allowance for FPs. However with the shorter sessions, the slopes are somewhat steeper in general at low FP fractions (for example,

pcorr has a slope of 0.96) in which case steps of 0.05 in increased fractional FP could result in an increase of nearly 0.05 in fractional TP.

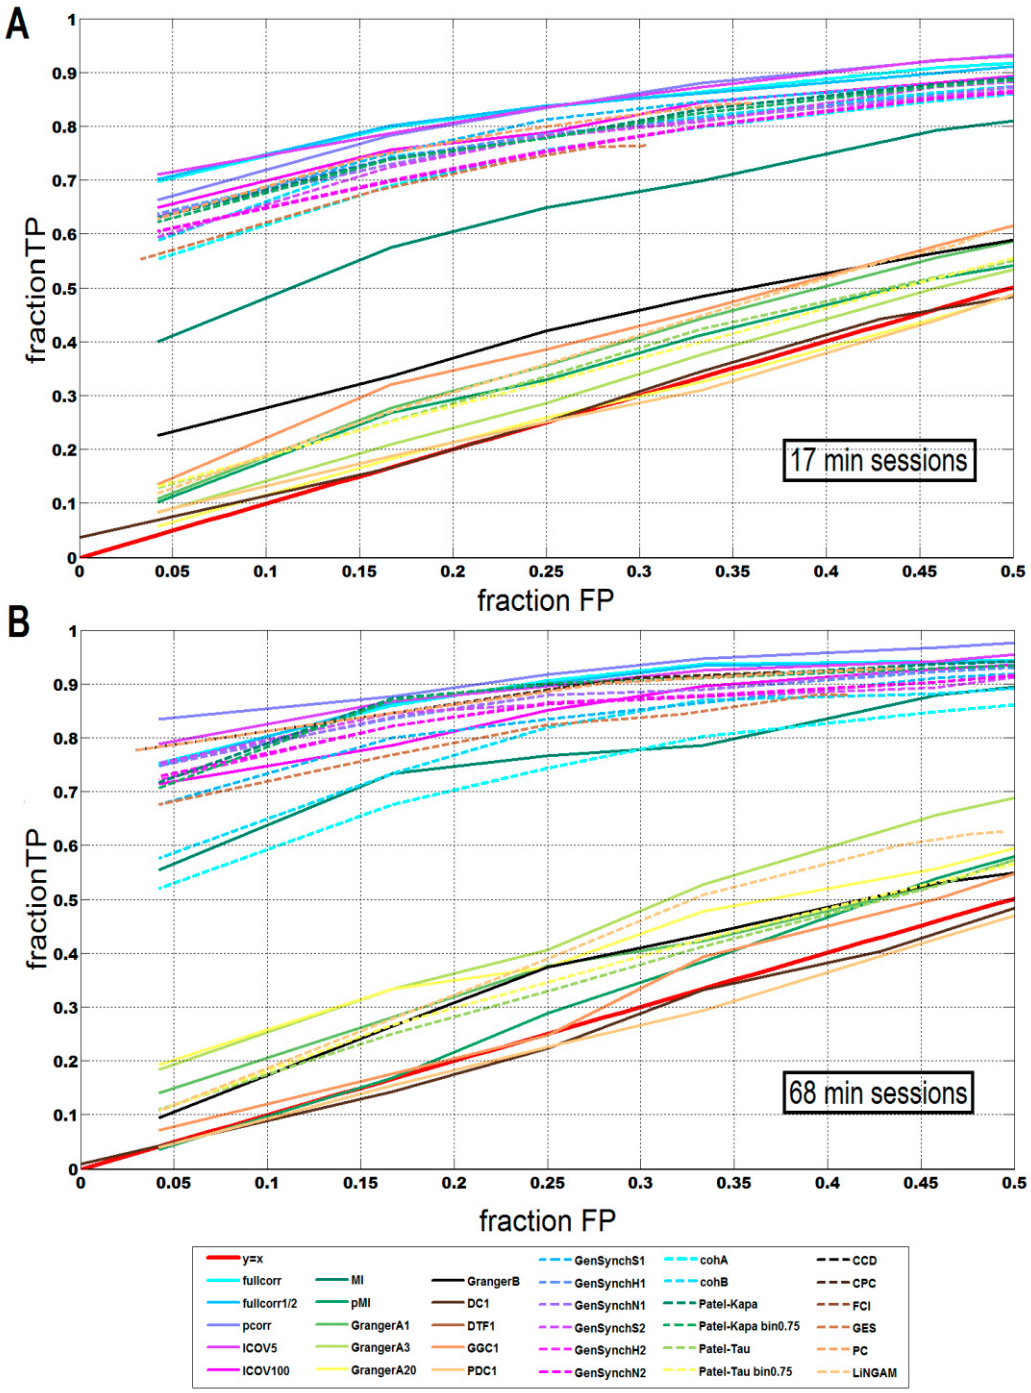

Figure 7. Receiver operating characteristic (ROC) curve of the models in the case of the 9 node network with 17 (A) and 68 (B) minute sessions. Plotted is the fraction of expected connections surpassing the unexpected connection threshold (true positives (TP)) versus the fraction of unexpected connections surpassing the threshold (false positives (FP)) (this value is generally equivalent to 1-threshold with the rare exception of when multiple connections with the same interaction strength exist at the threshold).

Discussion

C-sensitivity performances of the methods

The main analysis that we conducted was around a 9 node network, with 68 minute sessions and with the global signal regressed from the ROI time-series. These choices were made because we expected they would optimize results given preliminary analyses. In this analysis, we observed the best performances by the correlation methods, Petel's Kapa, the Bayesian Network models, and the GenSynch measures. Partial Correlation was expected to do well in separating the direct from indirect connections because it uses the information from all nodes in the network to assess whether a third node may be an intermediate in the communication between two nodes. Partial Correlation's results meet our

expectations, noting particularly in Fig. 3 that the interaction magnitudes of the unexpected connections scatter around zero correlation.

Where full correlation is concerned, the c-sensitivity results may deceivingly imply general success, but this is not the case as it does not do well at identifying only direct connections. Here it is successful because the interaction magnitudes of the direct connections were reliably greater than the magnitudes of the non-existing connections (Fig. 3), but its success depended on the calibration method we applied. The calibration method used here by setting a threshold at 95% of interaction measures of non-existing connections implicitly changes the utility of this method such that it becomes more sensitive to direct connections. Without prior knowledge of non-existing connections, full-correlation would end up with relatively high rates of false positives or low rates of true positives.

This appears to also be the case with the GenSynch measures (also only considering two nodes at a time), whose unexpected connections do not center near zero, but do maintain a sizeable separation in the interaction magnitudes of the direct, expected connections, and those expected to not be direct. The correlation methods and GenSynch methods, while differing in approach, have in common their simplicity, and the idea that if the signal magnitudes vary similarly over time in two nodes, there is a likely

connection existing between these nodes. Correlation considers given time points and examines the signal amplitude, while GenSynch considers given signal amplitudes and examines the time occurrences.

The Bayesian network models as well as Patel's Kapa and Tau measures employ Bayesian criteria which use information from previous time points to determine a relationship between nodes. The Tau measure doesn't perform on the same level as these other two models as it is conceived and optimized to be a measure of directionality, rather than of connection strength.

The Bayesian Network models in particular are very impressive with clear separation between expected and unexpected interaction magnitudes. These approaches incorporate a time influence consideration in order to determine a general likelihood of influence, not in order to observe a directly linked pattern of progression over time, as the lag-based methods do. This utilization of multivariate probabilistic analyses over a range of time points is able to take advantage of a larger pool of information all at once, more than most other models considered, which likely contributes to these models' successes. Given that the relationship between nodes over time is a factor in these models, a faster repetition time in fMRI data acquisition may improve the results of these methods.

We also note that as the interaction magnitudes obtained with Bayesian Network models for the expected and unexpected connections were very well separated in all tests, these models may be the best among those tested here for use when no prior knowledge of the network is available.

Mutual Information and Coherence are mid-performance models here. The coherence models calculate an interaction measure similarly to correlation, using cross wavelet power distributions. This similarity to correlation is a probable explanation for why it sees relative success. Mutual information also has its simplicity in common with Correlation and GenSynch. These simple approaches are quite intuitive as to why their criteria would correctly predict a connection.

While many of the lag based methods are also intuitive, they did not perform well in our analysis. We expect that we did not have the time resolution one would require in order to detect the flow of information in the tight, highly connected network we considered. Here we completed our analyses with data having a TR of 2 s, thus the period between time points in the data is far greater than the time required for flow of information through areas (on the order of 100ms) (Schmolesky et al., 1998). Different experimental parameters or different variations on the models may prove to be successful at establishing connections between brain regions.

In the case of the LiNGAM method, it may be simply that there was not enough data to allow it to reach its potential, although the session duration of 68 minutes used here is already in a non-practical, non-economical range. Smith et al. found that LiNGAM did poorly even with 60 minutes of simulated data, and only once there was 4 hours worth of data did it begin to see good c-sensitivity results. Given these observations, we expect that LiNGAM is not a good choice for resting-state networks analysis.

The Smith et al. simulation which has parameters closest to ours is simulation 6. In simulation 6, there were 10 nodes used, 60 minute fMRI sessions, a TR of 3 s, 1% noise, and variations in HRF (haemodynamic response function) delay of standard deviation 0.5 s. The top scoring models were pcorr, ICOV5 and 100, and the Bayes Net methods, followed closely by fullcorr, fullcorr1/2, the GenSynch methods, CohA and B, MI and pMI, and Patel's Kapa. All of these methods scored c-sensitivity above 0.8. The performances of the models applied to real data in our study compare quite well with their performances with simulated data.

Effect of session duration

With increasing session durations, the methods which tended to be more successful gained further success, and those methods which tended to be less successful showed inconsistent trends. These results are intuitive because an ineffective method will be ineffective regardless of the amount of data provided. Along the same lines, these findings strengthen our confidence in those methods we classify as successful. It is promising to see the effective models showing improvements, particularly when the amount of data is initially doubled from 8.5 to 17 minutes, with jumps of 0.2-0.3 in c-sensitivity of the best models. This is compared to other data doubling (17 to 34 minutes, and 34 to 68 minutes) which result in increases of less than 0.2 in c-sensitivity. We conclude that for any of these methods to be effective with resting state data obtained by conventional data acquisition methods at 3 Tesla from a network with more than just sparse connections, sessions of 8.5 minutes do not provide the models with sufficient information to make meaningful predictions. Given our results, the minimum amount of data, in an experiment with comparable parameters to ours, should be 15-20 minutes. Additional data beyond this will likely improve results further, but one may want to consider the trade-

off between small incremental improvements and freeing subject and magnet time for use on other tasks.

The above mentioned trend was also observed by Smith et al. in their simulations, where the best models showed greater success with 60 minute sessions as compared to 10 minute sessions. They also found that increasing the data further to 250 minute sessions allowed some of the less successful models to achieve high c-sensitivities, as in the case of the Granger methods which obtained up to 90% c-sensitivity. The importance of averaging has been emphasized also in task-based functional MRI: Saad et al. (2003) showed that brain regions considered to be inactive may pass the statistical threshold following considerable averaging.

Effect of complexity of network

In Fig. 5 we see that the complexity of the network, or the number of nodes input into the models, does not impact the ability of most models to separate direct and indirect connections. First, it is important to note that most of the models are bivariate, i.e. they consider two nodes' time-series at a time, and so the number of nodes in a system is indeed expected to have no effect. This expected result is what we see in models such as fullcorr and GenSynch. In multivariate models, though, additional nodes could have an effect on interaction predictions, due to the increased

number of degrees of freedom that multivariate methods need to take into account. Indeed, multivariate model saw decreased performance with increasing number of nodes, although the effect was moderate, especially for ICOV. This is likely the case because the connections we test are selected for their strength relative to other connections in the visual system. We hypothesized that any resting co-activations between regions outside of the quadrant would be significantly weaker than the interactions within a quadrant, based on the connections in the monkey model, thus the influence of out-of-quadrant nodes on the connections of interest should be minimal. Our results suggest that these hypotheses are indeed valid.

Effect of the global mean confound

The origin and significance of the global mean time-series as a component to overall measured fMRI time-series is currently not well understood and thus there is disagreement as to whether the mean should be regressed out, or if it should remain in the signal. Here we assessed our results with and without the removal of the global mean confound. Our hypothesis in this regard was that regressing-out the global effect, irrespective of its origin, enhances the detection of connections known to exist at the anatomical level (as observed in Fox et al., 2009; Carbonell et al., 2011), and that this step results in no detriment to the analysis. Our

results here corroborate these previous observations, as seen in Fig. 6. Thus, in order to reveal direct connections in a network, irrespective of the origin of the global signal, we recommend regressing-out the global effect by means of an appropriate method (Carbonell et al., 2011).

Other potential confounds

The simulations done by Smith et al. covered several potential confounds and demonstrated the expected effect if indeed they are present. Here we discuss how some of these confounds and simulations may relate to the results we observe with real data.

GenSynch

Smith et al. observed results quite similar to ours in their simulations 13 and 22, where the GenSynch methods perform on a similar level to the correlation methods. In simulation 13, they modeled the presence of backward or inhibitory connections between nodes and in simulation 22 they modeled non-stationary connection strengths. The GenSynch measures aim to model systems with greater complexity than many of the other models, taking into consideration clusters of similar signals, rather than just individual signals; this added complexity is a likely explanation for their success at modeling real networks in our study which may have any

number of confounds, such as backward connections and non-stationary connection strengths (Chang and Glover, 2010; Sotero et al., 2012).

The GenSynch measures were among the most sensitive to changes in strengths of connection, a conclusion made from the results of simulation 21 (Smith et al., 2011). When connections strengths were increased in simulation 15, the GenSynch methods saw improvements of approximately 15%. Given that the connections in lower visual cortex that we relied on are robust, this is another potential reason for GenSynch's heightened performance in our tests relative to the simulation tests.

External inputs

External inputs affecting the apparent interaction of nodes in the network by jointly activating unrelated nodes are an unlikely confound in our network. This is because of the elevated retinotopic nature of V1, V2, and V3 among areas of the visual system. Inputs into the visual cortex come from the lateral geniculate nucleus into V1 and from there are spread to other visual areas such as V2 and V3. That being said, since we are looking at resting state data, there is no prominent input arriving from the lateral geniculate and the result is that external inputs are not a major factor. It may be possible that information could be fed back through the higher visual areas not considered in our network or through areas in other

quadrants. We believe, however, that the impact of this feedback is not likely to be significant when compared to the strength of connections between V1, V2, and V3 within a quadrant. The reason for this belief is that V1, V2, and V3 show by far the finest retinotopic specificity when compared to all other areas. Therefore, the relative effect of non-retinotopically-specific feedback must be small.

Introducing Calibration in functional connectivity studies

As evident in Fig. 3, many unexpected connections are predicted by the most c-sensitive models with non-zero interaction magnitudes. This means that these models are not able to distinguish the true positives from the false positives without the use of additional statistical criteria, such as a threshold. Even when applying thresholding, an arbitrary choice of threshold will result in either a high rate of false positives or in missing direct connections, due to overlap in the distributions of false and true positives. Thus, in order to overcome this shortcoming, we propose to integrate a data-based calibration stage in all network studies, as we have done here using the false positive based percentile threshold. Here we have used a form of calibration in order to evaluate network analysis methods. However, we propose that similar calibration can be applied to all data and analysis types in order to determine a statistical threshold that will

minimize false positives and maximize true positives. A crucial condition is the prior knowledge of existence and non-existence of part of the connections for which data are available. In this context, we establish a set of connections expected to exist, and another set of connections we expect to not exist. The unexpected connection interaction magnitudes output by the model in use are sorted and the magnitude of a certain percentile of these connections is utilized as a threshold. This threshold defines the interaction magnitude above which all predicted connection magnitudes must be in order to be considered significant. The more expectations that can be set with confidence, the more robust this calibration will be. The percentile is chosen based on the confidence one requires. We note that calibration also overcomes variations in results due to data acquisition parameter differences.

Conclusion

Throughout tests done here with human resting state fMRI data, we find the most consistently successful and reliable methods in predicting direct neural networks to be Partial Correlation and Bayesian network models PC, CCD, FCI, or CPC. This being said, with the implementation of the calibration method we propose, Correlation, Regularized Inverse Covariance, and General Synchronization methods are also consistent in

their success and perform on a comparable level to Partial Correlation and BayesNets. In our 9 node, one visual quadrant test, we found the previously mentioned methods performed best, all predicting between 67 and 83% of true connections with a magnitude of interaction greater than the 95th percentile of the unexpected connection distribution. Importantly, we found that the fMRI session duration is very important to the successfulness of the network models, where longer sessions increasingly showed better results. Use of fMRI data with longer session length, reduced TR, and increased signal to noise ratio are all expected to increase the success rate of methods tasked such as here. We propose to integrate a calibration stage to functional connectivity studies, utilizing a threshold based on false positive statistical results for connections known to not exist. Such a calibration will reduce false positives and maximize true positives, making the results more reliable. This calibration is a vital step in any network study because false positives are a significant presence in the outputs of even the best models, and as such, we recommend implementing it in future studies of neural networks.

6- Discussion

In the first part of the discussion, we address areas we felt could use some additional, more detailed commenting on issues surrounding the manuscript presented in chapter 5. These comments go beyond the succinct discussion format in the manuscript. These include discussions on network selection, length of connections, ROI border definition, number of voxels in an ROI, and changes in results given a more conservative set of expected connections and more conservatively defined ROI borders. In the second part of the discussion, we discuss ideas raised by the thesis in a more general manner and in a broader context. This part includes a discussion on the Bayesian Network models in greater detail than in the manuscript, including an in depth look at the c-sensitivities and a comparison with Granger Causality, remarks on group analyses, and the significance of the work.

V1, V2 and V3 as a network

Here we explore the validity of various network modeling methods by evaluating their successes at predicting expected anatomical connections between different regions in the human visual cortex. We

chose retinotopic visual areas V1, V2, and V3 for pursuing this evaluation because of the in depth understanding we have of the anatomical connections within and between these areas based on studies of the macaque monkey brain. Although including ROIs from other areas could make the model more complete – the overall SNR would be worse, because other areas are smaller, and their retinotopy tuning is much broader. We preferred to use a 'clean and clear' network model based on V1, V2 and V3 to test the different methods.

Length of connections as a confound in resting-state functional connectivity

This discussion is included because of the fact that in our study, most of the expected connections are shorter than those in the unexpected connections set. Thus, it is reasonable to try and create a balance of longer and shorter connections and/or look at the effect of the length of a connection on the interaction strength in order to account for any impact this might have. The nature of our study is such that we selected a set of connections we could be very confident about in determining whether or not they would exist directly. This led to many close range connections that are expected and unexpected connections which were longer. We cannot test different or additional connections because current knowledge

does not allow classifying them confidently as expected or unexpected. Also, since retinotopy is a spatially continuous mapping from visual space to cortex, distance constitutes a confound when pursuing resting-state functional connectivity in the visual cortex. All of this being said, our aim here is to compare the performance of various network models, not to determine functional connectivity. Thus, we currently take no action in this regard, though we recognize that it is something to be aware of. A future study comparing the strengths of connections of different lengths is of interest, though it is outside of the scope of this Master's thesis.

On the horizontal connections and the scale of our ROIs

Horizontal connections could possibly introduce a confound in the way we have conducted our analyses. Amir et al. (1993), Rockland and Lund (1982), Rockland et al. (1982), Sesma et al. (1984), and Lyon et al. (1998) observe that intra-areal connections in early visual areas appear to be relatively localized and rarely reach distances of more than a few mm, at least in area V1. Thus, if we want to test for horizontal connections between eccentricity areas, the voxels we should consider would best be selected from a narrow strip of cortex along the eccentricity border. Given our methodology for normalizing the number of voxels in each ROI, not only are we not limited to that narrow strip near the border, we exclude a

number of voxels farthest from the centre point of the ROI. These considerations could bring into question whether the within-area connections are indeed predicted by known anatomy. Because of this, we have removed the horizontal connections from our set of ground truth expected connections and redone the analyses to observe any changes. Therefore, in our revised analysis, the expected and unexpected connections are based on inter-areal connections alone. Results with this new set of expected connections are presented in a section below along with a discussion of the relevance of the results.

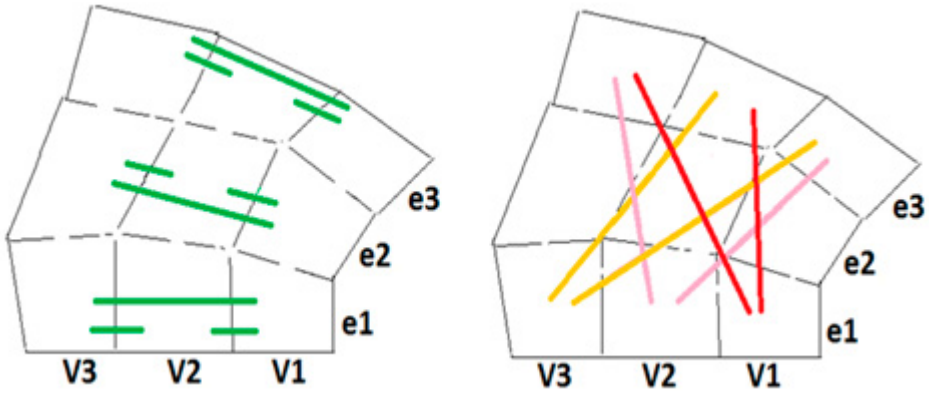

Figure 1. The reduced set of connections tested in the results below, horizontal connections within visual areas having been removed.

Definition of the ROIs

We took the occasion of rerunning the analyses to re-evaluate our definition of the ROIs. The original assignment of a voxel into a visual area was done such that the area which contained the highest weight of that

voxel was said to contain it. This meant that some of the voxels along the borders between areas would have been partially contained in the adjacent visual area. Given the results in Smith et al. (2011) regarding the blending of information from ROIs, or bad ROIs, we decided it was best to be stricter about the voxels we selected. The new ROIs are defined with a gap between visual areas V1, V2, and V3, by means of selecting voxels with polar angle within $\pm$ 39.5° of the oblique meridian in each quadrant of the visual field. This leaves a 5.5° gap in polar angle between each included voxel and the closest possible voxel in another area across the inter-areal border, with included voxels being exclusively contained in a single visual area. Between different eccentricity regions within the same visual area, the criteria are more relaxed, attributing a voxel to the region which contains the largest proportion of its volume. These relaxed criteria are necessary in order to maintain higher voxel counts within each ROI. Since we do not include within area connections in the set of expected connections, it is less important to have a strict gap between ROIs. Importantly, in practice, there was most often a gap between adjacent ROIs from within the same visual area due to the voxel selection process. Voxels are selected starting from the middle of the eccentricity region and then moving outward toward the edges until the total number of voxels to

be selected is met. In this way, most often only a few ROIs per subject have voxels very near the eccentricity borders. Below is a figure presenting just this for subject 5.

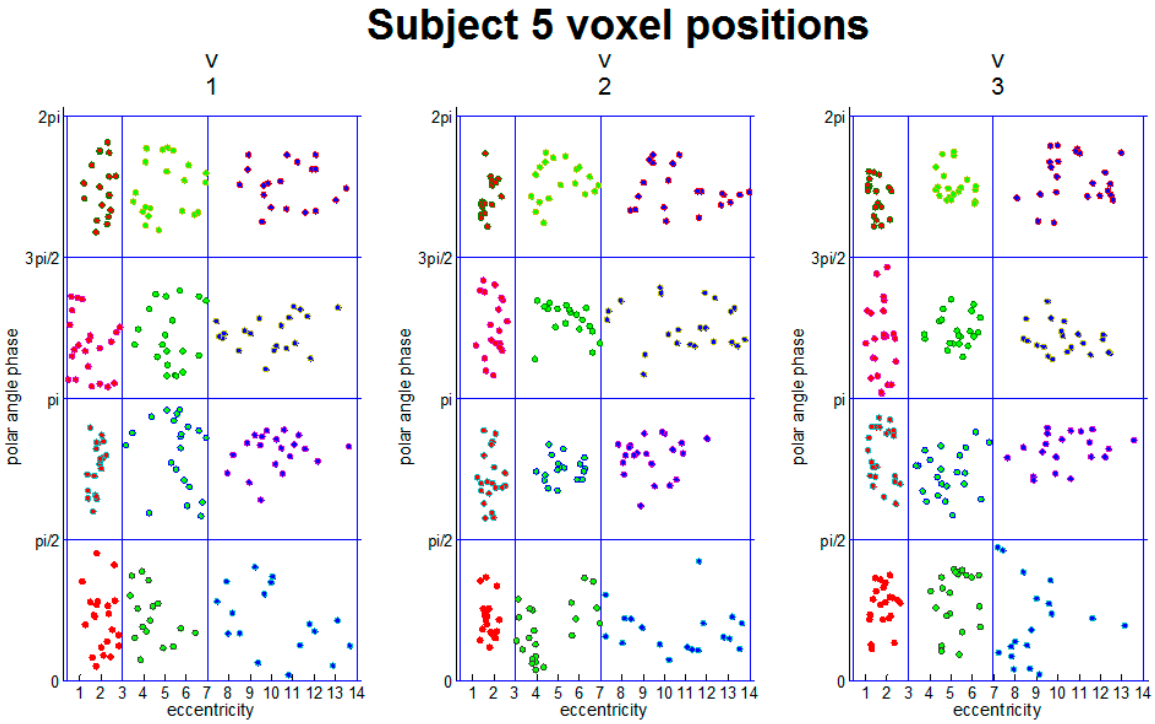

Figure 2. Positions of voxels selected for inclusion in ROIs for subject 5. Each plot represents a visual area. Within a plot, central eccentricity region voxels are in red, intermediate are in green, peripheral are in blue, and the eccentricity borders are represented by vertical lines. Horizontal lines separate quadrants. Notice how only the V2E2-V2E3 border in the 0 to pi/2 cluster (quadrant) has voxels near it on both sides.

Number of voxels in an ROI

During the construction of our new ROIs, we originally tried to select the smallest maximum voxel count across ROIs for a particular subject, however we found that with this new criterium of exclusivity of voxels in a visual area, that in some cases the voxel count was very low (under 10) and in others much higher (over 30). We sought to maximize the number of voxels for each subject as well as make the amount selected across subjects more comparable. In the end, we remedied low voxel counts by identifying the forth lowest in the set of maximum voxel counts for each ROI of a particular subject and followed this with the selection of that number of voxels from each ROI (of course this is with the exception of the 3 ROIs which had a lower maximum voxel count: their maximum was selected). This also led to somewhat more comparable numbers of selected voxels across subjects (25-39). The exceptions to this were subject 2, 3, and 6 who had lower 4^{th} lowest voxel counts of 23, 14, and 19 respectively.

All analyses have been redone with the new ROIs and the new set of expected connections. Below are new figures, showing the results with only the new set of expected connections, followed by the results with the new ROIs and the new set of connections.

Results from the analyses using the original ROIs, excluding horizontal connections in the set of expected connections

In this case, the trends in all tests remain the same. The change in results is a general decrease in c-sensitivity of about 0.1 for all models in all tests, and slightly steeper ROC curves. The fact that c-sensitivities decreased in general, indicates that the horizontal connections were found to influence the network structure more often than not, and that the inclusion of these connections did not cause less than perfect c-sensitivities. Instead, it was the inter-areal connections which were the major source of false negative results. Possibly this is also evidence that some slightly longer range horizontal connections do exist or that the few voxels which do lie closer to the edges of the eccentricity areas reflect strongly enough the horizontal connections that using ROIs which include some voxels outside the range of these connections does not mask the coactivation.

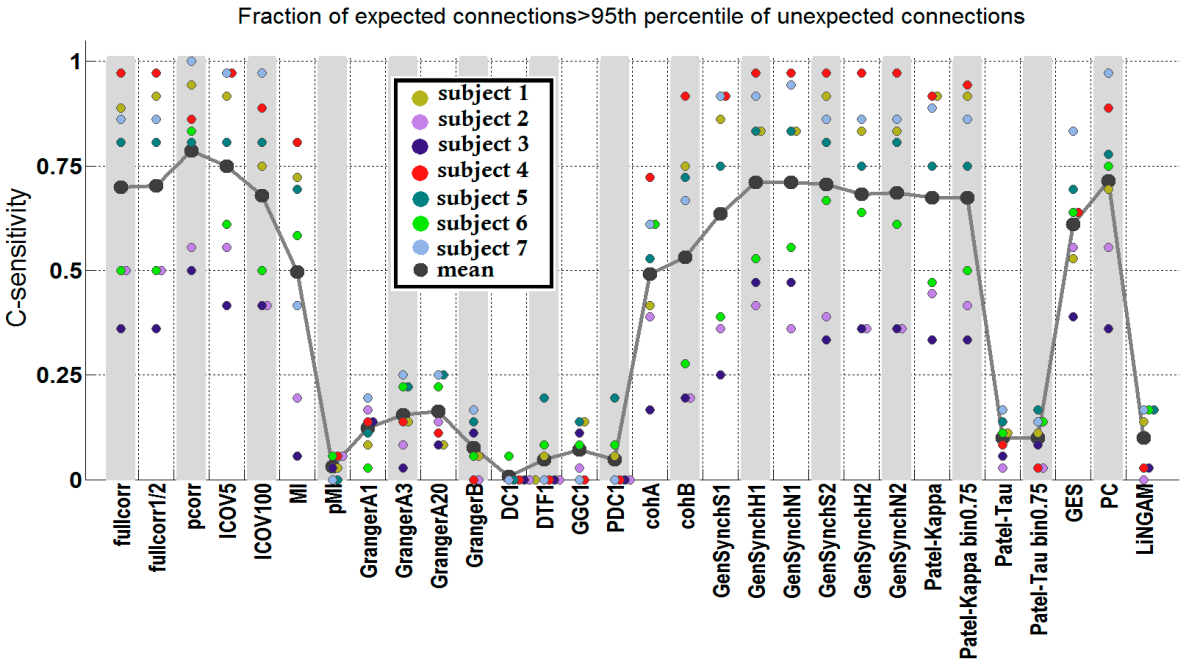

Figure 3. C-sensitivities of the models without the inclusion of horizontal connections in the set of expected connections.

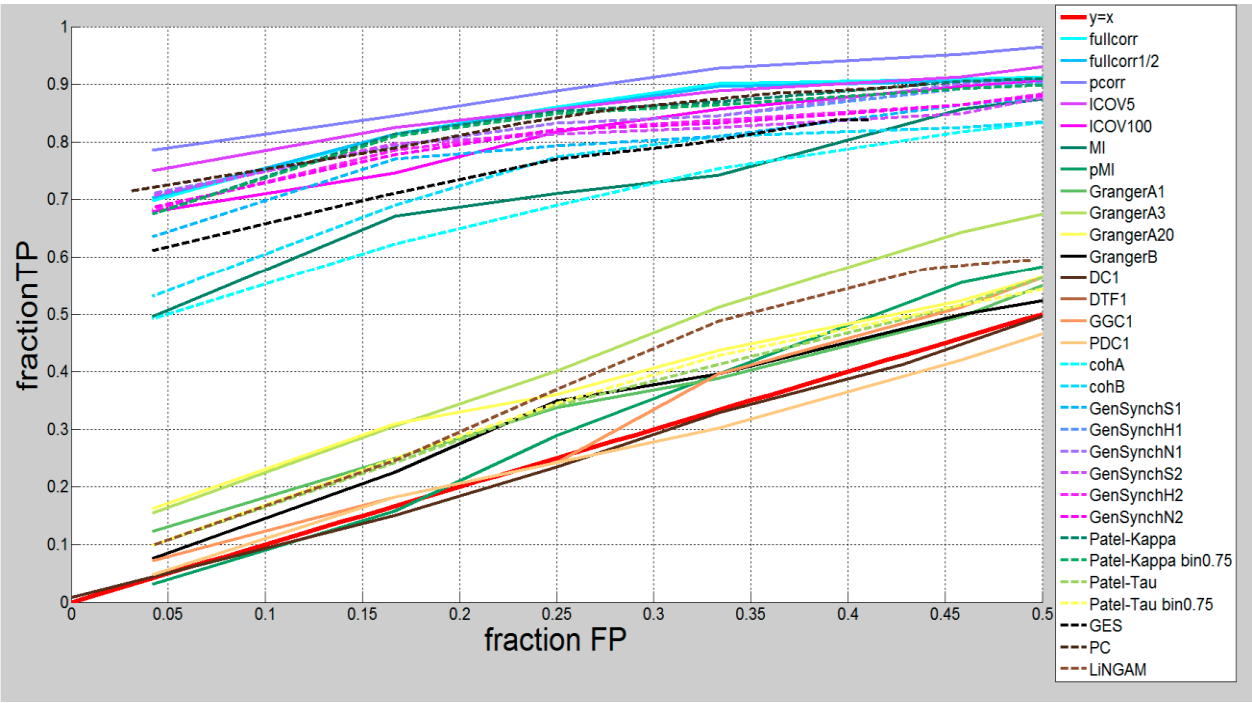

Figure 4. Receiver operating characteristic (ROC) curve of the models in the case of the 9 node network and 68 minute sessions without inclusion of horizontal connections in the set of expected connections.

Results from the analyses using new ROIs, excluding horizontal connections in the set of expected connections

In this case, with gaps between the ROIs and the exclusion of the horizontal connections from the set of expected connections, we noticed a clear separation in the performance of subjects. In the figure below, we see that subjects 2 and 3 have consistently low c-sensitivities, subject 6 has mid-range c-sensitivities, and the remaining subjects have very good c-sensitivities. Recall that subjects 2, 3, and 6 had lower 4th lowest maximum voxel counts than the other subjects. Subject 3 had a very low 4th lowest maximum voxel count relative to all other subjects as well as much lower voxel counts in the three ROIs with smaller maximum voxel counts (9, 10, 12). These 3 ROIs were also involved in a high proportion of tested connections (10 of 15) making it such that the low voxel numbers contributed highly to the measured results. While subject 2 had only somewhat fewer voxels selected from each ROI, the 3 ROIs which had fewer voxels were also involved in a higher proportion of the tested connections (8 of 15). Although subject 6 had a relatively low 4th lowest maximum voxel count, a smaller proportion of the connections tested involved the ROIs with lower voxel counts (4 of 15). The other subjects had high voxel counts (>22) even in the ROIs which had less than the 4th

lowest maximum voxel count, and so regardless of the participation of the ROIs in the tested connections, sufficient voxels were involved. The consistently poor results together with a reduced voxel count in ROIs important to the connections in our analyses, led us to the exclusion of subjects 2 and 3 in the tests done with the new ROIs. We consider the poor c-sensitivities of these subjects to be due to the lower numbers of voxels in the ROIs and thus their exclusion should not reflect any bias or data selection. We note that in the previous analyses with the original ROIs, the number of voxels contained in subject 2 and 3's ROIs were somewhat more comparable to the other subjects, and while they were still among the less c-sensitive subjects, they were not as notably substandard in performance. See figures 3 and 5 for the performances of each subject with the original ROIs and the new ROIs, respectively.

Care should be taken in ROI definition to include as much voxel information as possible in order to improve the sensitivity potential of the results. This concept holds in voxel or pixel based ROI analyses outside of those focussed on connectivity, as increased sensitivity is a well known benefit of having a greater number of elements to average in a group study (in this case, a group of voxels). This being said, there should still be a

strong confidence that voxels do indeed belong to the ROI, otherwise results may differ dramatically (Smith (JF) et al., 2011).

Considering the c-sensitivities of the models using only subject 1, 4, 5, 6 and 7, we find that the trends in all tests are similar to those with the original ROIs. The change in results observed relative to the results presented above with the new set of expected connections and the original ROIs is a general increase in c-sensitivity of about 0.1 for all models in all tests. This improvement may be in part explained by removal of the poor results from subjects 2 and 3 however there are other effects at play here evidenced by the differences in subject c-sensitivity distributions between the old and new ROI tests.

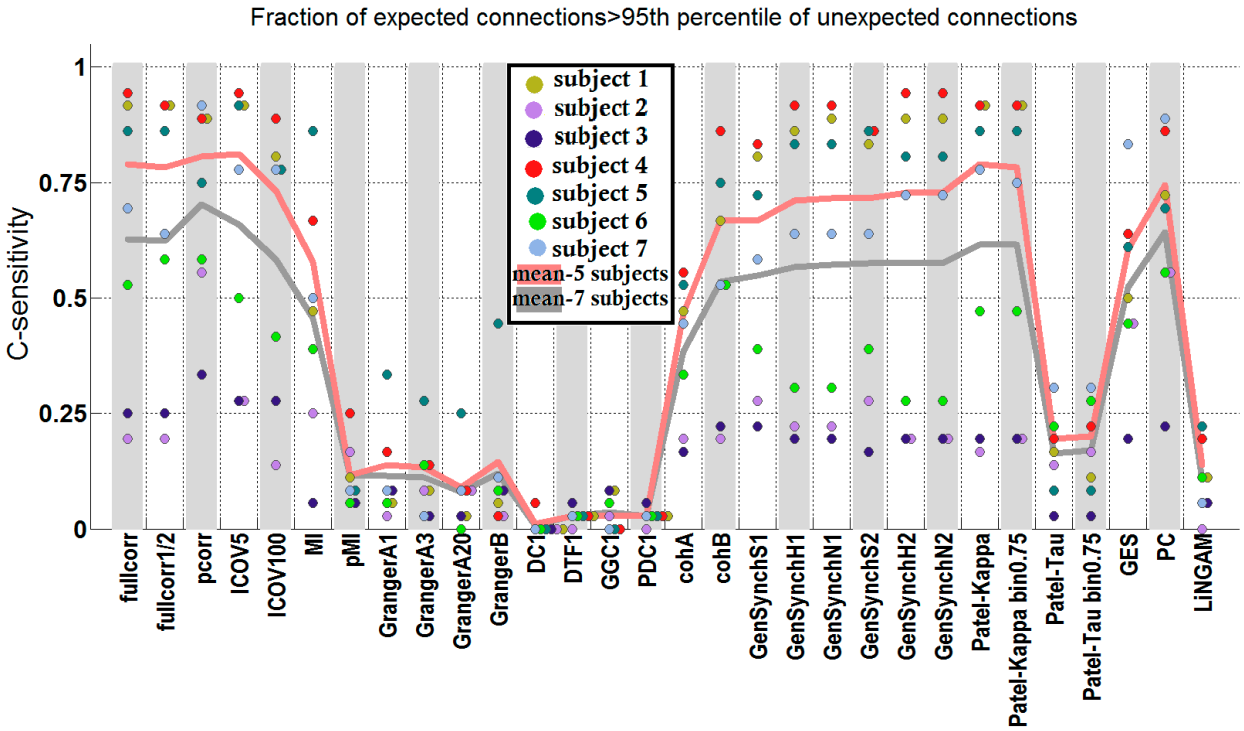

Figure 5. C-sensitivities of the models without the inclusion of horizontal connections in the set of expected connections and using the new ROIs. The mean is computed across all subjects (grey line) and across subjects 1, 4, 5, 6, and 7 (pink line).

One test which did provide results different form our tests with the original ROIs was the testing of the impact of the global signal. With the new ROIs (excluding subjects 2 and 3) we saw smaller change in the c-sensitivities of the models than we did with the old ROIs, the difference being of less than 0.25 for all but ICOV100 (recall that with the original ROIs many models saw drops greater than 0.3 in c-sensitivity). In most cases the removal of the 1st principle component resulted in somewhat better c-sensitivities than without its removal, this result being the same as we had originally observed.

The main purpose for the redefining of our ROIs was to be sure that we had no contamination of data from outside of the intended visual area. With this done, it appears that perhaps the impact of other confounds, such as the global signal, become less pronounced. Indeed in the tests done with shorter sessions, the drop in performance was less pronounced than with non-exclusive ROIs, particularly for the change from 8 to 4 runs where there was no drop in general. Our results further support the finding by Smith et al. (2011) that careful definition of ROIs should be a main concern in designing an ROI-based connectivity analysis.

C-sensitivity of Bayesian network models

The c-sensitivities for Bayesian Network models PC, CCD, CPC, and FCI (equivalent in the determination of existence of a connection), which had excellent separation in general of the expected and unexpected interaction magnitudes, are on a similar scale to the correlation methods which had much less impressive interaction magnitude mean separations. The BayesNet c-sensitivities are not greater because of the presence of a few higher-magnitude unexpected connections which drive up the unexpected connection-based threshold. These false positives (unexpected connections of magnitude greater than the threshold) were present in all subjects and the connections which occurred as false

positives varied. The higher threshold resulted in some expected connections being below the threshold.

Bayesian Network models and Granger Causality: a closer look

The contrast between the excellent performance of Bayesian network models and the poor performance of Granger causality models motivates a closer look at how BayesNet method PC determines the existence of a connection, followed by how Granger Causality does. We note that the two methods use amplitudes and time in order to determine connections; however, we found that their performances were different. Note that BayesNet methods PC, CCD, CPC, and FCI are equivalent in how they determine the existence of connection and differ only in their decisions on direction of connection, which we do not consider here, so we only discuss PC.

PC:

As mentioned in the manuscript, the Bayesian Network models assume that each node in a network has a set of parent nodes: some subset of nodes in the system which was activated before the node in question. The Parent subset is defined such that all earlier activated nodes not in the subset are conditionally independent from the node in question. It is in the stage where the conditional independence between

the nodes is determined that the existence of a connection is predicted. In the case of discrete data sets as we have used, the conditional independence of a set of nodes is determined using a g-square test. This statistical test looks at the frequencies of the variables having given amplitudes and determines conditional independence by comparing the expected frequency to the observed frequency at those amplitudes. I.e.,

$$G^2 = 2\sum(\text{Observed})\ln\left(\frac{\text{Observed}}{\text{Expected}}\right)$$

where Observed is the observed frequency of the amplitudes, Expected is the expected frequency, and the summation is over all possible amplitude combinations. For example, if we have 2 nodes, A and B, capable of having n and m different signal amplitudes respectively, we would look at:

1) the observed number of timepoints, F, where Amplitude(A)=$A_{(1)}$, Amplitude(A)=$A_{(2)}$, ... Amplitude(A)=$A_{(n)}$, and where Amplitude(B)=$B_{(1)}$, Amplitude(B)=$B_{(2)}$, ... Amplitude(B)=$B_{(m)}$;

2) The expected number of timepoints, E, where Amplitude(A)=$A_{(i)}$ and Amplitude(B)=$B_{(j)}$; $E_{ij} = F_i * F_j$ / (total number of timepoints from A and B)

3) The observed number of timepoints where Amplitude(A)=$A_{(i)}$ and Amplitude(B)=$B_{(j)}$, F_{ij} ;

and then we would calculate the above G^2 value summing over all combinations of i and j (Scheines et al., 1996).

We can now see that the utilization of the time dimension of the data is in the frequency of occurrence of the signal amplitudes we observe, not in terms of a progression of the signal as time advances.

Granger:

The Granger causality models determine the existence of a connection through the same calculation used to determine causality within a connection. Consider two nodes, X_1 and X_2. In this approach one examines the effect on the variance of a linear autoregressive model of our two node system. We could initially model each node to have timeseries:

$$X_1(t) = \sum_{j=1}^{p} A_{11,j} X_1(t-j) + E_1(t)$$

$$X_2(t) = \sum_{j=1}^{p} A_{22,j} X_2(t-j) + E_2(t)$$

.

We would then like to see if the variances, E_1 and E_2 decrease with the inclusion of information from earlier time points of the other node:

$$X_1(t) = \sum_{j=1}^{p} A_{11,j} X_1(t-j) + \sum_{j=1}^{p} A_{12,j} X_2(t-j) + E_1(t)$$

$$X_2(t) = \sum_{j=1}^{p} A_{21,j} X_1(t-j) + \sum_{j=1}^{p} A_{22,j} X_2(t-j) + E_2(t)$$

.

If the variance is indeed reduced by the inclusion of information from the other node, then these two nodes will have coefficients A_{12} or A_{21} significantly different from zero. If there is no significant change in the variance, the coefficients A_{12} or A_{21} will not be significantly different than zero, and no connection will be predicted between the nodes (Seth 2007). Thus, Granger causality utilizes the time dimension in the data to observe a direct progression over time in hopes of observing a causal relationship between a set of nodes.

Overall, although the two models, namely Bayesian networks and Granger causality make use of the same data, their approach to how to determine functional connectivity is different. This explains the contrast in performance we observed between these two methods when applied to our data from the visual cortex.

Group-wise Bayesian network approach and group-based analysis in general

Thus far, the thesis has addressed only how the methods apply to data on a single subject level, however group-based analyses are also an important generalization to consider. Popular group-based models which have been used to study resting sate functional connectivity include ICA and PCA, as seen in Yang et al. (2012) where ICA was applied to a group

of subjects in order to identify differences and classify the subjects into sub-groupings, and in Beu et al. (2009) where group-PCA was used to identify patterns of functional connectivity across the whole brain. Many of the methods we've tested can be extended for utilization in group-wise studies. One such model is BayesNet model GES, and so a natural curiosity is to see how well the model can perform with a group's data as opposed to a single subject's data. We have completed analyses with the group-wise version of GES, *Independent Multiple-sample Greedy Equivalence Search (*IMaGES) (Ramsey et al, 2010), using all subject data in the 9 node network (original ROIs), with the 8 runs concatenated and without the removal of the global signal. This particular test was chosen because it allowed a fair bit of room for improvement, whereas in the main test GES had already achieved a relatively high c-sensitivity (66%).The results are very promising, achieving a c-sensitivity much greater than all individual subject analyses (see figure 6 below). Further tests of this nature using IMaGES would be useful in verifying this apparent promise as with only one set of subjects to calculate a c-sensitivity for, we cannot compare the results across different datasets. Based on the results of this test and those of the single subject implementation, GES, this method shows strong potential for meaningful use in group-wise studies. Other

studies where it has been tested and showed promise include Ramsey et al. (2010) and Sun et al. (2012). Further work looking into the performances of network models in group-wise analyses would help to profile the potential of these models as well as the interpretability of resting state functional connectivity.

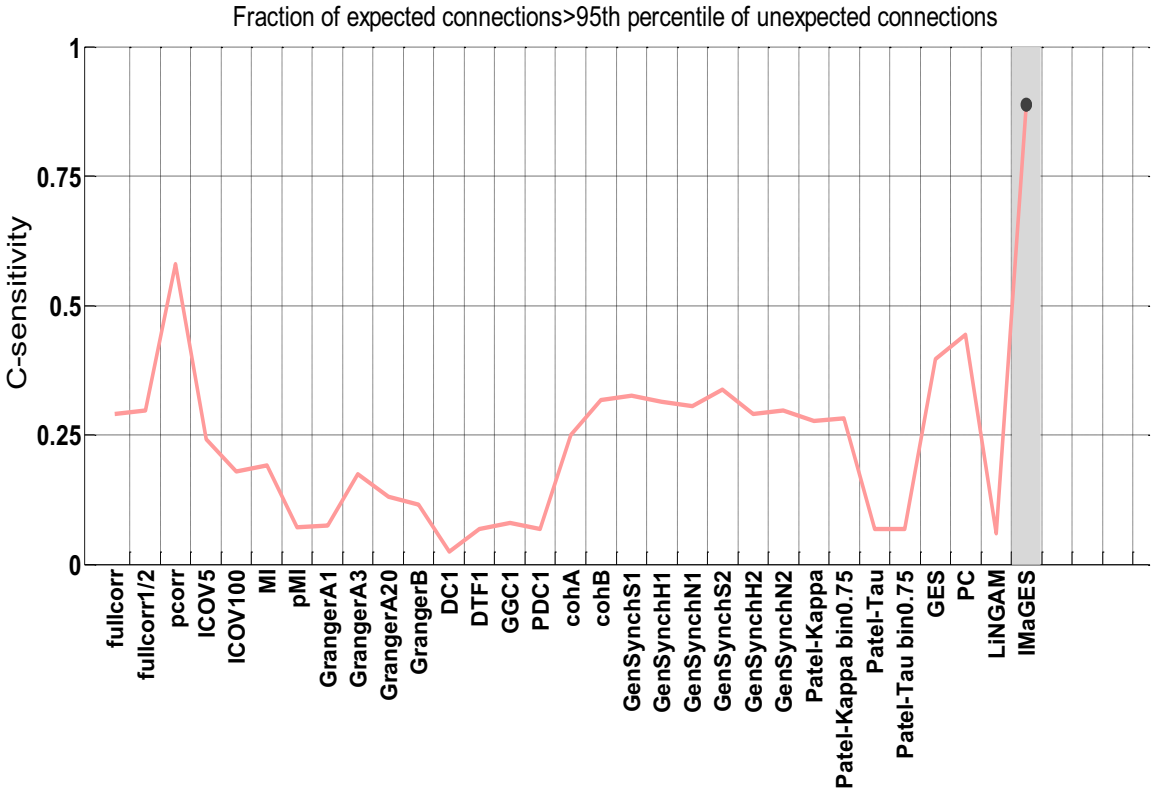

Figure 6. C-sensitivities of the models performed in a subject-wise manner juxtaposed next to the group-wise BayesNet model, IMaGES. The pink line plots the average c-sensitivities across subjects for each model given a 9 node network and 68 minutes of resting fMRI data without the removal of the global signal. The black dot plots the c-sensitivity of the IMaGES model using the data from all subjects.

Significance of the work

An important aspect of our work here is our proposition to include a calibration stage in network studies. This procedure allows for an informed data-based decision on statistical significance. The standards for calibration can be defined on a group level, subject level, or even a scan level, depending on the needs of the study. As the proposed calibration works on the strength of connectivity, it is transferable to studies done using data of many kinds (fMRI, EEG, MEG, PET, eCOG, etc) once it has been interpreted in terms of functional connectivity. We believe that the implementation of calibration to future studies will help us to achieve deeper insights into the brain's networks, contributing to the progress of theories in cognitive neuroscience, and widening the scope of possibilities of interpreting imaging data.

The calibration of connectivity measures in a network context as we have done here can also be a tool for separating indirect from direct connections. In nearly all tests done, fullcorr was able to perform on par with pcorr in the task of identifying a network of direct connections while avoiding false positive detection of a set of indirect connections. Traditionally fullcorr is unable to make this distinction. For example, if we consider two nodes which are functionally connected to a third node but

not functionally connected to each other without the influence of that third node, fullcorr would generally attribute a positive correlation between the two indirectly connected nodes (and pcorr would attempt to identify the influence of the third node by regression analysis). This being said, by measuring the correlation magnitudes of known indirect connections, we are able to set a meaningful threshold in order to calibrate the correlation predictions, and remove from our set of significant correlations many of the indirect connections, such as that mentioned above. The importance of the separation of direct from indirect connections lies in its role as a progression in top-down models of the brain. The more specific we can get about the components of the functional networks we observe, the better our models can become and the closer we come to understanding the brain in parts and in whole.

Resting state data acquisition is particularly interesting and important for implementation in a clinical setting for the simplicity of the instructions to patients and since resting state activations can be obtained even when the subject is under sedation. The better our techniques for interpreting resting state functional connectivity, the closer we get to establishing standard connectivity profiles that we can expect in healthy subjects, as well as in patients with any number of neurological

pathologies. Work has already begun of this nature and there are promising signs that resting state contrasts between healthy subjects and patients will allow for help in diagnoses (Zhang et al., 2012), localization of affected areas (Pannekoek et al., 2012), and monitoring the progress of therapies (Kell, 2012).

Continued work with models such as partial correlation and the Bayesian network models, with the integration of the calibration we've proposed, should positively contribute to future studies leading toward regular clinical utilization of resting state functional connectivity. Some intuitive future directions from this study include conducting similar tests and calibrations of models with different data types, such as magnetoencephalography (MEG) data, and in different brain networks where a set of known standard connections can be defined.

7- Summary and Conclusion

Network modelling methods can be helpful in investigations of brain networks. Functional connectivity in the resting state is reflective of the underlying anatomical connections and thus one way to evaluate network models is to compare their predicted networks using functional Magnetic Resonance Imaging data to a known anatomical network. In our studies of the human visual cortex, we compared the predicted fMRI-based networks to specific connections in monkey anatomy which are expected to exist in humans. We also considered a set of connections that do not exist in monkey anatomy, and expected that they should also not exist in the human network.

None of the models tested were able to predict the expected network while not predicting connections in our unexpected network. The presence of false positives is a major shortcoming of network modeling approaches as they are traditionally used. We propose to add a data-based calibration stage to connectivity analyses in order to minimize false positives while maintaining as many true positives as possible. The calibration should look at known expected and unexpected connections in a network related to the network of interest. Then, the calibration results

can be applied to the analysis of the unknown network of interest. With each of the known networks, one should conduct connectivity analyses with a model of choice. The interaction magnitude distribution of the unexpected network connections together with those of the expected network connections should be used to establish a threshold which will minimize the inclusion of false positive connections while minimizing losses of true positives. Once an appropriate threshold is set, this should be used to calibrate the interaction measures in the unknown network, thereby setting a standard for the connection magnitudes in an informed manner.

The best models for use on 3T resting state fMRI data with a TR of approximately 2s are partial correlation, Regularized Inverse Covariance, or Bayesian Network models PC, CCD, CPC, or FCI (which are equivalent in their determination of the connections existing in a network). Correlation and the General Synchronization measures can also be successful with the inclusion of the calibration stage in the analysis.

Longer fMRI sessions allow the models to converge closer to the true network; the shortest session which can provide reasonable results is 15 to 20 minutes in length in an experiment with a similar set of parameters to ours.

Network size (i.e. total number of nodes considered at one time) can impact the level of success of multivariate models such as partial correlation and the Bayesian Network models. The smallest contained set of nodes relevant to the network being tested will provide the best results as additional information from nodes unimportant to the network will have an effect similar to adding noise.

There has been some controversy as to whether or not a global signal should be removed from data timeseries in the analysis of functional brain activity, as it could perhaps contain important information relevant to the activity. We find that removing the global signal (approximated as the first principal component) leads to much better network predictions and so would recommend removing the global signal in resting state network analyses.

The correct separation of regions of interest such that information from areas which process information separately does not get incorporated into the region, is very important in functional connectivity analyses. Reliable detection of network edges relies heavily on these ROI definitions. Also, the number of voxels contained within the ROIs can strongly impact the sensitivity of a network analysis, such that lower voxel counts result in unreliable networks.

With the implementation of a data-based calibration stage to future network studies using functional brain activation timeseries, we hope to see an increase in network precision, and eventually to see our knowledge of brain networks on any spatial or temporal scale be made more refined and comprehensive. It would be of interest to test in a future study the generalization of our findings with respect to functional domains other than the visual system, such as the relatively well known domain of the motor system.

8- References

Aguirre, G., Zarahn E, D'Esposito M (1998). "The inferential impact of global signal covariates in functional neuroimaging analyses." Neuroimage **8**: 302-306.

Ames, A. I. (2000). "CNS energy metabolism as related to function." Brain Res. Rev **34**: 42-68.

Amir, Y., Harel M, Malach R (1993). "Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual cortex." J Comp Neurol. **334**(1): 19-49.

Angelucci, A., Levitt JB, Walton EJ, Hupe JM, Bullier J, Lund JS (2002). " Circuits for local and global signal integration in primary visual cortex." J Neurosci **22**: 8633-8646.

Angelucci, A., Schiessl I, Nowak L, McLoughlin N (2003). "Functional specificity of feedforward and feedback connections between primate V1 and V2." J Neurosci **22**: 8633-8646.

Arfanakis, K. e. a. (2000). "Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets." Magn. Reson. Imaging **18**(921-930).

Attwell, D. L., S. B (2001). "An energy budget for signalling in the grey matter of the brain." J. Cereb. Blood Flow Metab **21**: 1133-1145.

Baccalá, L. A., K. Sameshima (2001). "Partial directed coherence: a new concept in neural structure determination." Biological Cybernetics **84**(6): 463-474.

Bandettini, P., Wong EC, Hinks RS, Tikofsky RS, Hyde JS. (1992). "Time course EPI of human brain function during task activation." Magn Reson Med. **25**(2): 390-397.

Bandettini, P. A., Eric C. Wong, R. Scott Hinks, Ronald S. Tikofsky, James S. Hyde (1992). "Time course EPI of human brain function during task activation." Magnetic Resonance in Medicine **25**(2): 390-397.

Banerjee, O., El Ghaoui, L., d'Aspremont, A., Natsoulis, G. (2006). "Convex optimization techniques for fitting sparse Gaussian graphical models." Proceedings of the 23rd International Conference on Machine Learning **ACM**: 96.

Bartels, A. Z., S. (2005). "Brain dynamics during natural viewing conditions. A new guide for mapping connectivity in vivo." Neuroimage **24**: 339-349.

Bellec, P., Rosa-Neto P, Lyttelton OC, Benali H, Evans AC. (2010). "Multi-level bootstrap analysis of stable clusters in resting-state fMRI." Neuroimage **51**(3): 1123-1139.

Bernasconi, C., König P. (1999). "On the directionality of cortical interactions studied by structural analysis of electrophysiological recordings." Biol Cybern. **81**(3): 199-210.

Beu, M., Baudrexel S, Hautzel H, Antke C, Mueller HW. (2009). "Neural traffic as voxel-based measure of cerebral functional connectivity in fMRI." J Neurosci Methods. **176**(2): 263-269.

Bewick, V., Liz Cheek, Jonathan Ball (2005). "Statistics review 14: Logistic regression." Crit. Care **9**(1): 112-118.

Bianciardi, M., Masaki Fukunaga, Peter van Gelderen, Silvina G. Horovitz, Jacco A. de Zwart, Karin Shmueli, Jeff H. Duyn (2009). "Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study." Magnetic Resonance Imaging **27**(8): 1019-1029.

Birn, R. M., Diamond, J. B., Smith, M. A., Bandettini, P. A (2006). "Separating respiratory-variationrelated fluctuations from neuronal-activity-related fluctuations in fMRI." Neuroimage **31**(1536-1548).

Biswal, B., F. Zerrin Yetkin, Victor M. Haughton, James S. Hyde (1995). "Functional connectivity in the motor cortex of resting human brain using echo-planar mri." Magnetic Resonance in Medicine **34**(4): 537-541.

Buckner, R. L., J. Sepulcre, T. Talukdar, F.M. Krienen, H. Liu, T. Hedden, J.R. Andrews-Hanna, R.A. Sperling, K.A. Johnson (2009). "Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer's disease." J. neurosci **29**: 18060-11873.

Bullmore, E., Rabe-Hesketh S, Morris RG, Williams SC, Gregory L, Gray JA, Brammer MJ (1996). "Functional magnetic resonance image analysis of a large-scale neurocognitive network." Neuroimage **4**(1): 16-33.

Calabro, F., Vaina LM. (2012). "Interaction of cortical networks mediating object motion detection by moving observers." Exp Brain Res. **221**(2): 177-189.

Calhoun, V., Adali T, McGinty VB, Pekar JJ, Watson TD, Pearlson GD (2001). "fMRI activation in a visual-perception task: network of areas detected using the general linear model and independent components analysis." Neuroimage **14**(5): 1080-1088.

Carbonell, F., Bellec P, Shmuel A (2011). "Global and System-Specific Resting-State fMRI Fluctuations Are Uncorrelated: Principal Component Analysis Reveals Anti-Correlated Networks." Brain Connectivity **in press**.

Chang, C., Glover GH (2010). "Time-frequency dynamics of resting-state brain connectivity measured with fMRI." Neuroimage **50**: 81-89.

Chen, R., Resnick SM, Davatzikos C, Herskovits EH. (2012). "Dynamic Bayesian network modeling for longitudinal brain morphometry." Neuroimage **59**(3): 2330-2338.

Chickering, D. (2003). "Optimal structure identification with greedy search." J. Mach. Learn. Res **3**: 507-554.

Churchill, N., Oder A, Abdi H, Tam F, Lee W, Thomas C, Ween JE, Graham SJ, Strother SC. (2012). "Optimizing preprocessing and analysis pipelines for single-subject fMRI. I. Standard temporal motion and physiological noise correction methods." Hum Brain Mapp. **33**(3): 609-627.

Cole, M., Pathak S, Schneider W. (2010). "Identifying the brain's most globally connected regions." Neuroimage **49**(4): 3132-3148.

Cordes, D., Haughton VM, Arfanakis K, Wendt GJ, Turski PA, Moritz CH, Quigley MA, Meyerand ME. (2000). "Mapping functionally related regions of brain with functional connectivity MR imaging." AJNR Am J Neuroradiol. **21**(9): 1636-1644.

Dale, A., Fischl B, Sereno MI (1999). "Cortical surface-based analysis. I. Segmentation and surface reconstruction." Neuroimage **9**(2): 179-794.

Damien A. Fair, B. L. S., Alexander L. Cohen, Francis M. Miezin, Nico U.F. Dosenbach, Kristin K. Wenger, Michael D. Fox, Abraham Z. Snyder, Marcus E. Raichle, Steven E. Petersen (2007). "A method for using blocked and event-related fMRI data to study "restingstate" functional connectivity." Neuroimage **35**: 396-405.

David, O., Guillemain I, Saillet S, Reyt S, Deransart C, et al (2008). " Identifying neural drivers with functional MRI: An electrophysiological validation." PLoS Biol **6**(12): e315.

De Luca, M., Beckmann, C. F., De Stefano, N., Matthews, P. M., Smith, S. M (2006). "fMRI resting state networks define distinct modes of long-distance interactions in the human brain." Neuroimage **29**: 1359-1367.

Desjardins, A., Kiehl KA, Liddle PF (2001). "Removal of confounding effects of global signal in functional MRI analyses." Neuroimage **13**: 751-758.

Dhamala, M., G. Rangarajan, M. Ding (2008). "Analyzing information flow in brain networks with nonparametric Granger causality." Neuroimage **41**: 354-362.

Ding, M., Chen, Y., & Bressler, S.L. (2006). "Granger causality: Basic theory and application to neuroscience." In: Schelter. S., Winterhalder, N., & Timmer, J. Handbook of Time Series Analysis.

Dumoulin, S. O., Brian A. Wandell (2008). "Population receptive field estimates in human visual cortex." NeuroImage **39**: 647-660.

Duncan, R., Boynton GM. (2003). "Cortical magnification within human primary visual cortex correlates with acuity thresholds." Neuron **38**(4): 659-671.

Engel, S., Glover GH, Wandell BA (1997). "Retinotopic organization in human visual cortex and the spatial precision of functional MRI." Cereb Cortex **7**(2): 181-192.

Engel, S., Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN. (1994). "fMRI of human visual cortex." Nature **369**(6481).

Engel, S. A., G H Glover and B A Wandell (1997). "Retinotopic organization in human visual cortex and the spatial precision of functional MRI." Cereb Cortex. **7**(2): 181-192.

Esposito, F., Aragri A, Pesaresi I, Cirillo S, Tedeschi G, Marciano E, Goebel R, Di Salle F. (2008). "Independent component model of the default-mode brain function: combining individual-level and population-level analyses in resting-state fMRI." Magn Reson Imaging. **26**(7): 905-913.

Fischl, B., Sereno MI, Dale AM (1999). "Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system." Neuroimage **9**(2): 195-207.

Fornito, A., Harrison BJ, Zalesky A, Simons JS. (2012). "Competitive and cooperative dynamics of large-scale brain functional networks supporting recollection." Proc Natl Acad Sci U S A. **109**(31): 12788-12793.

Fox, A., Symons SP, Aviv RI, Howard P, Yeung R, Bartlett ES. (2009). "Should modeling methodology suppress anatomic excellence?" Stroke **40**(11): 3511-3517.

Fox, M., Zhang D, Snyder AZ, Raichle ME. (2009). "The global signal and observed anticorrelated resting state brain networks." J Neurophysiol. **101**(6): 3270-3283.

Fox, M. D., Corbetta, M., Snyder, A.Z., Vincent, J.L., & Raichle, M.E. (2006). "Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems." Proceedings of the National Academy of Sciences **103**: 10046-10051.

Fox, M. D., Michael Greicius (2010). "Clinical applications of resting state functional connectivity." Front. Syst. Neurosci **4**(19).

Fox, M. D. e. a. (2005). "The human brain is intrinsically organized into dynamic, anti-correlated functional networks." Proc. Natl Acad. Sci. USA **102**: 9673-9678.

Freiwald, W. A., P. Valdes, J. Bosch, R. Biscay, J.C. Jimenez, L.M. Rodriguez, V. Rodriguez, A.K. Kreiter, W. Singer (1999). "Testing non-linearity and directedness of interactions between neural groups in the macaque inferotemporal cortex." J. Neurosci. Methods **94**: 105-119.

Frenzel, S., Bernd Pompe (2007). "Partial Mutual Information for Coupling Analysis of Multivariate Time Series." PHYSICAL REVIEW LETTERS **99**(204101).

Friedman, J., Hastie, T., Tibshirani, R (2008). "Sparse inverse covariance estimation with the Graphical Lasso." Biostat **9**(3): 432-441.

Friston, K. J. (2011). "Functional and Effective Connectivity: A Review." BRAIN CONNECTIVITY **1**(1): 13-36.

Gates, K. M., P.C.M. Molenaar, F.G. Hillary, N. Ram, M.J. Rovine (2010). "Automatic search for fMRI connectivity mapping: an alternative to Granger causality testing using formal equivalences among SEM path modeling, VAR, and unified SEM." NeuroImage **50**: 1118-1125.

Gattass, R., Sheila Nascimento-Silva, Juliana G.M Soares, Bruss Lima, Ana Karla Jansen, Antonia Cinira M Diogo, Mariana F Farias, Marco Marcondes, Eliã P Botelho, Otávio S Mariani, João Azzi, and Mario Fiorani (2005). "Cortical visual areas in monkeys: location, topography, connections, columns, plasticity and cortical dynamics." Philos Trans R Soc Lond B Biol Sci. **360**(1456): 709-731.

Gavrilescu, M., Shaw ME, Stuart GW, Eckersley P, Svalbe ID, Egan GF. (2002). "Simulation of the effects of global normalization procedures in functional MRI." Neuroimage **17**: 532-542.

Geweke, J. (1982). "Measures of Conditional Linear Dependence and Feedback Between Time Series." J Am Stat Assoc. **77**: 304-313.

Gilbert, C., Wiesel TN (1979). "Morphology and intracortical projections of functionally characterized neurons in the cat visual cortex." Nature **280**: 120-125.

Girard, P., Salin PA, Bullier J (1991). "Visual activity in areas V3a and V3 during reversible inactivation of area V1 in the macaque monkey." J Neurophysiol. **66**(5): 1493-1503.

Glover, G. H., Li, T. Q. & Ress, D. (2000). "Image-based method for retrospective correction of physiological motion artifacts in fMRI: RETROICOR." Magn. Res. Med **44**: 162-167.

Golland, Y., Shlomo Bentin, Hagar Gelbard, Yoav Benjamini, Ruth Heller, Yuval Nir, Uri Hasson, Rafael Malach (2007). "Extrinsic and Intrinsic Systems in the Posterior Cortex of the Human Brain Revealed during Natural Sensory Stimulation." Cereb Cortex. **17**(4): 766-777.

Granger, C. W. J. (1969). "Investigating causal relations by econometric models and cross-spectral methods." Econometrica **37**(3): 424-438.

Grinsted, A., J. C. Moore, S. Jevrejeva (2004). "Application of the cross wavelet transform and wavelet coherence to geophysical time series." Nonlinear Processes in Geophysics **11**: 561-566.

Hampson, M., Olson, I. R., Leung, H.-C., Skudlarski, P. & Gore, J. C. (2004). "Changes in functional connectivity of human MT/V5 with visual motion input." Neuroreport. **15**: 1315-1319.

Harrell, F. E. (2001). "Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis." Springer Series in Statistics.

Havlicek, M., J. Jan, M. Brazdil, V.D. Calhoun (2010). "Dynamic Granger causality based on Kalman filter for evaluation of functional network connectivity in fMRI data." NeuroImage **53**: 1477-1491.

He, H., Sui J, Yu Q, Turner JA, Ho BC, Sponheim SR, Manoach DS, Clark VP, Calhoun VD. (2012). "Altered Small-World Brain Networks in Schizophrenia Patients during Working Memory Performance." PLoS One. **7**(6): e38195.

Heinzle, J., Thorsten Kahnt, John-Dylan Haynes (2011). "Topographically specific functional connectivity between visual field maps in the human brain." Neuroimage **56**(3): 1426-1436.

Heinzle, J., Wenzel MA, Haynes JD. (2012). "Visuomotor functional network topology predicts upcoming tasks." J Neurosci. **32**(29): 9960-9968.

Hesse, W., E. Moller, M. Arnold, B. Schack (2003). "The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies." J. Neurosci. Methods **124**: 65-77.

Hinds, O., Jonathan R. Polimeni, Niranjini Rajendran, Mukund Balasubramanian, Lawrence L. Wald, Jean C. Augustinack, Graham Wiggins, H. Diana Rosas, Bruce Fischl, and Eric L. Schwartz (2008). "The Intrinsic Shape of Human and Macaque Primary Visual Cortex." Cereb. Cortex **18**(11): 2586-2595.

Hinke, R., Hu X, Stillman AE, Kim SG, Merkle H, Salmi R, Ugurbil K. (1993). "Functional magnetic resonance imaging of Broca's area during internal speech." Neuroreport. **4**(6): 675-678.

HOCHBERG, Y., YOAV BENJAMINI (1990). "MORE POWERFUL PROCEDURES FOR MULTIPLE SIGNIFICANCE TESTING." STATISTICS IN MEDICINE **9**: 811-818.

Holmes, G. (1945). "The organisation of the visual cortex in man." Ferrier lecture. Proc R Soc Ser B **132**: 348-361.

Honey, C., Thivierge JP, Sporns O (2010). "Can structure predict function in the human brain?" Neuroimage **52**(3): 766-776.

Ing, A., Schwarzbauer C. (2012). "A dual echo approach to motion correction for functional connectivity studies." Neuroimage **epub**.

Inouye, T. (1900, 2000). "Visual disturbances following gunshot wounds of the cortical visual area based on observations of the wounded in the recent Japanese wars." Brain **-, 1**(-, 101): 1904-1905, 1123

Jiang, T., He Y, Zang Y, Weng X (2004). "Modulation of functional connectivity during the resting state and the motor task." Hum Brain Mapp. **22**(1): 63-71.

Jiang, T., He, Y., Zang, Y. & Weng, X (2004). "Modulation of functional connectivity during the resting state and the motor task." Hum. Brain Mapp. **22**: 63-71.

Kaminski, M., M. Ding, W. A. Truccolo, S. L. Bressler (2001). "Evaluating the causal relations in neural systems; Granger causality, directed transfer function and statistical assessment of significance." Biological Cybernetics **85**: 145-157.

Kell, C. (2012). "Resting-state MRI: A peek through the keyhole on therapy for stuttering." Neurology **epub**.

Koch, M. A., Norris, D. G., Hund-Georgiadis, M. (2002). "An investigation of functional and anatomical connectivity using magnetic resonance imaging." Neuroimage **16**(241-250).

Kwong, K., Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, et al (1992). "Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation.
." Proc Natl Acad Sci U S A. **89**(12): 5675-5679.

Kwong, K., Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, et al. (1992). "Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation." Proc Natl Acad Sci U S A. **89**(12): 5675-5679.

Larsson, J., Heeger DJ (2006). "Two retinotopic visual areas in human lateral occipital cortex." J Neurosci. **26**(51): 13128-13142.

Laurienti, P. (2004). "Deactivations, global signal, and the default mode of brain function." J Cogn Neurosci **16**: 1481-1493.

Lennie, P. (2003). "The cost of cortical computation." Curr. Biol. **13**: 493-497.

Li, R., Chen K, Fleisher AS, Reiman EM, Yao L, Wu X. (2011). "Large-scale directional connections among multi resting-state neural networks in human brain: a functional MRI and Bayesian network modeling study." Neuroimage **56**(3): 1035-1042.

Li, Z., Kadivar A, Pluta J, Dunlop J, Wang Z. (2012). "Test-retest stability analysis of resting brain activity revealed by blood oxygen level-dependent functional MRI." J Magn Reson Imaging. **36**(2): 344-354.

Liu, P., Zhang Y, Zhou G, Yuan K, Qin W, Zhuo L, Liang J, Chen P, Dai J, Liu Y, Tian J. (2009). "Partial correlation investigation on the default mode network involved in acupuncture: an fMRI study." Neurosci Lett. **462**(3): 183-187.

Long, X., Zuo XN, Kiviniemi V, Yang Y, Zou QH, Zhu CZ, Jiang TZ, Yang H, Gong QY, Wang L, Li KC, Xie S, Zang YF. (2008). "Default mode network as revealed with multiple methods for resting-state functional MRI analysis." J Neurosci Methods. **171**(2): 349-355.

Lowe, M. J., Dzemidzic, M., Lurito, J. T., Mathews, V. P. & Phillips, M. D (2000). "Correlations in low-frequency BOLD fluctuations reflect cortico-cortical connections." Neuroimage **12**(582-587).

Lund, T. E., Madsen, K. H., Sidaros, K., Luo, W., Nichols, T. E (2006). "Non-white noise in fMRI: does modelling have an impact?" Neuroimage **29**: 56-66.

Lutz, C. (2009). "Granger Causality Test." MATLAB Central.

Lyon, D. C., Neeraj Jain, Jon H. Kaas (1998). "Cortical connections of striate and extrastriate visual areas in tree shrews." THE JOURNAL OF COMPARATIVE NEUROLOGY **401**(1): 109-128.

Macey, P. M., Katherine E Macey, Rajesh Kumar, Ronald M Harper (2004). "A method for removal of global effects from fMRI time series." Neuroimage **22**(1): 360-366.

Macey, P. M., Macey, K. E., Kumar, R. & Harper, R. M. (2004). "A method for the removal of global effects from fMRI time series." Neuroimage **22**: 360-366.

Meek, C. (1995). "Causal inference and causal explanation with background knowledge." Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence: 403-410.

Mizuhara, H., Wang LQ, Kobayashi K, Yamaguchi Y (2004). "A long-range cortical network emerging with theta oscillation in a mental task." Neuroreport **15**(8): 1233-1238.

Moore, A. W. (2001). "Bayes Nets for representing and reasoning about uncertainty."

Morgan, V. L. P., R. R. (2004). "The effect of sensorimotor activation on functional connectivity mapping with fMRI." Magn. Reson. Med **22**: 1069-1075.

Ng, B., Abugharbieh R, Varoquaux G, Poline JB, Thirion B. (2011). "Connectivity-informed fMRI activation detection." Med Image Comput Comput Assist Interv. **14**((Pt 2)): 285-292.

O'Donnell, M., Edelstein WA. (1985). "NMR imaging in the presence of magnetic field inhomogeneities and gradient field nonlinearities." Med Phys. **12**(1): 20-26.

Ogawa, S., Lee TM, Kay AR, Tank DW. (1990). "Brain magnetic resonance imaging with contrast dependent on blood oxygenation." Proc Natl Acad Sci U S A. **87**(24): 9868-9872.

Ogawa, S., Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, Ugurbil K. (1992). "Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging.
." Proc Natl Acad Sci U S A. **89**(13): 5951-5955.

Pamilo, S., Malinen S, Hlushchuk Y, Seppä M, Tikka P, Hari R. (2012). "Functional Subdivision of Group-ICA Results of fMRI Data Collected during Cinema Viewing." PLoS One. **7**(7): e42000.

Pannekoek, J., Veer IM, van Tol MJ, van der Werff SJ, Demenescu LR, Aleman A, Veltman DJ, Zitman FG, Rombouts SA, van der Wee NJ. (2012). "Aberrant limbic and salience network resting-state functional connectivity in panic disorder without comorbidity." J Affect Disord. **epub**.

Patel, R., Bowman, F., Rilling, J (2006). "A Bayesian approach to determining connectivity of the human brain." Hum. Brain Mapp **27**: 267-276.

Perlbarg V, B. P., Anton JL, Pélégrini-Issac M, Doyon J, Benali H (2007). "CORSICA: correction of structured noise in fMRI by automatic identification of ICA components." Magn Reson Imaging **25**(1): 35-46.

Quigley, M. A. e. a. (2003). "Role of the corpus callosum in functional connectivity." Am. J. Neuroradiol **24**(208-212).

Quiroga, R. Q., A. Kraskov, T. Kreuz, P. Grassberger (2002). "Performance of different synchronization measures in real data: A case study
on electroencephalographic signals." PHYSICAL REVIEW E **65**(041903).

Raichle, F. M. D. a. M. E. (2007). "Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging." Nature **8**: 700-711.

Raichle, M. E. M., M. A. (2006). "Brain work and brain imaging." Annu. Rev. Neurosci. **29**: 449-476.

Ramsey, J., Hanson, S., Hanson, C., Halchenko, Y., Poldrack, R., Glymour, C. (2010). "Six problems for causal inference from fMRI." Neuroimage **49**(2): 1545-1558.

Ramsey, J., Zhang, J., Spirtes, P. (2006). "Adjacency-faithfulness and conservative causal inference." Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence.

Ramseya, J. D., S.J. Hanson, C. Hanson, Y.O. Halchenko, R.A. Poldrack, C. Glymour (2010). "Six problems for causal inference from fMRI." Neuroimage **49**(2): 1545-1558.

Reiss, P., Mennes M, Petkova E, Huang L, Hoptman MJ, Biswal BB, Colcombe SJ, Zuo XN, Milham MP. (2011). "Extracting information from functional connectivity maps via function-on-scalar regression." Neuroimage **56**(1): 140-148.

Richardson, T., Spirtes, P. (2001). "Automated discovery of linear feedback models." In: Glymour, C., Cooper, G. (Eds.), Computation, Causation, and Causality.

Rockland, K., Lund JS (1982). "Widespread periodic intrinsic connections in the tree shrew visual cortex." Science **215**: 1532-1534.

Rockland, K. S., J.S. Lund, and A.L. Humphrey (1982). "Anatomical banding of intrinsic connections in striate cortex of tree shrews (Tupaia glis)." Comp. Neurol **209**: 41-58.

Rodgers, J. L., W. Alan Nicewander (1988). "Thirteen Ways to Look at the Correlation Coefficient." The American Statistician **42**(1).

Roebroeck, A., Elia Formisano, Rainer Goebel (2009). "The identification of interacting networks in the brain using fMRI: Model selection, causality and deconvolution." Neuroimage.

Rosa, M. G., Paul R Manger (2005). "Clarifying homologies in the mammalian cerebral cortex: the case of the third visual area (v3)." Clinical and Experimental Pharmacology and Physiology **32**(5-6): 327-339.

Saad, Z. S., Ropella, K.M., DeYoe, E.A., Bandettini, P.A. (2003). "The spatial extent of the BOLD response." Neuroimage **19**(1): 132-144.

Salin, P., Bullier J (1995). " Corticocortical connections in the visual system: structure and function." Physiol Rev **75**: 107-154.

Salvador, R., Suckling J, Schwarzbauer C, Bullmore E. (2005). "Undirected graphs of frequency-dependent functional connectivity in whole brain networks." Philosophical Transactions of the Royal Society B **360**(1457): 937-946.

Scheines, R., Peter Spirtes, Clark Glymour, Christopher Meek, Thomas Richardson (1996). "TETRAD 3: Tools for Causal Modeling, User's Manual." from http://www.phil.cmu.edu/projects/tetrad/old/tet3/master.htm.

Scheinost, D., Benjamin J, Lacadie CM, Vohr B, Schneider KC, Ment LR, Papademetris X, Constable RT. (2012). "The intrinsic connectivity distribution: A novel contrast measure reflecting voxel level functional connectivity." Neuroimage **62**(3): 1510-1519.

Schmolesky, M. T., Youngchang Wang, Doug P. Hanes, Kirk G. Thompson,Stefan Leutgeb, Jeffrey D. Schall, Audie G. Leventhal (1998). "Signal Timing Across the Macaque Visual System." J Neurophysiol **79**: 3272-3278.

Schölvinck, M. L., Alexander Maier, Frank Q. Ye, Jeff H. Duyn, and David A. Leopold (2010). "Neural basis of global resting-state fMRI activity." PNAS **107**(22): 10238-10243.

Schwarz, A., McGonigle J. (2011). "Negative edges and soft thresholding in complex network analysis of resting state functional connectivity data." Neuroimage **55**(3): 1132-1146.

Sereno, M., Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB. (1995). "Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging." Science **269**(5212): 889-893.

Sesma, M. A., V.A. Casagrande, and J.H. Kaas (1984). "Cortical connections of area 17 in tree shrews." J. Comp. Neurol **230**(337-351).

Seth, A. (2007). "Granger causality." Scholarpedia **2**(7): 1667.

Seth, A. (2010). "A MATLAB toolbox for Granger causal connectivity analysis." Journal of Neuroscience Methods **186**: 262-273.

Shannon, C. (1948). "A mathematical theory of communication." Bell Syst. Tech. J. **27**: 379-423.

Shimizu, S., P.O. Hoyer, A. Hyvärinen, and A.J. Kerminen (2006). "A linear, non-gaussian acyclic model for causal discovery." Journal of Machine Learning Research **7**: 2003-2030.

Shmuel, A., David A. Leopold (2008). "Neuronal Correlates of Spontaneous Fluctuations in fMRI Signals in Monkey Visual Cortex: Implications for Functional Connectivity at Rest." Human Brain Mapping **29**: 751-761.

Shmuel, A., Maria Korman, Anna Sterkin, Michal Harel, Shimon Ullman, Rafael Malach, Amiram Grinvald (2005). "Retinotopic Axis Specificity and Selective Clustering of Feedback Projections from V2 to V1 in the Owl Monkey." J Neurosci **25**(8): 2117-2131.

Shulman, R. G., Rothman, D. L., Behar, K. L. & Hyder, F. (2004). "Energetic basis of brain activity: implications for neuroimaging." Trends Neurosci **27**: 489-495.

Smith, J. F., Ajay Pillai, Kewei Chen, Barry Horwitz (2011). "Effective Connectivity Modeling for fMRI: Six Issues and Possible Solutions Using Linear Dynamic Systems." Front Syst Neurosci **5**(104).

Smith, S. M., K.L. Miller, S. Moeller, J. Xu, E.J. Auerbach, M.W. Woolrich, C.F. Beckmann, M. Jenkinson, J. Andersson, M.F. Glasser, D. Van Essen, D. Feinberg, E. Yacoub, and K. Ugurbil. (2012). " Temporally-independent functional modes of spontaneous brain activity." PNAS **109**(8): 3131-3136.

Smith, S. M., Karla L. Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F. Beckmann, Thomas E. Nichols, Joseph D. Ramsey, Mark W. Woolrich (2011). "Network modelling methods for FMRI." Neuroimage **54**: 875-891.

Sotero, R., Shmuel A (2012). "Energy-based stochastic control of neural mass models suggests time-varying effective connectivity in the resting state." Journal of Computational Neuroscience **in press**.

Spirtes, P., C. Glymour, and R. Scheines (1991). "An Algorithm for Fast Recovery of Sparse Causal Graphs." Social Science Computer Review **9**: 62-72.

Spirtes, P., C. Glymour, R. Scheines (1993). "Causation, Prediction, and Search." Springer-Verlag Lecture Notes in Statistics **81**.

Stephana, K. E., Alard Roebroeckc (2012). "A short history of causal modeling of fMRI data." Neuroimage **62**(2): 856-863.

Sun, F., Miller LM, D'Esposito M. (2004). "Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data." Neuroimage. **21**(2): 647-658.

Sun, F. T., Miller, L. M., Rao, A. A. & D'Esposito, M. (2006). "Functional connectivity of cortical networks involved in bimanual motor sequence learning." Cereb. Cortex **17**: 1227-1234.

Sun, J., Hu X, Huang X, Liu Y, Li K, Li X, Han J, Guo L, Liu T, Zhang J. (2012). "Inferring consistent functional interaction patterns from natural stimulus FMRI data." Neuroimage **61**(4): 987-999.

Sun, J., Hu X, Huang X, Liu Y, Li K, Li X, Han J, Guo L, Liu T, Zhang J. (2012). "Inferring consistent functional interaction patterns from natural stimulus FMRI data." Neuroimage. **61**(4): 987-999.

Tanabe, J., Miller D, Tregellas J, Freedman R, Meyer FG. (2002). "Comparison of detrending methods for optimal fMRI preprocessing." Neuroimage **15**(4): 902-907.

Tohka, J., He Y, Evans AC. (2012). "The impact of sampling density upon cortical network analysis: regions or points." Magn Reson Imaging. **30**(7): 978-992.

Tomasi, D., N.D. Volkow (2010). "Functional connectivity density mapping." Proc. Natl. Acad. Sci. **107**: 9885-9889.

Tootell, R. B. H., Janine D. Mendola, Nouchine K. Hadjikhani, Patrick J. Ledden, Arthur. K. Liu, John B. Reppas, Martin I. Sereno, and Anders M. Dale (1997). "Functional Analysis of V3A and Related Areas in Human Visual Cortex." The Journal of Neuroscience **17**(18): 7060-7078.

Townsend, J., Torrisi SJ, Lieberman MD, Sugar CA, Bookheimer SY, Altshuler LL. (2012). "Frontal-Amygdala Connectivity Alterations During Emotion Downregulation in Bipolar I Disorder." Biol Psychiatry. **Epub**.

Triantafyllou, C., T R.D. Hoge, G. Krueger, C.J. Wiggins, A. Potthast, G.C. Wiggins, L.L. Walda (2005). "Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters." Neuroimage **26**: 243-250.

Turkheimera, F. E., C.B. Smitha, K. Schmidta (2001). "Estimation of the Number of "True" Null Hypotheses in Multivariate Analysis of Neuroimaging Data." Neuroimage **13**(5): 920-930.

Vaessen, M., Braakman HM, Heerink JS, Jansen JF, Debeij-van Hall MH, Hofman PA, Aldenkamp AP, Backes WH. (2012). "Abnormal Modular Organization of Functional Networks in Cognitively Impaired Children with Frontal Lobe Epilepsy." Cereb Cortex. **epub**.

Valdés-Sosa, P. A., Jose M Sánchez-Bornot, Agustín Lage-Castellanos, Mayrim Vega-Hernández, Jorge Bosch-Bayard, Lester Melie-García and Erick Canales-Rodríguez (2005). "Estimating brain functional connectivity with sparse multivariate autoregression." Phil. Trans. R. Soc. B **360**(1457): 969-981.

van den Heuvel, M., Mandl RC, Kahn RS, Hulshoff Pol HE (2009). "Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain." Hum Brain Mapp **30**: 3127-3141.

Wang, J., Aguirre GK, Kimberg DY, Detre JA. (2003). "Empirical analyses of null-hypothesis perfusion FMRI data at 1.5 and 4 T." Neuroimage **19**(4): 1449-1462.

Watkins, K., Cowey A, Alexander I, Filippini N, Kennedy JM, Smith SM, Ragge N, Bridge H. (2012). "Language networks in anophthalmia: maintained hierarchy of processing in 'visual' cortex." brain **135**(5): 1566-1577.

Wise, R. J. S., Ide, K., Poulin, M. J. & Tracey, I. (2004). "Resting state fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal." Neuroimage **21**: 1652-1664.

Worsley, K., Liao CH, Aston J, Petre V, Duncan GH, Morales F, Evans AC (2002). "A general statistical analysis for fMRI data." Neuroimage **15**(1): 1-15.

Xiong, J., Parsons LM, Gao JH, Fox PT. (1999). "Interregional connectivity to primary motor cortex revealed using MRI resting state images." Hum Brain Mapp. **8**(2-3): 151-156.

Xuan, X., Kevin Murphy (2007). "Modeling changing dependency structure in multivariate time series." Proceedings of the 24 th International Conference on Machine Learning.

Yang, J., Shu H. (2012). "The role of the premotor cortex and the primary motor cortex in action verb comprehension: Evidence from Granger causality analysis." Brain Res Bull. **88**(5): 460-466.

Yang, Z., Zuo XN, Wang P, Li Z, Laconte SM, Bandettini PA, Hu XP. (2012). "Generalized RAICAR: Discover homogeneous subject (sub)groups by reproducibility of their intrinsic connectivity networks." Neuroimage **63**(1): 403-414.

Zarahn, E., Aguirre, G. K. & D'Esposito, M. (1997). "Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions." Neuroimage **5**: 179-197.

Zhang, J. (2008). "On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias." Artif. Intell **172**: 1873-1896.

Zhang, J., Cheng W, Wang Z, Zhang Z, Lu W, Lu G, Feng J. (2012). "Pattern classification of large-scale functional brain networks: identification of informative neuroimaging markers for epilepsy." PLoS One. **7**(5): e36733.

Zhou, D., Wesley K. Thompson, Greg Siegle (2009). "MATLAB Toolbox for Functional Connectivity." Neuroimage **47**(4): 1590.