NeuroImage 149 (2017) 256–266

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

A variance components model for statistical inference on functional connectivity networks

Mark Fiecas[a,*], Ivor Cribben[b], Reyhaneh Bahktiari[c], Jacqueline Cummine[c]

[a] Division of Biostatistics, University of Minnesota, Minneapolis, MN 55455, USA
[b] Department of Finance and Statistical Analysis, Alberta School of Business, University of Alberta, Edmonton, AB, Canada T6G 2R6
[c] Department of Communication Sciences and Disorders, Faculty of Rehabilitation Medicine and Neuroscience and Mental Health Institute, University of Alberta, Edmonton, AB, Canada T6G 2G4

ARTICLE INFO

Keywords:
Functional connectivity networks
Temporal autocorrelation
Subject heterogeneity
Resting-state fMRI
Dyslexia

ABSTRACT

We propose a variance components linear modeling framework to conduct statistical inference on functional connectivity networks that directly accounts for the temporal autocorrelation inherent in functional magnetic resonance imaging (fMRI) time series data and for the heterogeneity across subjects in the study. The novel method estimates the autocorrelation structure in a nonparametric and subject-specific manner, and estimates the variance due to the heterogeneity using iterative least squares. We apply the new model to a resting-state fMRI study to compare the functional connectivity networks in both typical and reading impaired young adults in order to characterize the resting state networks that are related to reading processes. We also compare the performance of our model to other methods of statistical inference on functional connectivity networks that do not account for the temporal autocorrelation or heterogeneity across the subjects using simulated data, and show that by accounting for these sources of variation and covariation results in more powerful tests for statistical inference.

1. Introduction

Functional connectivity (FC) is the investigation of the temporal dependence in the fluctuations of the blood oxygenation level-dependent (BOLD) signals across different regions in the brain (Friston et al., 1993; Fiecas et al., 2013; Cribben and Fiecas, 2016). FC is sometimes referred to as the undirected association between two spatially distinct regions of the brain. FC has been studied using noninvasive imaging techniques such as functional magnetic resonance imaging (fMRI), electroencephalography (EEG) and magnetoencephalography (MEG) and using invasive techniques such as electrophysiology (ECoG).

It has been shown that FC can be disrupted for various neurological disorders including Alzheimer's disease (Buckner et al., 2009; Sorg et al., 2007), depression (Lui et al., 2011; Kaiser et al., 2015) and autism (Di Martino et al., 2014), compared to healthy subjects. An objective of FC analyses is to discover specific and stable neuroimaging biomarkers or phenotypes of neurological disorders using FC networks from resting-state data. Existing standard methods for conducting statistical inference on the FC networks between a group of patients and a group of healthy controls in a multi-subject resting-state fMRI data set can be broadly divided into two sets. The methods are based on techniques from complex network analysis that originated from graph theory. In this setup, a brain network is represented by a graph, in which the vertices of the graph correspond to regions-of-interest (ROIs), and edges correspond to the functional associations between the ROIs. The first set of methods directly considers the measured signal between the graph edges, whereas the second set of methods employs graph summary statistics. More specifically, the first set of methods begins by defining either functionally or anatomically the ROIs or the vertices in the network. To estimate the edges, we simply cross-correlate the BOLD signals obtained from these regions, for each subject. After applying the Fisher Z-transform to the sample correlation coefficients, we can use a two-sample t-test on the sample correlation coefficients to test the hypothesis of equal FC between the patients and the controls. The second set of methods also estimates the vertices and edges in the same fashion as the first set of methods (using various methods such as correlation, thresholded correlation, partial correlation or sparse partial correlation) for each subject, and then estimates graph summary statistics such as small–worldness, degree distribution, modularity to name just a few. Then to conduct statistical inference, we can test for differences between the patients and the controls groups on the graph summary statistics (Bullmore and Sporns, 2009; Zalesky et al., 2010, 2012; Ginestet et al., 2013).

* Corresponding author.
E-mail address: mfiecas@umn.edu (M. Fiecas).

http://dx.doi.org/10.1016/j.neuroimage.2017.01.051
Received 22 June 2016; Accepted 21 January 2017
Available online 24 January 2017

Other statistical methods also exist in this space. The exponential random graph modeling (ERGM) framework (Simpson et al., 2012) allows for the statistical comparison of groups of brain networks while accounting for topological differences across networks. However, it is difficult to incorporate topological features inherent in each individual network when comparing groups of networks (Simpson et al., 2011). Simpson et al. (2013) introduced two methods, the first of which compares the consistency of network organization between groups while the second compares the degree distributions (distribution of the number of connections each node has) between groups. In addition, Simpson and Laurienti (2015) introduced a two-part mixed-effects modeling framework that allows for testing for overall group differences in the strength and probability of network connections, group differences in network topology, and individual edge differences while accounting for the complex dependence structures of the networks. More recently, Pan et al. (2014) developed a permutation testing procedure called the aSPU method which was then used by Kim et al. (2015) to test for equality of FC networks. The latest developments in this area include the work by Fujita et al. (2015), who used the concept of correlation between graphs to identify differences in the interactions between controls and patients, and the work by Belilovsky et al. (2015), who introduced a hypothesis test for differences in Gaussian graphical models with an application to brain connectivity.

There are a number of issues with the standard sets of methods. First, the methods implicitly assume that the BOLD signals are temporally uncorrelated, which is not the case (Friston et al., 1995). Hence, it is important for researchers to understand how this assumption affects not only their FC estimates but also the hypothesis tests they carry out on their FC estimates. Second, the standard methods fail to account for the heterogeneity across subjects which also leads to spurious results. Finally, the standard methods either carry out pairwise edge tests between the groups of patients and controls or summarize the networks using graph summary statistics and then tests for equality between the groups. In both cases, the equality of the overall FC network is not being explicitly tested. Of course, one could use a Hotelling T^2 test, which is a multivariate version of the t-test, to test the equality of the entire network, but as we will observe later this test has very low power or is impossible to compute when the number of subjects is small.

In this work, we are particularly interested in comparing the FC networks between a group of patients and a group of healthy controls in a multi-subject resting-state fMRI data set. To this end, we introduce a variance components framework for modeling and conducting statistical inference on FC networks that directly takes into account the autocorrelation inherent in the ROI time series of each subject and subject heterogeneity. We show that, by accounting for these, our statistical tests can show a substantial improvement in performance with respect to statistical power. We focus on multi-subject FC for resting-state fMRI data, however, the statistical challenges we outline and methods we develop are applicable to modelling and performing statistical inference about FC in other neuroimaging modalities such as EEG, MEG and ECoG. We apply our variance components model to a resting-state fMRI study of reading in order to characterize the resting state networks that are related to reading processes. This topic is currently of great interest (Koyama et al., 2010, 2011). There is evidence for an intrinsically connected set of regions that includes the same brain areas that are shown to be active in task-based reading paradigms, including fusiform gyrus, superior temporal gyrus, temporo-parietal junction, precentral gyrus and inferior frontal gyrus (Koyama et al., 2013; Murdaugh et al., 2015). Indeed, researchers have shown that the strength of FC between these regions appears to be sensitive to differences between individuals with and without reading disability (Horowitz-Kraus et al., 2015). In addition, as reading performance increases following treatment, so too does the FC between visual regions and i) attentional networks, ii) executive networks, and iii) language networks in both typical and reading impaired children (Horowitz-Kraus et al., 2015; Murdaugh et al., 2015). Given the many confounding variables (e.g., previous diagnosis of ADHD, comorbid learning disabilities) and uncertainties about proper control groups (e.g., control groups matched by age and reading abilities) that are inherent in task-dependent brain imaging studies, characterization of the task-independent reading network serves as one approach to furthering our understanding of the underlying reading framework in both skilled and impaired readers.

This paper is organized as follows. In Section 2 we introduce the theoretical background of our variance components model and how to estimate the model's parameters. We also detail the resting-state fMRI data set on typical and reading impaired young adults and the simulation study. In Section 3 we show the performance of our model on these data sets. We end with a discussion and conclusions in Sections 4 and 5, respectively.

2. Materials and methods

2.1. A variance components model

Suppose we observe a p-variate time series or similarly a network consisting of p ROIs, across N subjects. Without loss of generality, suppose that the marginal time series across all subjects have been detrended. To quantify the strength of the FC, we calculate the sample correlations between the marginal time series observed from each ROI, for all subjects. Then each subject has $q = p(p-1)/2$ many ROI pairs of interest. Thus, testing on the FC network is conducted on N many q-variate vectors, where each vector comprises all pairwise sample correlations for each subject.

Each term of our proposed model is constructed as follows. Let r_{ij} denote the i-th sample correlation coefficient for the j-th subject and $\mathbf{Y} = (r_{11}, \ldots, r_{q1}, r_{12}, \ldots, r_{q2}, \ldots, r_{1N}, \ldots, r_{qN})$ be the vector of sample correlation coefficients stacked vertically across the subjects. Let the design matrix $\mathbf{X}$ be the $Nq \times q$ matrix of N many $q \times q$ identity matrices stacked vertically. The first error term, ϵ, has mean the zero vector and covariance matrix $\mathbf{\Sigma}$ and is used to model variability and covariability coming from the temporal autocorrelation in the ROI time series within each subject. To construct it, we first let ϵ_{ij} denote the i-th error term for the j-th subject and then $\epsilon_j = (\epsilon_{1j}, \ldots, \epsilon_{qj})' \sim N(\mathbf{0}, \Sigma_j)$. The diagonal elements of the covariance matrix Σ_j capture the variances of the correlation coefficients between each ROI for the j-th subject, and the off-diagonal elements capture the covariance between sample correlation coefficients observed between different pairs of ROIs within the j-th subject. We then construct $\mathbf{\Sigma}$ as a block-diagonal matrix with $\Sigma_1, \ldots, \Sigma_N$ along the diagonals. We simply stack the error term across all subjects, i.e., $\boldsymbol{\epsilon} = (\epsilon_1', \ldots, \epsilon_N')'$, so that this first error term has covariance matrix $\mathbf{\Sigma}$.

The second error term, ψ, has mean the zero vector and covariance matrix $\mathbf{\Psi}$ and controls for the heterogeneity of the subjects. To formulate it, let $\boldsymbol{\psi}_j \sim N(\mathbf{0}, \mathbf{\Psi}_0)$, where $\mathbf{\Psi}_0$ is a $q \times q$ scaled diagonal identity matrix. The covariance matrix $\mathbf{\Psi}$ is also a block-diagonal matrix with $\mathbf{\Psi}_0$ along the diagonals. Each element of $\mathbf{\Psi}_0$ represents the amount of variability that can be attributed to the sampling of the subjects across a population. Since $\mathbf{\Psi}_0$ is common across all subjects, we borrow information across all subjects to estimate this parameter. Similarly, for the second error term we simply stack across all subjects, i.e., $\boldsymbol{\psi} = (\psi_1', \ldots, \psi_N')'$, so that this error term has covariance matrix $\mathbf{\Psi}$.

To perform statistical inference on the FC network, our model accounts for i) the temporal autocorrelation in the marginal time series within each subject, ii) the covariance between the different pairs of sample correlation coefficients within each subject, and iii) the variability due to the sampling of the subjects from an underlying population. To this end, our proposed model takes the following form:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} + \boldsymbol{\psi}. \tag{1}$$

The parameter of interest that captures the true FC is the vector $\boldsymbol{\beta}$. We assume that the two error terms, $\boldsymbol{\epsilon}$ and $\boldsymbol{\psi}$, are independent of each