Published in final edited form as:

Biometrics. 2016 June ; 72(2): 596–605. doi:10.1111/biom.12433.

A Bayesian Hierarchical Framework for Modeling Brain Connectivity for Neuroimaging Data

Shuo Chen,

Department of Epidemiology and Biostatistics, University of Maryland, College Park, MD, 20742, USA

F. DuBois Bowman, and

Department of Biostatistics, Columbia University, Manhattan, NY, 10032, USA

Helen S. Mayberg

School of Medicine, Emory University, Atlanta, GA, 30322, USA

Shuo Chen: shuochen@umd.edu

Summary

We propose a novel Bayesian hierarchical model for brain imaging data that unifies voxel level (the most localized unit of measure) and region level brain connectivity analyses, and yields population level inferences. Functional connectivity generally refers to associations in brain activity between distinct locations. The first level of our model summarizes brain connectivity for cross-region voxel pairs using a two component mixture model consisting of connected and non-connected voxels. We use the proportion of connected voxel pairs to define a new measure of connectivity strength, which reflects the breadth of between-region connectivity. Furthermore, we evaluate the impact of clinical covariates on connectivity between region-pairs at a population level. We perform parameter estimation using Markov chain Monte Carlo (MCMC) techniques, which can be executed quickly relative to the number of model parameters. We apply our method to resting-state functional magnetic resonance imaging (fMRI) data from 32 subjects with major depression and simulated data to demonstrate the properties of our method.

Keywords

Bayesian hierarchical model; brain imaging; functional connectivity; MCMC; resting-state fMRI

1. Introduction

Functional brain connectivity refers to associations between neural processing units from distinct brain locations. Specific patterns of connectivity may be linked to corresponding actions, emotions, cognition, or resting-state (reflecting default mode of neural processing). Disruptions to these connectivity patterns are often associated with psychiatric and

Correspondence to: Shuo Chen, shuochen@umd.edu.

6. Supplementary Materials

Supplementary figures referenced in Sections (4.1, 4.2) and the software are available with this article at the Biometrics website on Wiley Online Library.

neurological disorders. Functional magnetic resonance imaging (fMRI) enables the assessment of functional connectivity by calculating the temporal coherence of neural activity measures arising from spatially distinct regions of the brain. There are several connectivity metrics available to measure the coherence of the temporal profiles. Examples include Pearson correlation; spectral coherence, which measures the coherence between the two temporal profiles in the frequency domain and mutual information, which measures the mutual dependence between the random variables giving rise to the observed two time series (Sun *et al.*, 2004). Modeling connectivity directly on fMRI time series often leads to a prohibitive computational burden, and thus most methods first calculate the connectivity metrics, then model these summaries to draw statistical inferences across subjects (Turk-Browne, 2013; Zatorre, 2013).

There are numerous intrinsic challenges to evaluating functional connectivity in the brain. Typical neuroimaging data in standard 2mm Montreal Neuroimaging Institute (MNI) space usually include roughly 300,000 voxels within the brain (on a regular lattice of $91 \times 109 \times 91$ voxels). Therefore, there are more than 40 billion voxel pairs for a single whole brain connectivity analysis. It is desirable for connectivity analysis to provide spatially detailed inferences between all voxel pairs (Turk-Browne, 2013). However, such ultra-high dimensionality poses a formidable computational load. Zalesky *et al.*, 2010 summarized the numbers of regions (or voxels) used for analysis in recent highly cited functional connectivity studies. A common strategy is to choose a small set of seed voxels among all voxels and calculate the associations between each seed voxel and all of the remaining voxels (Cordes *et al.*, 2001). A second approach is to define regions of interest (ROIs) using an existing anatomical parcellation of the brain, for example automated anatomical labeling (AAL) or the Brodmann brain atlas (Tzourio-Mazoyer *et al.*, 2002; Garey, 1994) and to quantify connectivity at the region level rather than voxel level (Greicius *et al.*, 2003; Zalesky *et al.*, 2012b). This approach usually first summarizes the temporal profiles of all voxels within each region by taking averages (or similar dimension reduction) and then estimating the connectivity by using the regional summaries. In addition, independent component analysis (ICA) has been applied to assess connectivity by separating spatially independent sources of signals (Calhoun *et al.*, 2001).

Recently, there has been a focus on the application of graph theory metrics to depict complex brain networks (Bullmore and Sporns, 2009; Simpson *et al.*, 2013; Bowman *et al.*, 2014). Graph theory metrics reflect important topological properties of brain networks. For example, the “small-worldness” property indicates that most neural units (voxels) are often directly connected to neighbors (e.g. high positive Pearson correlation), but distant units can be connected through only a small number of steps via highly connected nodes (‘hub’ nodes) (Achard *et al.*, 2005). The “small-worldness” property is common in highly organized networks such as the internet, vehicle traffic, and social networks. Nevertheless, those metrics usually only function as summary statistics rather than as motifs for establishing brain connectivity.

The seed voxel approach is simple for computation, but it is inherently limited in scope as it ignores functional connections other than the ones associated with the seeds. The region based method neglects the variations within regions, which may lead to substantial bias and

information loss, and these problems are exacerbated as the size of an ROI increases. Moreover, region-based methods lack detail on the particular voxels that are connected between pairs of regions. ICA is often conducted on the voxel level where detected components may be difficult to interpret and may span multiple brain regions. Bayesian modeling has been applied to estimate localized brain activity with corresponding group/task differences and to perform region level connectivity analysis using task-induced fMRI data (Bowman *et al.*, 2008; Patel *et al.*, 2006; Xue *et al.*, 2015). In summary, most existing statistical methods for brain connectivity analysis are conducted either on the region level or on the voxel level alone. There is an unmet need to develop analytical methods to unify connectivity analyses at these two levels of spatial units. We provide a new statistical framework that unifies voxel and region level connectivity analyses.

We propose a novel Bayesian hierarchical model for inferring brain connectivity. We use Pearson correlation, which is the most widely used connectivity metric in the neuroimaging community as the input data for our method (Zalesky *et al.*, 2012a; Zatorre, 2013; Greicius *et al.*, 2003). The base level of our framework assumes that the connections between voxel pairs from two brain regions follow a mixture distribution with two modes reflecting connected and non-connected voxel pairs (i.e. high and low connectivity metrics). The assumption follows naturally from the “small-worldness” property of brain networks, since two distant ROIs may be well connected through a small proportion of efficiently connected nodes (represented by the mixture component with higher connectivity metric values). The rest of the nodes in these two regions may not have direct across-region connections, but they could communicate via those efficiently connected voxels (connected component). Our model utilizes the connectivity information from all cross-region voxel pairs. From parameters in our mixture model, we introduce a new statistic called connectivity breadth, which reflects a measure of connectivity strength. We adjust for an individual’s covariate effects such as age, gender, and psychiatric conditions in a generalized linear model (GLM).

The hierarchical Bayesian model unifies voxel level and region level connectivity by recognizing the distribution of voxel connectivity between regions, which allows us to fully utilize rich voxel level connectivity information while keeping the more interpretable region level connectivity findings. At the voxel level, our model provides inferences to identify voxels (within a region) acting as ‘hubs’ by exhibiting many connections to another region. In addition, our model identifies cross-region voxel pairs that exhibit strong connectivity. Thus, two regions may be functionally connected via a hub in at least one of the regions or by two high-traffic ports directly linking the regions. Our method could have important applications, for example to reveal differences in connectivity properties (or other potential biomarkers) between patients with Parkinson’s disease and healthy controls. We estimate the parameters from our model using Markov chain Monte Carlo (MCMC) methods and draw inferences based on the joint posterior probability distribution for all of the parameters.

2. Motivating Example

There is currently a major focus in the neuroimaging literature on resting-state functional connectivity, which examines connectivity properties during an undirected mental state without the explicit performance of tasks. We consider resting-state fMRI data from 32

patients with major depressive disorder (MDD) participating in a depression treatment biomarker study (Dunlop *et al.*, 2012). The sample includes 12 males and 20 females ages 24 to 57 years. We have data from baseline fMRI scans acquired prior to treatment randomization, a baseline 17 item Hamilton Depression Rating Scale (HAM-D17), and covariates age and gender. The patient cohort was divided into two groups based on the HAM-D17 score, as suggested by Maruish (1999) and Schutte and John (1995). The resulting classification includes 11 mildly depressed subjects, with HAM-D17 greater than 15 and less than 18, and 21 severely depressed subjects with HAM-D17 18 or above. For each subject, 150 fMRI volumes were obtained in 7.5 minutes during a resting-state, in which subjects are left to think for themselves while visually focusing on a cross displayed on a monitor in the scanner. We performed standard data preprocessing steps, including motion correction to offset subject movement in the scanner, slice timing correction to adjust for the fact that each 3D scan representing a single time point is actually acquired as a series of 2D slices over time, normalization to adopt a common reference space across subjects, and spatial smoothing. Also, we apply de-trending and demeaning steps to remove low frequency trends unrelated to neural processing (e.g. due to scanner hardware) and to adjust for spatial variation in brain tissue properties that may affect the fMRI signal. Pre-whitening is performed to remove the temporal correlation of the time course at each voxel.

Models of depression emphasize dysfunctional interactions among specific cortical and limbic brain areas. Particularly, the interactions between subdivisions of frontal cortex and cingulate gyrus have been repeatedly reported in depression studies by using multiple complementary imaging methods (Seminowicz *et al.*, 2004; Mayberg *et al.*, 2005; Greicius *et al.*, 2007; and Fox *et al.*, 2012). We focus on the interactions between some of those reported ROIs (based on the Brodmann atlas) in the prefrontal cortex (PFC) including medial frontal cortex (mF10), orbital frontal cortex (OF11), and lateral prefrontal cortex (latF9); subgenual cingulate cortex (Cg25) and anterior cingulate (Cg24). We show these regions of interest in Figure 1. The goal is to assess functional connections between the five ROIs and to make inferences about how depression severity influences them.

Exploring the distribution of voxel-level functional connections (using correlations for example) between ROIs, we note the appearance of a mixture distribution with two components: a larger component of relatively weakly connected voxel pairs and a smaller component of more strongly connected voxel pairs (see Figure 2). Moreover, each subject has a different number of connected voxel pairs. The size of the connected component represents the connectivity breadth, and reflects aspects of the the “small-worldness” property. Therefore, we develop a method to study such connectivity breadth and to identify what factors may influence this connectivity breadth.

3. Methods

We formulate a model that builds on the calculated voxel-level functional connections within ROIs and between ROIs. ROIs are labeled as $g = 1, \dots, G$, where we may set G as high as 116 for the AAL map and 48 for the Brodmann map (for our data example $G = 5$). We use $z_{ig\bar{v}}$ to denote connectivity (here, we consider correlations with a Fisher transformation) between the cross-region pair of voxels v and v' denoted by v, v' for regions $g (v \in g)$ and $g' (v' \in g')$.

$\in g'$), labeled by g ; and for subject i ($i = 1, \dots, n$). Throughout, we use the two-dot notation to indicate pairs of voxels or regions. For a region pair g , the total number of voxel pairs is V_g . Thus, if regions g and g' include V_g and $V_{g'}$ voxels respectively, then $V_{g \cdot g'} = V_g \times V_{g'}$.

3.1 Bayesian Hierarchical Model

We use a Bayesian hierarchical framework to jointly model functional connections between pairs of voxels and pairs of ROIs along with clinical covariate effects while tracking the variability at each level (see Figure 3).

3.1.1 Distribution for connectivity—The first level of our model specifies functional links between ROIs for each subject via a distribution of connections from cross-region voxel pairs. A driving heuristic is that according to the “small-worldness” property of brain networks, two regions can be efficiently connected through only a small proportion of connected voxels, and this property is also borne out by descriptive assessments of our data. In the two component mixture model, we limit our attention to only non-negative correlations, following a strong precedent in the neuroimaging literature, which facilitates interpretations. (Weissenbacher *et al.*, 2009; Fox *et al.*, 2009).

At the first level, we assume that $z_{ig\ddot{v}}$ follows a mixture distribution with two modes representing connected and non-connected voxel pairs. The two components vary around means μ_{0ig} (non-connected) and μ_{1ig} (connected), with the variances σ_{0ig}^2 and σ_{1ig}^2 respectively. We set $\mu_{1ig} > \mu_{0ig}$ to aid identifiability. Specifically, the model is given by

$$z_{ig\ddot{v}} | \omega_{ig\ddot{v}}, \mu_{0ig}, \sigma_{0ig}^2, \mu_{1ig}, \sigma_{1ig}^2 \sim \omega_{ig\ddot{v}} N(\mu_{1ig}, \sigma_{1ig}^2) + (1 - \omega_{ig\ddot{v}}) N(\mu_{0ig}, \sigma_{0ig}^2),$$

$$\omega_{ig\ddot{v}} | \pi_{ig} \sim \text{Bernoulli}(\pi_{ig}), \text{ and correspondingly } \sum_{\ddot{v}} \omega_{ig\ddot{v}} = \Omega_{ig} | \lambda_{ig} \sim \text{Poisson}(\lambda_{ig}), \quad (1)$$

where $\lambda_{ig} = V_g \pi_{ig}$.

The latent indicator variable $\omega_{ig\ddot{v}}$ classifies a voxel pair as either connected or non-connected. Our interest lies in detecting connectivity breadth between ROIs, which is represented by the average number of connected voxel pairs $\lambda_{ig} = V_g \pi_{ig}$. Note that $\omega_{ig\ddot{v}}$ in line 2 may be correlated with $\omega_{ig\ddot{v}}$; and then the indicator variables of all cross-region voxel pairs are a set of dependent Bernoulli trials (connected 1 or not connected 0). Moreover, we

let the marginal sum $\Omega_{ig} = \sum_{\ddot{v}} \omega_{ig\ddot{v}}$ follow a Poisson distribution and link the hyperprior parameters with λ_{ig} because i) the information of parameters from Bernoulli distribution and Poisson distribution are almost equivalent ($\lambda_{ig} = V_g \pi_{ig}$) and ii) Ω_{ig} is more sensible and computationally convenient for group level inferences. It has been shown that the Poisson distribution has a bounded approximation for the sum of dependent Bernoulli trials (Chen, 1975). The approximation should be very close in our application since V_g is on the order of hundreds of thousands in many cases. Therefore, using independent or dependent correlation structure between all voxel pairs $\omega_{ig\ddot{v}}$ does not impact the marginal parameter estimation of λ_{ig} . Our approximation is analogous to the areal/lattice spatial data analysis in disease

mapping, where each incident (dependent on each other) follows Bernoulli distribution yet within an area (e.g. county) the marginal sum incident count is often used as an outcome variable for association analysis.

Why two component mixture?: We adopt the two component mixture model reflecting that voxels exhibit connectivity or not, rather than specifying a finite mixture model with $k > 2$ or an infinite mixture model. Our formulation mirrors the use of two component mixtures in other massive hypothesis testing settings, e.g. local *fdr* (Efron, 2004; George and McCulloch, 1993; Do *et al.*, 2005; Müller *et al.*, 2007). The common target is to identify the true positive elements from a mixture model consisting of a dominating null distribution and a small number of positive events (here, functional connections). Whereas the local *fdr* method tends to focus on control rates at specific z values or events, we are interested in the sum of latent indicator variables Ω_{ig} . Marginally,

$E(\Omega_{ig}) = \sum_v^{V_g} \int P(z_{igv}) f(z_{igv}) dz_{igv} = V_g \pi_{ig} = \lambda_{ig}$, where $f(z_{igv}) = \pi_{ig} f_1(z_{igv}) + (1 - \pi_{ig}) f_0(z_{igv})$; f_0 and f_1 represent the probability density functions associated with the null and positive components, and $P(z_{igv})$ is the probability of “true positive”, which is identical to $1 - fdr$. The marginal variable Ω_{ig} is subject to less variance and is invariant to the choice of distributions for f_0 and f_1 .

In addition, the conceptualization as a two component mixture mitigates the cross-subject identifiability issues. If each subject includes k ($k > 2$) components, we ought to ensure that not only all components (with different center parameters) are identifiable, but also that the components can be matched across subjects for group-level inferences. The matching procedure pose challenges with respect to numerical optimization.

3.1.2 Prior specifications—The second level of the hierarchical model in (2) expresses a prior belief that each subject’s mixture component means μ_{0ig} and μ_{1ig} arise from normal distributions with means u_{0g} and u_{1g} and variances τ_{0g}^2 and τ_{1g}^2 respectively. We anticipate that the mean parameter u_{1g} for the cross-region voxel pairs from the connected component will be relatively large. Thus, we empirically inform u_{1g} by calculating the mode of voxel pair functional connections within an ROI at geometric distances up to 30 mm for each subject, which seems to represent a reasonable starting point since voxel pairs within anatomically-defined regions generally exhibit high association.

$$\begin{aligned} \mu_{0ig} | u_{0g}, \tau_{0g}^2 &\sim N(u_{0g}, \tau_{0g}^2), \\ \mu_{1ig} | u_{1g}, \tau_{1g}^2 &\sim N(u_{1g}, \tau_{1g}^2), \\ \sigma_{0ig}^{-2} &\sim \text{Gamma}(a_0, b_0), \\ \sigma_{1ig}^{-2} &\sim \text{Gamma}(a_1, b_1), \\ \log(\lambda_{ig}) &= \beta_{g0} + \beta_{gi} x_i, \text{ where } \lambda_{ig} = E(\Omega_{ig} | \beta_{g0}, \beta_{gi}), \quad (2) \end{aligned}$$

We inform the mode of u_{0g} by using the mode of voxel pair functional connections between ROIs across all subjects which is close to zero. The variances τ_{0g}^2 and τ_{1g}^2 reflect the confidence of the prior mean specification, and we set τ_{0g}^2 and τ_{1g}^2 as a relatively small value

as the mean priors are informative. For the first level variance parameters σ_{0ig}^2 and σ_{1ig}^2 , we assume Gamma distributions with hyperprior parameters $a_0 = 0.1$, $b_0 = 0.005$, $a_1 = 0.1$, $b_1 = 0.005$, reflecting vague priors.

We model the region pair mean connectivity breadth λ_g through a GLM with a *log* link to accommodate the Poisson distribution. The parameters β_g of the GLM capture covariate effects such as disease status, age, and treatment group. The priors for β_g are set vaguely enough to let the large amount of observed data weigh heavily in determining the posterior distribution, specifically using

$$\beta_{g0} \sim N(0, \tau_{0b}^2), \quad \beta_g \sim N(0, \tau_b^2). \quad (3)$$

3.2 Estimation and posterior inference

We apply MCMC algorithms to estimate the parameters in our hierarchical probability model. The priors for the main effects β in the GLM are not conjugate, so we use a Gibbs sampler with Metropolis updates where needed for posterior sampling. The full conditional distributions are given by the following:

$$\begin{aligned} \omega_{ig\bar{v}} &\sim \text{Bernoulli} \left(\frac{\pi_{ig} N(z_{ig\bar{v}} | \mu_{1ig}, \sigma_{1ig}^2)}{\pi_{ig} N(z_{ig\bar{v}} | \mu_{1ig}, \sigma_{1ig}^2) + (1 - \pi_{ig}) N(z_{ig\bar{v}} | \mu_{0ig}, \sigma_{0ig}^2)} \right) \\ \mu_{0ig} &\sim \text{Normal} \left(\Psi_{0ig} \{ \tau_{0ig}^{-2} u_{0ig} + (V_g - \Omega_{ig}) \sigma_{0ig}^{-2} \bar{z}_{0ig} \}, \Psi_{0ig} \right) \\ \mu_{1ig} &\sim \text{Normal} \left(\Psi_{1ig} \{ \tau_{1ig}^{-2} u_{1ig} + \Omega_{ig} \sigma_{1ig}^{-2} \bar{z}_{1ig} \}, \Psi_{1ig} \right) \\ \sigma_{0ig}^{-2} &\sim \text{Gamma} \left\{ a_0 + (V_g - \Omega_{ig})/2, \left(\frac{1}{b_0} + \frac{1}{2} \sum_{\forall \omega_{ig\bar{v}}=0} (z_{ig\bar{v}} - \mu_{0ig})^2 \right)^{-1} \right\} \\ \sigma_{1ig}^{-2} &\sim \text{Gamma} \left\{ a_1 + (\Omega_{ig})/2, \left(\frac{1}{b_1} + \frac{1}{2} \sum_{\forall \omega_{ig\bar{v}}=1} (z_{ig\bar{v}} - \mu_{1ig})^2 \right)^{-1} \right\} \\ \beta_{g0} &\propto \prod_i e^{(\beta_{g0} + \beta_g x_i) \Omega_{ig}} \exp(-e^{(\beta_{g0} + \beta_g x_i)}) N(\beta_{g0} | 0, \tau_{0b}^2) \\ \beta_g &\propto \prod_i e^{(\beta_{g0} + \beta_g x_i) \Omega_{ig}} \exp(-e^{(\beta_{g0} + \beta_g x_i)}) N(\beta_g | 0, \tau_b^2), \end{aligned} \quad (4)$$

where $\Psi_{0ig} = ((V_g - \Omega_{1g}) \sigma_{0ig}^2 + \tau_{0ig}^2)$, $\Psi_{1ig} = (\Omega_{ig} \sigma_{1ig}^{-2} + \tau_{1ig}^{-2})$, $\bar{z}_{1ig}$ is the mean for all cross-region voxel pairs $\bar{v}$ from the connected component, i.e. $\omega_{ig\bar{v}} = 1$, $\bar{z}_{0ig}$ is the mean for all $\bar{v}$ of $\omega_{ig\bar{v}} = 0$, and $\pi_{ig} = e^{\beta_{g0} + \beta_g x_i} / V_g$.

When implementing the MCMC procedure, we first sample the latent indicator variable, then sample the four parameters of the Gaussian mixture model. In formula 4 line 1, each latent indicator variable $\omega_{ig\bar{v}}$ for the corresponding voxel pair $\bar{v}$ is conditional on the region level parameter π_{ig} , which is equivalent to λ_{ig} / V_g . λ_{ig} is determined by β_g (along with the subject covariates) and β_{g0} ; and both β_g and β_{g0} are sampled by incorporating $\Omega_{ig} = \sum_{\bar{v}}^{V_g} \omega_{ig\bar{v}}$ (see the last two lines of formula 4). Note that Ω_{ig} is updated by taking the summation of updated $\omega_{ig\bar{v}}$ rather than sampled directly from $\text{Poi}(\lambda_{ig})$. Therefore, similar to other mixture

model settings, our sampling scheme includes an implicit constraint that each ω_{igv} is dependent on $[\omega_{igv}, \omega_{igv}, \dots]$ because the update of $\pi_{ig}(\lambda_{ig}/V_g)$ depends on connections of all voxel pairs between a region pair i.e. $\pi_{ig} = \lambda_{ig}/V_g = g^{-1}(\beta_{g0}, \beta_g)$ and the updates of β_{g0} and β_g are both dependent on Ω_{ig} (the sum of $[\omega_{igv}, \omega_{igv}, \omega_{igv}, \dots]$).

At the second level, the GLM parameters are updated through Metropolis steps with input of Ω_{ig} following a Poisson (*log* link). We favor a Poisson distribution, rather than a Binomial, because with Ω_{ig} and V_g being on the order of thousands and millions, respectively, the ratio function of the Metropolis step may easily generate zeros in the denominator if the *logit* link is used. In addition, since the number of voxel pairs V_g is often large, the inferences for π_{ig} can be directly computed as λ_{ig}/V_g from the posterior simulations. The ratios of β s in the Metropolis step can be expressed as:

$$r_{\beta} = \exp \left\{ \sum_i^n [\Omega_{ig}(\beta'_{ig} - \beta_{ig}) + \exp(-\eta'_{ig}) - \exp(-\eta_{ig})] + \log(p_0(\beta'_{ig}|0, \tau_b^2)) - \log(p_0(\beta_{ig}|0, \tau_b^2)) \right\},$$

where β' are proposed values, p_0 indicates the prior distribution, and we accordingly obtain η_{ig} from $\eta_{ig} = \beta_{g0} + \beta_g x_i$. The proposed value will be accepted with probability $\min(1, r_{\beta})$. All other parameters can be sampled from known distributions through the Gibbs sampler as shown in (4).

Implementation of our Bayesian model takes samples from the joint posterior distribution for all of the model parameters, which enables statistical inferences regarding connectivity and brain network properties. First, we can make inferences about voxel-level connectivity for each subject. For example, we can target inferences pertaining to the probability of a cross region voxel pair being connected, say $P_{ig}(v)$, by calculating the proportion of $\omega_{ig}(v)$ equaling one. Secondly, we can identify which voxels operate as hubs, with a hub defined here as a voxel that connects with numerous cross-region voxels. Thirdly, we can determine which specific pairs of voxels are highly connected to each other. We can also make inferences about region-level connectivity. For example, we can draw inference about the connectivity breadth of a region pair by calculating the posterior probability of $e^{\beta_{g0} + \beta_g x_i}$. Relevant to our upcoming data analysis, our modeling framework yields inferences at the population level, and we can investigate the impact of covariate effects such as age, disease groups, and treatment responses on connectivity breadth between ROIs.

4. Results

We apply our Bayesian hierarchical model to resting-state fMRI data from the depression study and empirically evaluate our model using simulated data. We compare our method with the region average method which is the most commonly used group-wise region level analysis approach under the predefined brain ROI atlases.

4.1 MDD Study Using Resting-state fMRI Data

In the MDD study, our goal is to explore the neural circuits associated with medial and lateral regions of prefrontal cortex (e.g. latF9, mF10, OF11) and effective subdivisions of the

anterior Cingulate Cg25 and Cg24 as described in section 2. We employ the commonly used Pearson correlation coefficient as a measure of the temporal coherence between fMRI time series from distinct voxels and further apply the Fisher normalizing transformation. As demonstrated in Figure 2, the functional connections appear to follow a mixture distribution. We apply our Bayesian hierarchical model to the transformed correlation metrics to investigate functional connectivity. For illustration, we highlight results for the region pair mF10 and Cg25.

Inferences for voxel pair connectivity: After applying our Bayesian model, we estimate the probability of functional connectivity between each cross-region voxel pair. Then, we identify voxels in one brain region that are highly connected to voxels in the other region, representing hubs at the population level. In Figure 4a, each row represents a voxel in mF10, each column represents a subject, and the intensity indicates the corresponding number of functionally connected voxels in Cg25. For convenience, we have ordered the voxels in mF10 according to the number of functionally connected voxels (across subjects), with the most highly connected voxels appearing at the top and with the least connected voxels appearing at the bottom. We note that there are several mF10 voxels that are highly connected to region Cg25 across all 32 subjects, while most voxels are weakly connected. These highly connected voxels serve as “hubs” connecting the two regions since they reveal coherent activity with many other voxels. Reciprocally, Figure 4b shows the voxels in Cg25 connected to mF10. Specifically, we consider the voxels that are cumulatively connected to roughly 30% or more of the voxels in the other region. For example, in considering hubs linking voxels in mF10 to those in Cg25, we identify voxels with degree > 60 (the degree of a voxel equals to the number of its connected voxels) and voxels in Cg25 with degree > 150 . The joint posterior distribution also allows us to identify particular voxel pairs connecting the two regions. Figure 4c shows the highly connected voxel pairs between the two regions, and Figure 4d illustrates locations of the functionally connected voxel pairs linking regions mF10 and Cg25. The posterior also yields inferences to identify highly connected voxels for each subject, and note from Figure 4 that connections between voxels vary across subjects. Our framework has particular utility in reforming models of depression relevant to focal neuromodulation where targeting of discrete pathways may be critical (Mayberg *et al.*, 2005; Gutman *et al.*, 2009; and Fox *et al.*, 2012).

Inferences for region pair connectivity: At the region level, we summarize the breadth of functional connectivity, stratified by different depression severity groups, by counting the number of the functionally connected voxel pairs. In Table 1, we note that the region pairs mF10 and Cg25, OF11 and Cg25, and Cg24 and Cg25 are relatively highly connected, but not latF9 and Cg25C. There are marked differences in connectivity breadth between mildly and severely depressed patients for mF10 and Cg25 (7.5%), latF9 and Cg25 (6%), and OF11 and Cg25 (10%) with more severely depressed patients having increased functional connections. These findings may reveal that with increased depression severity, the brain locations for emotion processing interact more with the prefrontal cortex neural units. There is very little difference, however, in extent of functional connections between groups for Cg24 and Cg25 (2% difference), with severely depressed patients exhibiting slightly reduced connectivity breadth.

For comparison, we apply a commonly used approach that calculates correlations between the regional temporal profiles obtained as an average of all temporal profiles in mF10 and Cg25 for all subjects. There is no difference between the two groups with t test p value at 0.57 (Supplementary Figure 2). We further address differences between our proposed method and the regional representative approach in a simulation study. We conclude that our model is able to reveal the subtle difference driven by the relatively few connected voxel pairs, and accordingly to detect the group difference. Overall, we find that that cg25 and frontal cortices appeared hyper-connected in more depressed subjects, which is consistent with previous research findings about the frontal-limbic circuitry (Greicius *et al.*, 2007; Greicius, 2008; Zeng *et al.*, 2012). Besides the similar region-level connectivity findings above, we provide additional voxel level inferences that previous methods do not permit (e.g. Figure 4 results); for instance we can identify those highly-connected voxel pairs that effectively connect a pair of regions.

We check our MCMC runs to ensure the stability of the results and the convergence of the chains. However, the massive number of voxel pairs prevents formal use of MCMC convergence tools for all model parameters. We monitor the trace plots for all region level parameters and for parameters corresponding to select voxel pairs. For each MCMC, we sample 20,000 iterations in total and set the first 500 iterations as burn-in. The resulting Markov chain converges within 500 burn-in samples (Supplementary Figure 1).

4.2 Simulation Study

We conduct a simulation study to further examine properties of our proposed connectivity modeling framework. We generate resting-state fMRI time series based on one region pair with sizes of a $7 \times 7 \times 7$ voxel cube and a $5 \times 5 \times 5$ voxel cube, from 20 subjects (10 controls and 10 cases). Therefore, for each subject we simulate $7^3 + 5^3 = 468$ fMRI time series, which allows us to calculate the connectivity metrics of $V_g = 7^3 \times 5^3 = 42875$ between-region voxel pairs. First, we randomly generate the numbers of connected cross-region voxel pairs using Poisson distributions with mean parameter:

$$\Omega_{ig} = \exp(\beta_{g0} + \beta_{gi} x_i).$$

We use $\beta_{g0} = 8$ and $\beta_{gi} = 0.9$ to indicate that the controls have more breadth in connectivity than cases do. Second, we construct the correlation matrix R (468×468) for all voxels in the two regions, which determines the correlations between the fMRI time series. We set each group with different R s since we use different Ω_{ig} and $V_g \sim \Omega_{ig}$ to indicate the numbers of highly correlated (e.g. correlation coefficient=0.6) and uncorrelated (e.g. correlation coefficient=0) between-region voxel pairs for each group. For the within region voxel pairs, we use correlations on the geometric distances between voxels so that voxels close to each other are more highly correlated and distant voxel pairs are less correlated. We slightly adjust R to ensure that it is positive definite. Based on different R s of case and control groups, we generate 468 resting-state fMRI time series for each of the 20 subjects by using multivariate normal distribution with mean zero and covariane of R (Supplementary Figure 3 demonstrates the resting-state fMRI time series for 10 voxels in one subject). We assume that the pre-whitening has been performed, so there is no temporal autocorrelation within

each voxel. We also apply different variance parameters of multivariate normal distribution (power magnitude) of the BOLD signal and band filtering to remove the low and high frequency components. The tuning of these simulation parameters does little change to the resulting distribution of correlation (connectivity metric) which is the input of our method.

Based on the simulated resting-state fMRI time series, we calculate correlation coefficients of the between region voxel pairs for each subject (see histograms in Supplementary Figure 4), where the controls have a larger component of connected voxel pairs. Our main goal is to detect a group effect and to check the number of voxel pairs that are determined to be connected for each subject. We apply our proposed hierarchical Bayesian framework to the simulated data.

Based on trace plots of parameters for the group effects from our MCMC samples of 20,000 iterations (Supplementary Figure 5), we note that the chains converge after around 300 iterations and henceforth tightly vary around the true values. The results shown in Table 2 reveal that our model captures the group effect accurately. Thus, this simulation supports that the proposed method and estimation algorithm can detect differential connectivity breadth when they exist. The small bias of β s may come from the false negative classification of the mixture model due to the heavy inclination of the predominant null component.

As a comparison, we use the regional representative method described in the analysis of our depression data, to detect the connectivity difference between the two groups. We first calculate the regional temporal profiles for each subject using the average of all of the fMRI time series in each region. Then, we measure the correlation coefficients between the two regional fMRI time series for each subject. We conduct a t-test to assess the null hypothesis that the two regions have the same connectivity between the two groups, and the result shows non-significant difference with $p = 0.673$. We note that our method detects the informative group difference that the region representative method fails to identify. Our method also provides richer information about voxel level connectivity. There are many methods that quantify connectivity, in a general sense, between pairs of ROIs across different clinical groups. However, as we limit our attention to compare connectivity between one pair of ROIs, we only compare our approach with the region representative method.

5. Discussion

Our proposed method is the first model that is able to use all available voxel pair connectivity information, while also capturing connectivity between region pairs. Given the massive number of voxel pairs across the entire brain (tens of billions), it is imperative for any voxel level approach to reduce the number of pairs considered, and we do so by focusing on particular ROIs. Another attribute of our model is that it naturally lends itself to network graph theoretical properties characterizing brain connectivity. Our Bayesian model yields inferences on numerous quantities reflecting brain connectivity. Our model enables the examination of covariate effects, and we illustrate this by considering how depression severity impacts the extent of brain connectivity. Estimation of regional brain connectivity is

robust to potential bias due to imperfect image registration because the posterior inferences are not dependent on the specific locations of voxel pairs between a region pair but only on the marginal connectivity between a region pair. At the voxel level, we may also detect inter-region “hub” voxels as those that are connected to many other voxels with high probabilities, and the voxel pair connectivity pathways that connect the two regions with high probabilities across subjects. The relative size of the connected component, i.e. the proportion of connected voxel pairs, could arguably be affected by the physical distance between the pair of ROIs, e.g. yielding increased connectivity breadth between two adjacent regions. However, increased connectivity breadth between adjacent (or relatively close) regions arguably may have a valid neurophysiological basis. Regardless, the influence of distance should have little bearing on our statistical inferences, for instance the between group comparisons, since all subjects share the identical spatial adjacency information.

The mixture model in the first level determines the number of connected voxel pairs in a fashion that is similar to *local fdr*, rather than choosing hard cutoff points, which improves the accuracy. Although we choose Gaussian distributions for both components of the mixture model (which seems to work well for our data example), we could adopt more flexible distributions (e.g. a gamma distribution for the skewed connected component). Both the MDD study and our simulation study results indicate that the MCMC parameter estimation is satisfactory, and the trace plots of the chains behave normally. Although there are numerous voxel pairs per subject and many subjects are included, the computational load is manageable. All the programming including MCMC and imaging processing was implemented in MATLAB. With a sample size of 20 subjects and more than 40,000 voxel pairs, the MCMC procedure takes around 6 minutes for the simulation study and on average 15 minutes for the example data to run for one region pair on an Intel i7 3.4 GHz CPU and 24G memory PC. When modeling a larger number of region pairs, we would leverage the appropriate parallelization algorithms for faster implementation. It will cost approximately 60 hours to calculate all region pairs of the Brodmann atlas in 2mm standard space by using parallel computing.

In this article, our model does not explicitly account for the correlation between different pairs of voxel connections. In contrast to activity analyses, where spatial correlations are more readily addressed (Bowman *et al.*, 2008; Derado *et al.*, 2010), connectivity studies typically consider the association between time courses of two brain units as input data. Thus, the dependence structure between two connectivity metrics may rely on the spatial information of four nodes. Numerous complexities are involved in modeling this spatial dependence structure, and the commonly used assumptions for spatial dependence (e.g. second-order stationarity) may need to be cautiously tested and verified. Thus, while accounting for the spatial dependence structure between connections of brain areas (including region level and voxel level) is certainly an important consideration, the topic will require in depth investigation in future work. Our proposed model is also applicable to other functional connectivity metrics (e.g. mutual information and spectral coherence), and other brain connectivity measures such as structural connectivity from DTI data.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was partially supported by the National Institutes of Health grant U18 NS082143.

References

- Achard S, Salvador R, Whitcher B, Suckling J, Bullmore E. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *Journal of Neuroscience*. 2005; 26:63–72. [PubMed: 16399673]
- Bowman FD, Caffo B, Bassett SS, Kilts C. A Bayesian hierarchical framework for spatial modeling of fMRI data. *NeuroImage*. 2008; 39(1):146–156. [PubMed: 17936016]
- Bowman FD. Imaging Analysis. *Annual Review of Statistics and Its Application*. 2014; 1:61–85.
- Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*. 2009; 10:186–198. [PubMed: 19190637]
- Calhoun VD, Adali T, Pearlson GD, Pekar JJ. A method for making group inferences from fMRI data using ICA. *Human Brain Mapping*. 2001; 14(3):140–151. [PubMed: 11559959]
- Chen LH. Poisson approximation for dependent trials. *The Annals of Probability*. 1975:534–545.
- Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Meyerand ME. Frequencies contributing to functional connectivity in the cerebral cortex in rs data. *American Journal of Neuroradiology*. 2001; 22(7):1326–1333. [PubMed: 11498421]
- Derado G, Bowman FD, Kilts C. Modeling the spatial and temporal dependence in fMRI data. *Biometrics*. 2010; 66:949–957. [PubMed: 19912175]
- Do KA, Muller P, Tang PA. Bayesian mixture model for differential gene expression. *Journal of the Royal Statistical Society: Series C*. 2005; 54(3):627–44.
- Dunlop BW, Binder B, Cubells F, Goodman G, Kelley E, Kinkead B, et al. Predictors of Remission in Depression to Individual and Combined Treatments (PREdict): Study Protocol for a Randomized Controlled Trial. *Trials*. 2012; 13(1):106. [PubMed: 22776534]
- Efron B. Large-scale simultaneous hypothesis testing. *Journal of American Statistical Association*. 2004; 99:96–104.
- Fox MD, Zhang D, Snyder Z, Raichle E. The global signal and observed anticorrelated rs brain networks. *Journal of Neurophysiology*. 2009; 101(6):3270–3283. [PubMed: 19339462]
- Fox MD, Buckner RL, White MP, Greicius MD, Pascual-Leone A. Efficacy of transcranial magnetic stimulation targets for depression is related to intrinsic functional connectivity with the subgenual cingulate. *Biological Psychiatry*. 2012; 72(7):595–603. [PubMed: 22658708]
- George EI, McCulloch RE. Variable Selection via Gibbs Sampling. *Journal of American Statistics Association*. 1993; 88(423):881–889.
- Greicius MD, Flores BH, Menon V, Glover H, Solvason B, Schatzberg F. rs functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus. *Biological Psychiatry*. 2007; 62(5):429–37. [PubMed: 17210143]
- Greicius M. Resting-state functional connectivity in neuropsychiatric disorders. *Current Opinion in Neurology*. 2008; 21(4):424–430. [PubMed: 18607202]
- Greicius MD, Krasnow B, Reiss AL, Menon V. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*. 2003; 100(1):253–258.
- Gutman D, Holtzheimer P, Behrens T, Johansen-Berg H, Mayberg H. A tractography analysis of two deep brain stimulation white matter targets for depression. *Biological Psychiatry*. 2009; 65(4):276–82. [PubMed: 19013554]
- Garey, L. Brodman's localisation in the cerebral cortex: the principles of comparative localisation based on cytoarchitectonics. Springer; London: 1994.

- Jirsa, K.; McIntosh, R. Handbook of Brain Connectivity. Springer; 2007.
- Maruish, MR. The Use of Psychological Testing for Treatment Planning and Outcomes Assessment. Lawrence Erlbaum Associates; 1999.
- Mayberg HS, Lozano AM, Voon V, McNeely HE, Seminowicz D, Hamani C, et al. DBS for treatment-resistant depression. *Neuron*. 2005; 45(5):51–60.
- Müller, P.; Parmigiani, G.; Rice, K. FDR and Bayesian multiple comparisons rules. *Proc. Valencia/ISBA 8th World Meeting on Bayesian Statistics*; 2006.
- Patel RS, Bowman FD, Rilling JK. A Bayesian approach to determining connectivity of the human brain. *Human Brain Mapping*. 2006; 27(3):267–276. [PubMed: 16092131]
- Simpson S, Bowman FD, Laurienti P. Analyzing Complex Functional Brain Networks: Fusing Statistics and Network Science. *Statistics Surveys*. 2013; 7:1–36. [PubMed: 25309643]
- Sporns, O. Networks of the Brain. MIT Press; 2010.
- Schutte, NS.; John, M. Sourcebook of Adult Assessment Strategies. New York: 1995.
- Seminowicz A, Mayberg H, Segal Z, Rafi-Tari S. Limbic-frontal circuitry in major depression: a path modeling metanalysis. *Neuro Image*. 2004; 22(1):409–18. [PubMed: 15110034]
- Sun F, Miller L, D’Esposito M. Measuring interregional functional connectivity using coherence and partial coherence analyses of fmri data. *Neuro Image*. 2004; 21:647–658. [PubMed: 14980567]
- Turk-Browne NB. Functional Interactions as Big Data in the Human Brain. *Science*. 2013; 342(6158): 580–584. [PubMed: 24179218]
- Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Mazoyer B, Joliot M. Automated anatomical labelling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single subject brain. *Neuro Image*. 2002; 15:273–289. [PubMed: 11771995]
- Weissenbacher A, Kasess C, Gerstl F, Lanzenberger R, Moser E, Windischberger C. Correlations and anticorrelations in resting-state functional connectivity MRI: a quantitative comparison of preprocessing strategies. *Neuro Image*. 2009; 47(4):1408–1416. [PubMed: 19442749]
- Xue W, Bowman FD, Pileggi AV, Mayer AR. A multimodal approach for determining brain networks by jointly modeling functional and structural connectivity. *Frontiers in Computational Neuroscience*. 2015:9. [PubMed: 25698965]
- Zatorre RJ. Predispositions and plasticity in music and speech learning: neural correlates and implications. *Science*. 2013; 342(6158):585–589. [PubMed: 24179219]
- Zalesky A. Whole-brain anatomical networks: does the choice of nodes matter? *Neuro Image*. 2010; 50(3):970–983. [PubMed: 20035887]
- Zalesky A, Fornito A, Bullmore E. On the use of correlation as a measure of network connectivity. *Neuro Image*. 2012; 60(4):2096–2106. [PubMed: 22343126]
- Zalesky A, Cocchi L, Fornito A, Murray MM, Bullmore E. Connectivity differences in brain networks. *Neuro Image*. 2012; 60(2):1055–1062. [PubMed: 22273567]
- Zeng LL, Shen H, Liu L, Wang L, Li B, Fang P, et al. Identifying major depression using whole-brain functional connectivity: a multivariate pattern analysis. *Brain*. 2012; 59:1498–1507. [PubMed: 22418737]

Figure 1. Sagittal view of the brain with selected regions of interest (ROI). Region sizes are given as the number of voxels in 4mm, space and the ROI centers are given in Talarach space. This figure appears in color in the electronic version of this article.

Figure 2.

Histograms of voxel pair functional connections (correlation via Fisher's transformation) for two example subjects. The histograms demonstrate asymmetry due to the highly connected component in the right tail (the two components demonstrated in the figure are estimated by Gaussian mixture model). The two modes center at -0.1 and 0.8 .

Figure 3.

Diagram of the Bayesian hierarchical model. From the bottom to top are voxel pair data (observed and augmented), region level parameters, and hyperpriors.

Figure 4.

Voxel pair connectivity inferences: (a) the heatmap intensity reflects the number of functionally connected voxels in Cg25 to each voxel (row) of brain region mF10 across subjects (columns). (b) shows the number of functionally connected voxels in mF10 to each voxel (row) in Cg25 across 32 Subjects. The regional voxels (rows) in (a) and (b) are ordered such that the voxel with the most connections is listed first and the voxel with the least connections is listed last. Note that some voxels in mF10 and Cg25 consistently show a large number of connections to voxels in the other region across the 32 subjects. These highly connected voxels serve as neural processing hubs since they reveal coherent activity with many other voxels. (c) is a heatmap showing voxel pairs that are highly connected between mF10 and Cg25 and the intensity is the average posterior probability that a cross-region voxel pair is connected (averaged across subjects). (d) right is a 3D plot showing voxel pairs connected with high probabilities ($P_g(v) > 0.9$ for across subject average) falling in regions mF10 and Cg25. The squares and crosses represent the voxels in mF10 and Cg25 respectively. (d) left is enlarged and rotated for better visualization of links between voxels. This figure appears in color in the electronic version of this article.

Table 1

Patient group comparisons for breadth of connectivity between specified regions (entries represent number of connected voxel pairs.)

Parameters	Median (π_c %)	std	25 pct	75 pct	Mode	Difference(CI)	% Difference
Cg25 & mF10 (122,591 voxel pairs)							
(Mildly depressed $exp^{(h_0+h_1)}$)	36526 (29.8)	143	36466	36526	36600	2879	7.5 %
(Severely depressed $exp^{(h)}$)	39423(32.1)	155	39325	39422	39437	(2705,3084)	
Cg25 & latF9 (119,426 voxel pairs)							
(Mildly depressed $exp^{(h_0+h_1)}$)	8654 (7.2)	26	8334	8921	8688	488	6 %
(Severely depressed $exp^{(h)}$)	9124(7.6)	31	8968	9306	9132	(440,518)	
Cg25 & OF11 (214,376 voxel pairs)							
(Mildly depressed $exp^{(h_0+h_1)}$)	45632 (21.3)	318	45273	45987	45686	4729	10 %
(Severely depressed $exp^{(h)}$)	50361 (23.5)	387	50186	50626	50357	(4461,4891)	
Cg25 & Cg24 (32,705 voxel pairs)							
(Mildly depressed $exp^{(h_0+h_1)}$)	10162 (31.1)	71	10083	10239	10168	-149	-2 %
(Severely depressed $exp^{(h)}$)	10013 (30.5)	68	9972	10045	10018	(-120,-171)	

Summaries of distributions of posteriors in the simulation study with posterior median, standard deviation credible interval, and Ω_0 and Ω_1 represent the group average number of connected voxel pairs in cases and controls respectively.

Table 2

Parameters	Median	std	25 pct	75 pct	Mode
Ω_0	2785	12.38	2591	2889	2781
Ω_1	6778	15.32	6691	7153	6776
β_0 (s)	7.95	0.022	7.91	8.01	7.97
β_1 (0.9)	0.89	0.009	0.88	0.90	0.89