A Longitudinal Model for Functional Connectivity Networks Using Resting-State fMRI

Brian Hart[1], Ivor Cribben[2], Mark Fiecas[1], and for the Alzheimer's Disease Neuroimaging Initiative *[1]

[1]University of Minnesota
[2]University of Alberta

Abstract

Many neuroimaging studies collect functional magnetic resonance imaging (fMRI) data in a longitudinal manner. However, the current network modeling literature lacks a general framework for analyzing functional connectivity (FC) networks in fMRI data obtained from a longitudinal study. In this work, we build a novel longitudinal FC network model using a variance components approach. First, for all subjects' visits, we account for the autocorrelation inherent in the fMRI time series data using a non-parametric technique. Second, we use a generalized least squares approach to estimate 1) the within-subject variance component shared across the population, 2) the FC network, and 3) the FC network's longitudinal trend. Our novel method for longitudinal FC networks seeks to account for the within-subject dependence across multiple visits, the variability due to the subjects being sampled from a population, and the autocorrelation present in fMRI data, while restricting the number of parameters in order to make the method computationally feasible and stable. We develop a permutation testing procedure to draw valid inference on group differences in baseline FC and change in FC over time between a set of patients and a comparable set of controls. To examine performance, we run a series of simulations and apply the model to longitudinal fMRI data collected from the Alzheimer's Disease Neuroimaging Initiative database.

Keywords: fMRI; Functional connectivity; Longitudinal; Temporal autocorrelation.

1 Introduction

1.1 Resting-State fMRI

Resting-state functional magnetic resonance imaging (fMRI) captures a series of images of the brain in subjects who are not given a particular task to perform while in the scanner. The scanner repeatedly captures blood oxygenation level dependent (BOLD) signals at hundreds of thousands of locations within the brain, creating a time series of images of the brain. By capturing the BOLD signal of the resting brain, resting-state fMRI provides an opportunity for researchers to examine the functional connectivity (FC) network not tied to a particular task. We define FC as the temporal dependence, measured through cross-correlations, in the BOLD signals between brain regions [Friston et al., 1993]. Identifying group differences in FC can help better understand the underlying neurological process of a disease and its progression. Observed group differences can also potentially form biomarkers to be used for early detection and treatment of neurological disorders [Fox and Raichle, 2007].

1.2 Alzheimer's Disease Neuroimaging Initiative Data

A subset of the Alzheimer's Disease Neuroimaging Initiative (ADNI) data was used to demonstrate a practical application of our model. The data consists of longitudinal resting-state fMRI images collected at baseline, 3 months from baseline, 6 months from baseline, 12 months from baseline, and annually thereafter. There are two groups of interest, the cognitively normal (CN) group and the Alzheimer's Disease (AD) group. We focused our attention on late-onset AD and included only patients who were 65 years of age or older at baseline [van der Flier

*Data used in preparation of this article were obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: `http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf`

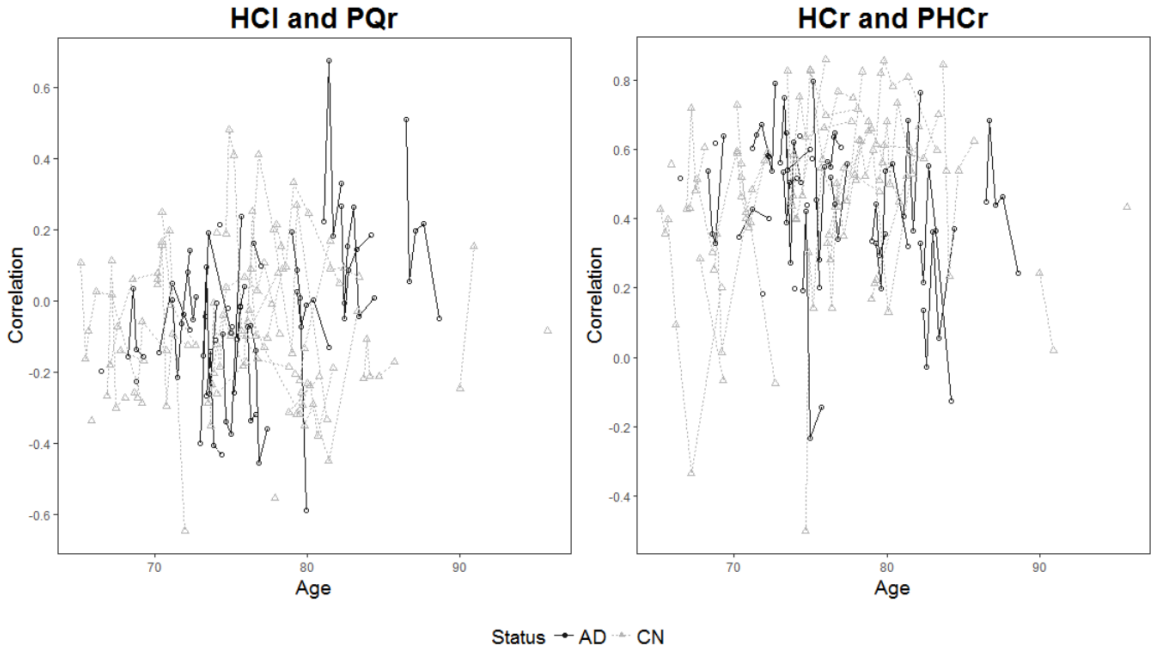

Figure 1: Spaghetti plots of the correlation between two ROIs against age. Each point represents a visit, and each line represents a subject. The ROIs represented in these plots are the left and right hippocampus (HCl and HCr), right precuneus (PQr), and right parahippocampus (PHCr).

et al., 2011, Holland et al., 2012]. To better separate the AD and CN groups, only patients who remained in one group for the entirety of the follow-up were considered in our analysis. The remaining CN group consists of 111 visits from 30 patients (17 females and 13 males) with each patient having between 1 and 6 visits. The AD group consists of 79 visits from 26 patients (11 females and 15 males) with each patient having between 1 and 5 visits. The average age was 75.9 for the CN group with a range of 65.2 to 95.7, while the AD group average age was 76.7 with a range of 66.5 to 88.6.

1.3 Clinical Motivations

Previous works have demonstrated the utility of FC network analysis. For example, past research has identified altered FC between healthy aging patients and those with AD. Even among CN individuals, FC demonstrates aging effects that are heterogeneous between different brain regions [Chen et al., 2016]. Chase [2014] and Hafkemeijer et al. [2012] showed altered FC patterns beyond healthy aging in patients with dementia and AD. Other works, including Ren et al. [2016] and Wang et al. [2007] have noted abnormal FC in various stages of AD. Wang et al. [2012] even demonstrated the impact of family history of AD on FC. In addition, Xiang et al. [2013] showed decreased FC from CN patients to mild cognitively impaired (MCI) patients to AD patients, and Li et al. [2015] found decreased network FC for CN patients who progressed to MCI over the following 24 months compared to non-progressers. These previous works, however, used *cross-sectional models*. Our work will build off of these results on the clinical utility of FC as a biomarker for AD.

1.4 Methodological Considerations

Of the studies mentioned in Section 1.3, only Ren et al. [2016] utilized truly longitudinal fMRI data. Aging effects are often measured by comparing young and ederly groups rather than following one group of subjects over time. A comprehensive longitudinal model that tests for differences in baseline and trend is needed to verify and expand on the previous results. Finn and Constable [2016] demonstrated that CN patients have distinct brain signatures in fMRI images, implying that separate scans from a single individual exhibit dependence. This finding can be leveraged in a longitudinal framework to better model aging effects. Our new methodology analyzes fMRI data at the region of interest (ROI) level. Thus, before fitting the model, the investigator must select a number of ROIs to include in the FC network analysis. We motivate the challenges of longitudinal FC network analysis with a preliminary examination of the ADNI data. Figure 1 shows spaghetti plots of the FC between the preprocessed

fMRI time series obtained from two ROI pairs for the AD and CN groups. The clustering of points within each line shows the within-subject dependence. In addition, there is considerable within-subject and within-group noise present in the estimates of FC. What is not evident from the figure is that the time series from which these correlations were obtained exhibit autocorrelation that contributes to the overall variability in FC. To add another level of complication, the figure depicts the marginal relationship between two ROIs, but to properly model the entire network we need a joint model that considers the network of all possible pairwise groupings of the chosen ROIs.

While methods such as Fiecas et al. [2017] exist for the cross-sectional case, no modeling framework exists for longitudinal FC network analysis. We fill this gap in the literature by proposing a novel longitudinal model and inference procedure that considers the network of all possible pairwise groupings of the chosen ROIs in resting-state fMRI data. Our longitudinal variance components FC network model accounts for the within-subject dependence across multiple visits, the variability due to the subjects being sampled from a population, and any autocorrelation present in fMRI time series. We also propose an efficient permutation based inference procedure that allows for valid hypothesis testing of group differences in baseline FC networks and FC network aging effects.

The remainder of the paper is laid out as follows. Section 2 formally introduces the model, including the estimation and inference procedures. Section 3 explains the design of the simulation study and discusses the results. Section 4 more fully describes the ADNI data, specifies the models fit to the data, and interprets the results of the FC analysis. We close with a discussion of future work and a conclusion in Section 5. R code for the methods proposed in this paper may be found at `https://github.com/mfiecas/longitudinalFC`.

2 Materials and Methods

2.1 Model Specification

Suppose we have a cohort of N individuals and let P denote the number of ROIs selected for a FC network analysis. We collect a P-variate, fMRI time series of length T from the preprocessed fMRI images of each of the N subjects at each visit. Let the subscripts i and j denote subject and visit, respectively. A particular subject returns for J_i total visits, and the cohort has a total of $J = \sum_{i=1}^{N} J_i$ visits. Let Y_i represent the vector of sample correlation coefficients for subject i of length QJ_i, where $Q = P(P-1)/2$ is the number of ROI pairs. Within Y_i, the Q correlations from the first visit, Y_{i1}, are followed by the Q correlations from the second visit, Y_{i2}, and so on until the J_i-th visit. The full response vector Y is formed by stacking the N different Y_i vectors. Our longitudinal model for the FC network is a linear model with baseline effect β_0 and longitudinal trend β_1, where each of these model parameters are vectors of length Q. We denote the time at visit j for subject i as v_{ij}. The vector v_i is formed by stacking the J_i distinct $v_{ij} \otimes \mathbb{1}_Q$ vectors for subject i, where $\mathbb{1}_Q$ is a vector of ones of length Q. Likewise v is formed by stacking the N distinct v_i vectors. Depending on the nature of the data and the research questions at hand, v_{ij} can be set to the visit number, the time since baseline, or the patient's age. Then, denoting element-wise multiplication with $*$, our model has the following linear form:

$$Y = \mathbb{1}_J \otimes \beta_0 + v * (\mathbb{1}_J \otimes \beta_1) + \varepsilon, \text{ where Var } (\varepsilon) = \Sigma + \Psi. \tag{1}$$

The key element in our longitudinal linear model is the variance structure of the error term. We separate the error variance into two components, Σ and Ψ. Σ accounts for the within-visit variance and the temporal autocorrelation in the fMRI time series, and Ψ accounts for the variability and covariability arising from the heterogeneity across subjects and the within-subject covariation coming from the longitudinal design. Σ is block diagonal where each $Q \times Q$ block, Σ_{ij}, accounts for the within-visit variance present in visit j for subject i for the Q pairs of ROIs. Ψ is also block diagonal with a $QJ_i \times QJ_i$ block for participant i. These diagonal blocks do not differ between subjects except through their dimensions which depends on the number of visits for each subject. Let Ψ_{diag} be an arbitrary diagonal block of Ψ. We then further break Ψ_{diag} into two components, Ψ_0 and Ψ_1. Ψ_0 is a $Q \times Q$ block that is repeated along the diagonal of each Ψ_{diag}. This term models the within-visit covariability not captured by Σ. Ψ_1 is a $Q \times Q$ block that populates the off diagonal blocks of Ψ_{diag}, modeling the within-subject, across-visit covariability coming from the longitudinal design.

We write Equation 1 in the form of a linear model with a vector response, allowing us to use existing methods for estimating the parameters and for statistical inference. To this end, our model can also be written in the standard linear model form with a design matrix X. Let $X_{ij} = [1 \quad v_{ij}] \otimes I_Q$, where I_Q is the $Q \times Q$ identity matrix. To form X_i, the portion of the design matrix specific to subject i, we stack the J_i individual X_{ij}. Likewise, to form X we stack the N individual X_i. If we define β as a vector of length $2Q$ where the first Q elements are β_0 and the last Q elements are β_1, then Equation 1 can be written as $Y = X\beta + \varepsilon$.

Figure 2 shows a workflow chart of the procedure used to estimate the model parameters and test hypotheses which are subsequently described in more detail.

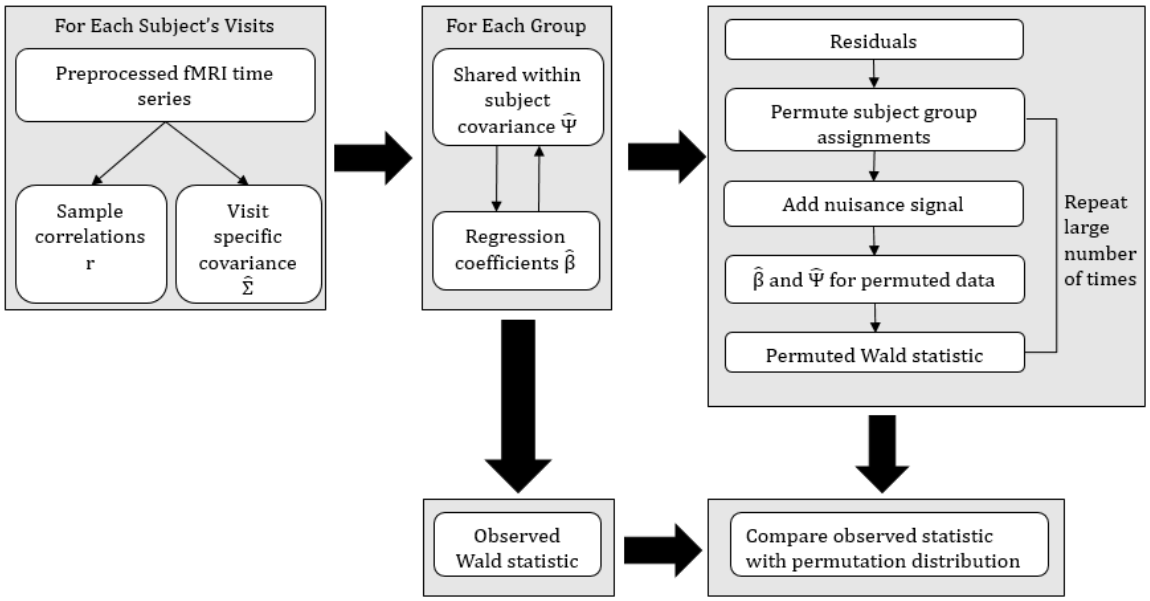

Figure 2: A workflow chart of the estimation and inferential procedure of our variance components model.

2.2 Estimating Within Visit Covariance

We start by estimating the sample correlation coefficient between all ROI pairs for all visits and their corresponding variances and covariances. Let $(w_{1t}, \ldots, w_{Pt})'$ for $t \in \{1, 2, \ldots, T\}$ be the time series of preprocessed BOLD signals from P ROIs for a single visit. Here the first subscript on w indicates the ROI and the second subscript indicates the time within the time series. Then for the p-th and q-th ROIs, where $p, q \in \{1, 2, \ldots, P\}$, we have $r_{pq} = \sum_{t=1}^{T}(w_{pt} - \bar{w}_p)(w_{qt} - \bar{w}_q)/\sqrt{\sum_{t=1}^{T}(w_{pt} - \bar{w}_p)^2 \sum_{t=1}^{T}(w_{qt} - \bar{w}_q)^2}$. We now seek a consistent estimator of the covariance between correlations of two ROI pairs, (p, q) and (p', q'). The estimator of Melard et al. [1991] and Roy [1989] offers a nice solution which accounts for the autocorrelation present in fMRI time series and is consistent under mild conditions. First, define $\hat{\gamma}_{pq}(u) = \sum_{t=1}^{T-u}(w_{pt} - \bar{w}_p)(w_{q,t+u} - \bar{w}_q)/T$. Then, letting $h(\cdot)$ be the modified Bartlett window with bandwidth $b(T)$, we set $\hat{\Theta}(p, q, p', q') = \sum_{u=-T+1}^{T-1} h^2(u)\hat{\gamma}_{pq}(u)\hat{\gamma}_{p'q'}(u)$. Using $\hat{\gamma}$ and $\hat{\Theta}$, we then let $\hat{\Delta}(p, q, p', q') = \hat{\Theta}(p, q, p', q')/\sqrt{\hat{\gamma}_{pp}(0)\hat{\gamma}_{qq}(0)\hat{\gamma}_{p'p'}(0)\hat{\gamma}_{q'q'}(0)}$. Finally, our estimator takes the following form:

$$\begin{aligned}\widehat{Cov}(r_{pq}, r_{p'q'}) = [&0.5 r_{pq} r_{p'q'}\{\hat{\Delta}(p, p', p, p') + \hat{\Delta}(p, q', p, q') + \hat{\Delta}(q, p', q, p') + \hat{\Delta}(q, q', q, q')\} \\ &- r_{pq}\{\hat{\Delta}(p, p', p, q') + \hat{\Delta}(q, p', q, q')\} - r_{p'q'}\{\hat{\Delta}(q, p', p, p') + \hat{\Delta}(q, q', p, q')\} \\ &+ \hat{\Delta}(p, p', q, q') + \hat{\Delta}(q, p', p, q')]/T. \end{aligned} \tag{2}$$

We use Equation 2 to populate the cells of the each Σ_{ij} block to get our estimates $\hat{\Sigma}_{ij}$.

2.3 Estimating Between Subject Covariance, Ψ, and β

Utilizing a generalized least squares (GLS) approach, we now proceed with the estimation of the between subject covariance Ψ and the regression coefficients β, conditional on the previously estimated within-visit covariances, $\hat{\Sigma}_{ij}$.

Although the framework allows for many different structures for Ψ, we assume a block compound symmetry structure so that all diagonal blocks, Ψ_0, are equal and all off diagonal blocks, Ψ_1, are equal. The block compound symmetry assumption keeps the parameter space to a reasonable size, but note that one could easily consider other forms of Ψ, such as an autoregressive structure, with minimal modification to the estimation procedure. We use the ordinary least squares estimator $\hat{\beta} = (X'X)^{-1}X'Y$ to provide a good starting estimate of β. We then

update the two components of Ψ using the following formulas:

$$\hat{\Psi}_0 = \frac{1}{\sum_{i=1}^{N} J_i}\left\{\sum_{i=1}^{N}\sum_{j=1}^{J_i}(Y_{ij} - X_{ij}\hat{\beta})'(Y_{ij} - X_{ij}\hat{\beta}) - \hat{\Sigma}_{ij}\right\}, \quad \text{and} \tag{3}$$

$$\hat{\Psi}_1 = \frac{1}{\sum_{i=1}^{N} \frac{J_i(J_i-1)}{2}}\sum_{i=1}^{N}\sum_{j\neq k}(Y_{ij} - X_{ij}\hat{\beta})'(Y_{ik} - X_{ik}\hat{\beta}). \tag{4}$$

Note how these estimators resemble the mean squared error. For $\hat{\Psi}_0$, the sum of the squared residuals for all visits are summed. We then subtract the previously estimated $\hat{\Sigma}_{ij}$ and divide the remaining covariance matrix by the total number of visits to get $\hat{\Psi}_0$. We estimate Ψ_1 in a similar fashion. In this case, we calculate and average the cross products of the residuals from all pairs of visits for each subject. We do not subtract a $\hat{\Sigma}$ term here because $\hat{\Sigma}$ has been set to 0 for the off diagonal blocks occupied by Ψ_1. These simple estimators, obtained through matching moments, significantly reduce computation in comparison to a maximum likelihood approach. To increase model parsimony, different structures can be considered for Σ_{ij}, Ψ_0, and Ψ_1. For example, to enforce a diagonal structure, set all off-diagonal elements to 0, or to enforce a compound symmetry structure, set all diagonal elements to the average of the diagonal elements and likewise for the off-diagonal elements [Laird, 2004]. Changing the form of Σ_{ij}, Ψ_0, and Ψ_1 allows you to fit the model with flexible variance assumptions as is often done in traditional generalized least squares linear models.

With an estimate of Ψ obtained using Equations 3 and 4, we can now use the standard GLS formula to update the regression coefficients as follows: $\hat{\beta} = \{X'(\hat{\Sigma} + \hat{\Psi})^{-1}X\}^{-1}X'(\hat{\Sigma} + \hat{\Psi})^{-1}Y$. At this point we have two choices: iteratively update $\hat{\Psi}$ and $\hat{\beta}$ until convergence (full convergence), or accept the estimates (one-step) and proceed with the inferential procedure. The one-step estimator provides a significant advantage in computing time as $\hat{\Psi}$ and $\hat{\beta}$ must be estimated for each permutation of the inference procedure [Ganjgahi et al., 2015]. Additionally, Amemiya [1977] proved that the one-step GLS estimator maintains consistency.

We estimate $\hat{\beta}, \hat{\Psi}$, and $\hat{\Sigma}$ for each group (CN and AD for ADNI) separately using this estimation procedure. Superscripts on the parameter estimates denote the group (e.g. $\hat{\beta}^{G_1}, \hat{\Psi}^{G_1}$, and $\hat{\Sigma}^{G_1}$ are the estimates for group 1).

2.4 Inference

Two general hypothesis tests are of interest in a longitudinal FC model: the group difference in the baseline FC, and the group difference in the longitudinal trend in FC. For each hypothesis, we would like to test for the significance of the group difference in both the global FC network and the local ROI pair FC. We refer to the vector wide test of a difference in the parameter vector β_0 or β_1 as a global test and refer to a test of a group difference in a single element of β_0 or β_1 as a local test. To accomplish our hypothesis testing objectives, we use the Wald statistic, $\{C(\hat{\beta}^{G1} - \hat{\beta}^{G2})\}'[C\{\widehat{\text{Var}}(\hat{\beta}^{G1}) + \widehat{\text{Var}}(\hat{\beta}^{G2})\}C']^{-1}C(\hat{\beta}^{G1} - \hat{\beta}^{G2})$, and adjust the contrast matrix, C, depending on the hypothesis of interest. For instance, to test for a global difference in β_0, we replace all β terms with the β_0 vector for the proper group and set the contrast matrix, C, to the $Q \times Q$ identity matrix.

We estimate the variance of each group's regression coefficients using $\widehat{\text{Var}}(\hat{\beta}) = \{X'(\hat{\Sigma} + \hat{\Psi})^{-1}X\}^{-1}$. We now detail a permutation procedure for all inference.

1. Calculate residuals from the fitted model for each subject: $e_i = Y_i - X_i\hat{\beta}^G$ for subject i in group G.
2. Permute group assignments of e_i.
3. Add the nuisance signal back to e_i based on new permuted group assignments G^*. For the main effect (intercept) tests we add in the longitudinal trends by setting $e^*_{ij} = e_{ij} + v_{ij}\hat{\beta}_1^{G^*}$. Likewise, for the interaction (slope) tests we set $e^*_{ij} = e_{ij} + \hat{\beta}_0^{G^*}$.
4. Refit the model on e^*, the permuted, adjusted, and stacked residuals from step 3.
5. Calculate a new Wald statistic for the fitted values of $\hat{\beta}^{G^*}$ and $\hat{\Psi}^{G^*}$.

We repeat steps 2 through 5 a large number of times to create a permutation distribution to be used as a reference distribution of the originally calculated test statistic. The number of permutations determines the precision of the p-value and should be chosen to be large enough to offer sufficient precision after any multiple comparisons adjustment. In our particular data example we chose to run 5,000 permutations. Because the obtained p-values are estimated discrete values, we additionally use the permutation p-value correction procedure of Phipson et al. [2010]. To account for the fact that $2Q$ local hypotheses are tested simultaneously, we then apply the false discovery rate (FDR) controlling procedure of Benjamini and Hochberg [1995] to the corrected p-values from the local tests. The Phipson et al. [2010] correction helps avoid unadjusted p-values with value 0 which may improperly maintain significance after a multiple comparisons correction.

Chung and Romano [2013] showed that studentized test statistics, such as the proposed Wald statistic, allow for valid inference in many permutation test settings. A recent comparison of the performance of different permutation strategies by Winkler et al. [2014] showed that the Ter Braak permutation testing procedure we use maintains nominal Type I error and is fairly robust [Ter Braak, 1992]. This method offers the additional advantage that the data only needs to be permuted once and the model only fit twice at each iteration of the permutation test to test all local and global hypotheses. Testing all hypotheses under a single permutation schedule greatly reduces the computational burden of the testing procedure.

3 Simulation Study

3.1 Simulation Setup

A series of simulations were designed with different data generating mechanisms to assess model performance. In all scenarios each time series contained 120 time points and had an autocorrelation structure that followed a first-order autoregressive process with an AR parameter of 0.3. A multivariate time series was simulated for each subject at three visits. For each visit, the Q correlations were simulated from a multivariate normal distribution where the mean and variance varied by group based on the simulation setting. For group 1, the mean vector was always assumed to be 0 and the covariance matrix was the same across all simulation settings. The simulations used P of either 3, 5, or 10 as the dimension of the multivariate normal distribution. For the 3 and 5 dimension settings only the first element of the group 2 mean vector was allowed to vary by simulation setting, while the other elements were set to match group 1. For the 10 dimension settings the first 5 elements of the group 2 mean vector varied by simulation setting and the other elements were again set to match group 1. 1,000 Monte Carlo simulations were run for all simulation settings with 3 and 5 dimensions, and 500 simulations were run for the 10 dimensional simulation settings. Group sizes of 15 and 30 were considered. The true variance of the correlations was either equal for the two groups or the group 2 variance was double the group 1 variance. 500 permutations were used for the permutation test for all settings. The group size, number of visits, and time series length were selected to reflect values found in typical fMRI studies. The effect size of 0.1 is realistic for fMRI longitudinal scenarios. Considering a smaller effect size would lead to very similar conclusions with slight decreases in power across all models considered.

As mentioned earlier, there are no other competitor models which consider FC networks longitudinally. Because of this, we have chosen to fit three versions of our model with different variance assumptions and estimation methods so they can be compared to each other. The first model considered was a full convergence model which iterated between $\hat{\Psi}$ and $\hat{\beta}$ until convergence. It assumed an unstructured Σ_{ij} and compound symmetry for Ψ_0 and Ψ_1. This model is referred to as the full convergence full variance model. The second was a one-step model which stops after one iteration of solving for $\hat{\Psi}$ and $\hat{\beta}$. It also assumed an unstructured Σ_{ij} and compound symmetry for Ψ_0 and Ψ_1 and is referred to as the one-step convergence full variance model. The final model was a one-step model which assumed a diagonal structure for Σ_{ij} and scaled identity structures for Ψ_0 and Ψ_1. This last model is referred to as the one-step convergence reduced variance model.

3.2 Simulation Results

Tables 1 and 2 show the simulation study results. Table 1 shows the global and local Type I error for the main effect and interaction across all simulations. The reported global test results are the average global Type I errors across 500 or 1,000 Monte Carlo runs. The local test results are the average Type I errors of the unadjusted p-values for all null hypotheses across the 500 or 1,000 Monte Carlo runs. While the local p-values would be adjusted in practice, the numbers in the table provide easy reference to a nominal Type I error of 0.05. Table 2 shows the average global power and average local power using false discovery rate adjusted p-values. All p-values were corrected using the method of Phipson et al. [2010]. Additional simulation study results can be found in Web Table 1 of web-based supplementary materials. Table 1 shows roughly nominal Type I error rates for all three model specification. While there was some slight inflation in all three models, especially for the 10 ROI simulations, the inflation was attenuated by the increase in sample size from 15 to 30 per group. Table 2 demonstrates adequate power, both locally and globally for all three models. As expected, power increased with larger group size and decreased with a larger true group 2 variance.

There was no consistent difference in performance between the three models in terms of power or Type I error across the different simulation settings. The two full variance settings match the true model of the simulated data, yet the reduced variance model did not suffer in comparison. The reduced variance model may have offered similar performance because the smaller parameter space allowed for improved estimation. The reduced variance model did not capture the full true variance, but it still performed well by allowing the FC for each ROI pair to be correlated across multiple visits for a given subject.

Table 1: Type I error rates for simulation study for all globally null simulation settings. Type I errors for the main effect (group difference in intercepts) and interaction effect (group difference in longitudinal slopes) are reported both globally and locally. The global Type I errors are averaged across all simulations. The local Type I errors reported are unadjusted and averaged across all simulations and all null ROI pairs.

		Convergence Type:			**Full**		**One-Step**			**One-Step**	
		Variance Structure:			**Full**		**Full**			**Reduced**	
	Group Size	**Variance**	β_0	β_1	**3 ROIs**	**5 ROIs**	**3 ROIs**	**5 ROIs**	**10 ROIs**	**3 ROIs**	**5 ROIs**
Main Effect Global Test	**15**	**Equal**	**0**	**0**	0.061	0.052	0.057	0.055	0.066	0.056	0.054
			0	**0.1**	0.065	0.065	0.067	0.071	0.064	0.067	0.063
		Group 2	**0**	**0**	0.058	0.049	0.055	0.053	0.054	0.056	0.048
		Double	**0**	**0.1**	0.068	0.058	0.069	0.056	0.068	0.070	0.060
	30	**Equal**	**0**	**0**	0.045	0.049	0.046	0.050	0.066	0.043	0.047
			0	**0.1**	0.046	0.050	0.054	0.047	0.066	0.048	0.049
		Group 2	**0**	**0**	0.057	0.054	0.058	0.053	0.062	0.058	0.060
		Double	**0**	**0.1**	0.049	0.066	0.055	0.059	0.078	0.058	0.053
Main Effect Local Test	**15**	**Equal**	**0**	**0**	0.071	0.061	0.073	0.061	0.066	0.066	0.062
			0	**0.1**	0.064	0.067	0.067	0.069	0.069	0.064	0.067
		Group 2	**0**	**0**	0.064	0.063	0.061	0.062	0.071	0.064	0.064
		Double	**0**	**0.1**	0.068	0.062	0.069	0.063	0.069	0.065	0.063
	30	**Equal**	**0**	**0**	0.050	0.059	0.048	0.058	0.058	0.048	0.057
			0	**0.1**	0.053	0.058	0.051	0.058	0.056	0.052	0.057
		Group 2	**0**	**0**	0.063	0.055	0.062	0.055	0.057	0.060	0.056
		Double	**0**	**0.1**	0.054	0.059	0.055	0.058	0.059	0.053	0.057
Interaction Global Test	**15**	**Equal**	**0**	**0**	0.062	0.054	0.052	0.057	0.072	0.049	0.052
			0.1	**0**	0.046	0.054	0.044	0.049	0.076	0.044	0.049
		Group 2	**0**	**0**	0.057	0.067	0.050	0.062	0.070	0.051	0.069
		Double	**0.1**	**0**	0.079	0.071	0.078	0.061	0.090	0.073	0.063
	30	**Equal**	**0**	**0**	0.054	0.059	0.058	0.061	0.078	0.056	0.058
			0.1	**0**	0.052	0.067	0.049	0.063	0.070	0.051	0.051
		Group 2	**0**	**0**	0.068	0.065	0.071	0.061	0.072	0.075	0.057
		Double	**0.1**	**0**	0.056	0.041	0.057	0.049	0.068	0.059	0.040
Interaction Local Test	**15**	**Equal**	**0**	**0**	0.059	0.062	0.058	0.062	0.067	0.059	0.064
			0.1	**0**	0.055	0.060	0.054	0.057	0.071	0.057	0.060
		Group 2	**0**	**0**	0.063	0.065	0.065	0.061	0.072	0.064	0.063
		Double	**0.1**	**0**	0.074	0.067	0.071	0.061	0.067	0.074	0.064
	30	**Equal**	**0**	**0**	0.050	0.055	0.049	0.054	0.059	0.049	0.055
			0.1	**0**	0.054	0.059	0.053	0.058	0.056	0.058	0.059
		Group 2	**0**	**0**	0.061	0.059	0.062	0.057	0.060	0.063	0.060
		Double	**0.1**	**0**	0.057	0.055	0.057	0.055	0.058	0.058	0.055

Table 2: The power calculations for the simulation study. Power results for the main effect (group difference in intercepts) and interaction effect (group difference in longitudinal slopes) are reported both globally and locally. The global power results are averaged across all simulations. The local power results reported are unadjusted and averaged across all simulations and all non-null ROI pairs.

		Convergence Type:			**Full**		**One-Step**			**One-Step**	
		Variance Structure:			**Full**		**Full**			**Reduced**	
	Group Size	**Variance**	β_0	β_1	**3 ROIs**	**5 ROIs**	**3 ROIs**	**5 ROIs**	**10 ROIs**	**3 ROIs**	**5 ROIs**
Main Effect Global Test	**15**	**Equal**	**0.1**	**0**	0.382	0.223	0.369	0.229	0.482	0.372	0.211
			0.1	**0.1**	0.389	0.214	0.390	0.203	0.466	0.391	0.204
		Group 2	**0.1**	**0**	0.300	0.196	0.294	0.185	0.356	0.303	0.185
		Double	**0.1**	**0.1**	0.328	0.180	0.321	0.170	0.386	0.316	0.177
	30	**Equal**	**0.1**	**0**	0.674	0.491	0.679	0.489	0.902	0.682	0.469
			0.1	**0.1**	0.684	0.468	0.686	0.470	0.884	0.670	0.456
		Group 2	**0.1**	**0**	0.582	0.383	0.581	0.385	0.836	0.579	0.367
		Double	**0.1**	**0.1**	0.582	0.383	0.580	0.380	0.800	0.572	0.371
Main Effect Local Test	**15**	**Equal**	**0.1**	**0**	0.412	0.413	0.413	0.379	0.270	0.427	0.458
			0.1	**0.1**	0.424	0.406	0.417	0.372	0.233	0.446	0.410
		Group 2	**0.1**	**0**	0.294	0.293	0.296	0.267	0.206	0.292	0.307
		Double	**0.1**	**0.1**	0.326	0.304	0.291	0.253	0.160	0.330	0.324
	30	**Equal**	**0.1**	**0**	0.807	0.822	0.810	0.815	0.778	0.811	0.813
			0.1	**0.1**	0.817	0.800	0.818	0.790	0.759	0.816	0.786
		Group 2	**0.1**	**0**	0.689	0.696	0.693	0.701	0.670	0.687	0.697
		Double	**0.1**	**0.1**	0.707	0.691	0.708	0.693	0.638	0.704	0.694
Interaction Global Test	**15**	**Equal**	**0**	**0.1**	0.757	0.525	0.765	0.492	0.742	0.771	0.550
			0.1	**0.1**	0.767	0.537	0.753	0.494	0.728	0.783	0.557
		Group 2	**0**	**0.1**	0.706	0.502	0.694	0.444	0.726	0.719	0.515
		Double	**0.1**	**0.1**	0.674	0.460	0.679	0.422	0.668	0.667	0.483
	30	**Equal**	**0**	**0.1**	0.983	0.881	0.983	0.877	0.942	0.986	0.888
			0.1	**0.1**	0.977	0.899	0.980	0.894	0.936	0.977	0.906
		Group 2	**0**	**0.1**	0.959	0.842	0.958	0.831	0.922	0.954	0.838
		Double	**0.1**	**0.1**	0.950	0.839	0.952	0.821	0.890	0.945	0.828
Interaction Local Test	**15**	**Equal**	**0**	**0.1**	0.866	0.885	0.868	0.853	0.728	0.868	0.875
			0.1	**0.1**	0.876	0.861	0.872	0.850	0.740	0.876	0.883
		Group 2	**0**	**0.1**	0.843	0.831	0.834	0.800	0.664	0.850	0.843
		Double	**0.1**	**0.1**	0.808	0.793	0.811	0.761	0.612	0.818	0.805
	30	**Equal**	**0**	**0.1**	0.997	0.997	0.997	0.996	0.982	0.998	0.994
			0.1	**0.1**	0.992	0.993	0.993	0.991	0.983	0.994	0.991
		Group 2	**0**	**0.1**	0.991	0.987	0.991	0.985	0.954	0.990	0.986
		Double	**0.1**	**0.1**	0.989	0.984	0.988	0.982	0.947	0.989	0.986

Overall, the full convergence results did not consistently improve on the one-step estimator results. With negligible gains in terms of power and Type I error, the additional computational resources used by the full-convergence model did not offer any practical advantage. A primary difference in the three models was the computational time. The full convergence model took, on average, over 2.5 times longer to run without seeing any boost in performance. For reference, using a 3.7 GHz Quad-Core Intel Xeon with 16GB ram, the average times to fit the one-step full variance model with 30 subjects per group for 3, 5, and 10 ROIs were 3.8 seconds, 54.9 seconds, and 87.2 minutes, respectively. These results show that the time increases quickly with the dimension of the model. The computational time is largely driven by the permutation procedure. Thus, if a larger number of permutations is desired for the testing procedure, then the computational time will see a corresponding increase.

4 ADNI Data Analysis

4.1 ADNI Preprocessing

Data used in the preparation of this article were obtained from the ADNI database (`http://adni.loni.usc.edu`). ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging, positron emission tomography, other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of MCI and early AD. For up-to-date information, see `www.adni-info.org`.

We preprocessed the ADNI data using both FSL (version 5.0.9, `https://fsl.fmrib.ox.ac.uk/`) and AFNI (version AFNI_17.0.15, `https://afni.nimh.nih.gov/`). The preprocessing steps were as follows. We 1) applied motion correction to the images using FSL's mcflirt (rigid body transform; cost function normalized correlation; reference volume the middle volume) and then 2) normalized the images into the Montreal Neurological Institute space using FSL's flirt (affine transform; cost function correlation ratio). We used FSL's fast to 3) obtain a probabilistic segmentation of the brain to obtain white matter and cerebrospinal fluid (CSF) probabilistic maps, thresholded at 0.75. Using FSL's fslmaths, we 4) spatially smoothed the volumes using a Gaussian kernel with FWHM=5 mm. We used AFNI's 3dDetrend to 5) remove nuisance signals, namely the six motion parameters, white matter and CSF signals, and the global signal. Finally, 6) the linear trend was removed from each time series using linear regression and a 4th order Butterworth low-pass filter with a 0.1 Hertz cutoff was applied to each fMRI time series.

We used the Automated Anatomical Label (AAL) atlas to subdivide the brain into 116 anatomical regions [Tzourio-Mazoyer et al., 2002]. We define the ROI level time series for a given region as the average of the time series from each voxel (3D location) within that region of the brain. We then selected $P = 10$ ROIs for analysis of the ADNI data based on previous literature which has shown differences in FC between AD and CN patients in the default mode network (DMN) and hippocampi [Supekar et al., 2008, Greicius et al., 2004, Sorg et al., 2007]. The ten regions we selected were the left and right hippocampus (HC), parahippocampus (PHC), posterior cingulate (PCC), precuneus (PQ), and prefrontal cortex (PFC). In all results that follow an l suffix for an ROI denotes the left side of the brain and an r suffix denotes the right side. Because a full brain analysis of all 116 regions is not currently feasible using our method, investigators should select a set of ROIs for their particular dataset and research question based on expert knowledge and literature review.

4.2 ADNI Analysis

Four models were fit to the ADNI data with differing assumptions. Model 1 is a one-step estimation model which assumes compound symmetry structure for Ψ_0 and Ψ_1 and unstructured Σ_{ij}. Model 2, makes the same assumptions for Ψ_0, Ψ_1, and Σ_{ij} but uses the full convergence estimator. Model 3 is a one-step estimation model assuming scaled identity structures for Ψ_0 and Ψ_1 and a diagonal structure for Σ_{ij}. Finally, Model 4 uses one-step estimation, assumes a diagonal structure for Ψ_0, sets all elements of Ψ_1 to 0, and assumes a diagonal structure for Σ_{ij}. This final model is equivalent to a massive univariate approach which ignores the within-subject dependence. 5,000 permutations were run for all models. The intercept of each model represents the FC network of each group at age 65.

4.3 ADNI Results

Table 3 shows results from the global hypothesis tests and all local hypothesis tests that were significant before p-value adjustment for all four models. Neither the overall main effect or interaction term were found to significantly differ between groups in the global tests in any of the four models considered. The only ROI pair level difference that remained significant after p-value adjustment and correction in any of the models was the difference in the CN and AD group longitudinal slopes in the FC between the left HC and the right PCC in Models 1 and 2. These models conclude that the FC between HCl and PCCr declines at a significantly quicker rate in the AD

population than in their CN counterparts. The estimated Model 1 group intercepts, group longitudinal trends, group differences in intercepts and longitudinal trends, and $-\log_{10}$ p-values after correction and adjustment from local hypothesis tests are presented graphically in Figure 3. Similar figures for Models 2, 3, and 4 can be found in Web Figures 1, 2, and 3, respectively, of the web-based supplementary materials. The four models fit to the ADNI

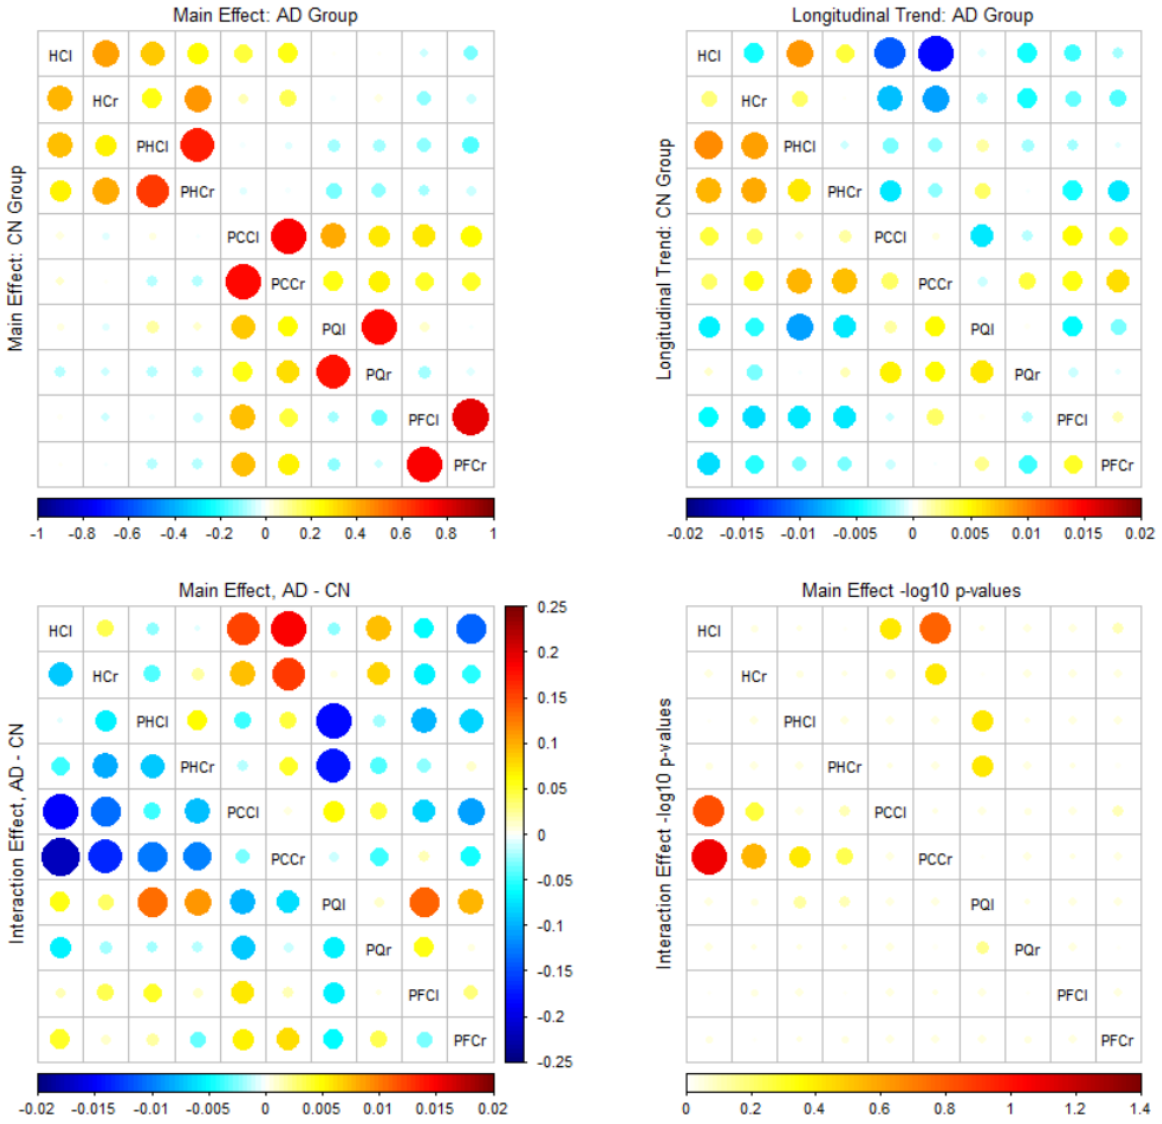

Figure 3: Model 1 results. Top left: A plot of the estimated intercept terms for the CN group (bottom left triangle) and AD group (top right triangle). Top right: A plot of the estimated slope terms for the CN group (bottom left triangle) and AD group (top right triangle). Bottom left: a plot of the group differences (AD estimates - CN estimates) for the estimated intercepts (top right triangle) and slopes (bottom left triangle). Bottom right: A plot of the $-\log_{10}$ corrected and adjusted p-values from all local hypothesis tests of group differences (AD estimates - CN estimates) for the estimated intercepts (top right triangle) and slopes (bottom left triangle).

data present slightly different results. The difference between Model 1 and Model 2 is very minimal. The nearly identical results provide further support of the conclusion from the simulation study that the full convergence and one-step estimator models lead to very similar estimates and inference. Some more pronounced differences in results arise when Model 1 is compared with Models 3 and 4. These differences make sense considering that Model 4 does not account for the within-subject dependence of the FC and thus appears to suffer from diminished power to detect group differences.

The more interesting differences exist in comparing Models 1 and 3, but some common patterns run throughout both sets of results. In both models, many of the local hypotheses that were significant prior to the FDR correction

Table 3: Hypothesis tests on the ADNI data. Global tests and all local tests with unadjusted p-values of < 0.05 are shown for Models 1-4. The numbers in italics are from the global hypothesis tests.

	β_{CN}	β_{AD}	Test Statistic	Unadjusted p-value	Adjusted p-value
Model 1: One-step, Compound Symmetry Ψ_0 and Ψ_1, and Unstructured Σ					
Main Effects			*42.86*	*0.154*	
HCl and PCCl	0.037	0.189	3.39	0.024	0.315
HCl and PCCr	0.043	0.232	5.26	0.004	0.108
HCr and PCCr	0.011	0.168	3.63	0.018	0.315
PHCl and PQl	0.089	-0.093	4.83	0.030	0.315
PHCr and PQl	0.041	-0.136	4.59	0.031	0.315
Interactions			*41.85*	*0.197*	
HCl and PCCl	0.004	-0.011	5.85	0.002	0.090
HCl and PCCr	0.003	-0.014	7.82	0.001	0.045
HCr and PCCr	0.005	-0.009	4.49	0.009	0.209
PHCl and PCCr	0.008	-0.002	2.67	0.032	0.315
Model 2: Full Convergence, Compound Symmetry Ψ_0 and Ψ_1, and Unstructured Σ					
Main Effects			*42.06*	*0.156*	
HCl and PCCl	0.037	0.189	3.35	0.024	0.317
HCl and PCCr	0.043	0.232	5.14	0.004	0.111
HCr and PCCr	0.010	0.168	3.57	0.019	0.317
PHCl and PQl	0.089	-0.093	4.79	0.030	0.317
PHCr and PQl	0.042	-0.136	4.56	0.029	0.317
Interactions			*41.26*	*0.199*	
HCl and PCCl	0.004	-0.011	5.77	0.002	0.081
HCl and PCCr	0.003	-0.014	7.65	¡0.001	0.027
HCr and PCCr	0.005	-0.009	4.41	0.010	0.218
PHCl and PCCr	0.008	-0.002	2.61	0.032	0.317
Model 3: One-step, Scaled Identity Ψ_0 and Ψ_1, and Diagonal Σ					
Main Effects			*47.75*	*0.312*	
HCl and PCCl	0.058	0.228	3.95	0.021	0.367
HCl and PCCr	0.068	0.269	5.51	0.005	0.150
PHCl and PQl	0.125	-0.074	5.37	0.033	0.367
PHCr and PQl	0.056	-0.164	6.57	0.009	0.207
PCCl and PFCr	0.381	0.240	2.86	0.033	0.367
Interactions			*54.90*	*0.238*	
HCl and PCCl	0.002	-0.013	5.14	0.005	0.150
HCl and PCCr	0.000	-0.017	6.49	0.017	0.150
PHCl and PCCr	0.007	-0.003	2.75	0.037	0.367
PHCl and PQl	-0.011	0.003	4.71	0.049	0.403
PHCr and PQl	-0.007	0.006	3.66	0.047	0.403
PQl and PQr	0.007	0.000	1.31	0.035	0.367
Model 4: One-step, Scaled Identity Ψ_0, Zero Ψ_1, and Diagonal Σ					
Main Effects			*62.65*	*0.412*	
HCl and PCCr	0.078	0.281	8.41	0.013	0.513
PHCr and PQl	0.040	-0.162	6.61	0.023	0.513
PCCl and PFCr	0.402	0.223	4.35	0.018	0.513
Interactions			*68.63*	*0.317*	
HCl and PCCl	0.000	-0.013	5.60	0.036	0.592
HCl and PCCr	-0.001	-0.018	9.11	0.006	0.513
PHCl and PCCr	0.006	-0.004	3.83	0.040	0.592

appear between the HC/PHC and the PCC. These group differences strengthen the one local hypothesis that is significant after FDR correction from Models 1 and 2 which shows a significantly larger decrease in FC between the HCl and PCCr in the AD group than in the CN group. While the significant results disappear after FDR correction for Model 2, the fact that many other HC/PHC connections with the PCC see a similar pattern helps to indicate differing baseline and longitudinal trend effects in the FC of the two groups. This clustering of group differences can be seen in Figure 3 with the smallest p-values (red and orange circles) appearing between the HC/PHC and PCC. Wang et al. [2006], Sorg et al. [2007], and Greicius et al. [2004] all noted decreased FC between the HC and PCC in patients with AD in analyses of cross-sectional data. Similar results from Supekar et al. [2008] showed decreased clustering coefficients for the HC. Our analysis confirms these results with the addition of a longitudinal component to the analysis. Our results not only conclude that AD and CN patients have differing FC between the HC and PCC, as the previous works have shown, but we also more clearly describe the differences in baseline and longitudinal trend in FC between these two regions.

5 Discussion and Conclusions

We have introduced a novel variance components longitudinal model to estimate and draw inference on the group differences in FC using resting-state fMRI data. The model properly accounts for the correlation inherent in repeated measures data and the autocorrelation present in fMRI time series. A permutation testing procedure performs global and local two-sample testing in an efficient manner. The linear model framework and utilization of generalized least squares estimators offers great simplicity and a large number of natural extensions. This work fills a current gap in the literature by providing a general framework for estimation and hypothesis testing of longitudinal FC data.

As a practical example, we applied the method to resting-state fMRI data from the ADNI database. Our analysis found a faster decline in FC between the HCl and the PCCr in AD patients compared to the CN controls. This finding confirms the results of previous studies and helps solidify the central roles of the hippocampus and default mode network in AD.

Our familiar linear model framework allows for easy adoption and understanding of the model and its results. Additionally, the linear model framework offers many natural extensions. One could easily include terms for additional covariates such as scanner effect or gender. Different structures for the variance components could also be implemented to capture a wider range of possible correlation structures.

Our current method has the advantage of allowing for joint modeling of a complete FC network rather than taking a massive univariate approach. We see this joint modeling as a significant step forward, but complete brain analyses are still not yet feasible due to high computational demands of a model fit to many ROIs and the limited sample size of many fMRI studies. Here we have fit models to 10 ROI networks, but many brain atlases include more than 100 regions. In the future, some form of regularization could be introduced into the model to allow for analysis of an entire brain atlas. The ability to run a full brain FC network analysis would alleviate the issues arising from ROI selection.

The selection of the proper structure for the variance components deserves more attention. While a block compound symmetry structure for Ψ has a natural interpretation similar to that of a random intercept, there are certainly other viable structures. Choosing between structures is not a trivial task. One way to alleviate the model selection dilemma is to introduce a more robust sandwich type estimator of $\widehat{\mathrm{Var}}(\hat{\beta})$, in which case incorrect specification of the variance would lead to valid inference with only a reduction in power.

6 Acknowledgments

Data used in preparation of this article were obtained from the ADNI database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: `http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf`.

References

Takeshi Amemiya. A note on a heteroscedastic model. *Journal of Econometrics*, 6(3):365–370, 1977.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300, 1995.

Alex Chase. Altered functional connectivity in preclinical dementia. *Nature Reviews Neurology*, 10(11):609, 2014.

Pin-Yu Chen, Jeng-Min Chiou, Ya-Fang Yang, Yu-Ting Chen, Hsin-Long Hsieh, Yu-Ling Chang, and Wen-Yih I Tseng. Heterogeneous aging effects on functional connectivity in different cortical regions: A resting-state functional mri study using functional data analysis. *PLOS ONE*, 11(9):e0162028, 2016.

EunYi Chung and Joseph P Romano. Exact and asymptotically robust permutation tests. *The Annals of Statistics*, 41(2):484–507, 2013.

Mark Fiecas, Ivor Cribben, Reyhaneh Bahktiari, and Jacqueline Cummine. A variance components model for statistical inference on functional connectivity networks. *NeuroImage*, 149:256–266, 2017.

Emily S Finn and R Todd Constable. Individual variation in functional brain connectivity: implications for personalized approaches to psychiatric disease. *Dialogues in Clinical Neuroscience*, 18(3):277–287, 2016.

Michael D Fox and Marcus E Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8(9):700–711, 2007.

KJ Friston, CD Frith, PF Liddle, and RSJ Frackowiak. Functional connectivity: the principal-component analysis of large (pet) data sets. *Journal of Cerebral Blood Flow & Metabolism*, 13(1):5–14, 1993.

Habib Ganjgahi, Anderson M Winkler, David C Glahn, John Blangero, Peter Kochunov, and Thomas E Nichols. Fast and powerful heritability inference for family-based neuroimaging studies. *NeuroImage*, 115:256–268, 2015.

Michael D Greicius, Gaurav Srivastava, Allan L Reiss, and Vinod Menon. Default-mode network activity distinguishes alzheimer's disease from healthy aging: evidence from functional mri. *Proceedings of the National Academy of Sciences of the United States of America*, 101(13):4637–4642, 2004.

Anne Hafkemeijer, Jeroen van der Grond, and Serge ARB Rombouts. Imaging the default mode network in aging and dementia. *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease*, 1822(3):431–441, 2012.

Dominic Holland, Rahul S Desikan, Anders M Dale, Linda K McEvoy, Alzheimers Disease Neuroimaging Initiative, et al. Rates of decline in alzheimer disease decrease with age. *PLOS ONE*, 7(8):e42325, 2012.

Nan Laird. Analysis of longitudinal and cluster-correlated data. In *NSF-CBMS Regional Conference Series in Probability and Statistics*, pages i–155. JSTOR, 2004.

Yuxia Li, Xiaoni Wang, Yongqiu Li, Yu Sun, Can Sheng, Hongyan Li, Xuanyu Li, Yang Yu, Guanqun Chen, Xiaochen Hu, et al. Abnormal resting-state functional connectivity strength in mild cognitive impairment and its conversion to alzheimers disease. *Neural Plasticity*, 2016, 2015.

Guy Melard, Marianne Paesmans, and Roch Roy. Consistent estimation of the asymptotic covariance structure of multivariate serial correlations. *Journal of Time Series Analysis*, 12(4):351–361, 1991.

Belinda Phipson, Gordon K Smyth, et al. Permutation p-values should never be zero: calculating exact p-values when permutations are randomly drawn. *Stat Appl Genet Mol Biol*, 9(1):Article39–Article39, 2010.

Ping Ren, Raymond Y Lo, Benjamin P Chapman, Mark Mapstone, Anton Porsteinsson, and Feng Lin. Longitudinal alteration of intrinsic brain activity in the striatum in mild cognitive impairment. *Journal of Alzheimer's Disease*, 54(1):69–78, 2016.

Roch Roy. Asymptotic covariance structure of serial correlations in multivariate time series. *Biometrika*, 76(4): 824–827, 1989.

Christian Sorg, Valentin Riedl, Mark Mühlau, Vince D Calhoun, Tom Eichele, Leonhard Läer, Alexander Drzezga, Hans Förstl, Alexander Kurz, Claus Zimmer, et al. Selective changes of resting-state networks in individuals at risk for alzheimer's disease. *Proceedings of the National Academy of Sciences*, 104(47):18760–18765, 2007.

Kaustubh Supekar, Vinod Menon, Daniel Rubin, Mark Musen, and Michael D Greicius. Network analysis of intrinsic functional brain connectivity in alzheimer's disease. *PLoS Comput Biol*, 4(6):e1000100, 2008.

Cajo JF Ter Braak. Permutation versus bootstrap significance tests in multiple regression and anova. In *Bootstrapping and related techniques*, pages 79–85. Springer, 1992.

Nathalie Tzourio-Mazoyer, Brigitte Landeau, Dimitri Papathanassiou, Fabrice Crivello, Olivier Etard, Nicolas Delcroix, Bernard Mazoyer, and Marc Joliot. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. *Neuroimage*, 15(1):273–289, 2002.

Wiesje M van der Flier, Yolande AL Pijnenburg, Nick C Fox, and Philip Scheltens. Early-onset versus late-onset alzheimer's disease: the case of the missing apoe 4 allele. *The Lancet Neurology*, 10(3):280–288, 2011.

Kun Wang, Meng Liang, Liang Wang, Lixia Tian, Xinqing Zhang, Kuncheng Li, and Tianzi Jiang. Altered functional connectivity in early alzheimer's disease: a resting-state fmri study. *Human Brain Mapping*, 28(10): 967–978, 2007.

Liang Wang, Yufeng Zang, Yong He, Meng Liang, Xinqing Zhang, Lixia Tian, Tao Wu, Tianzi Jiang, and Kuncheng Li. Changes in hippocampal connectivity in the early stages of alzheimer's disease: evidence from resting state fmri. *NeuroImage*, 31(2):496–504, 2006.

Liang Wang, Catherine M Roe, Abraham Z Snyder, Matthew R Brier, Jewell B Thomas, Chengjie Xiong, Tammie L Benzinger, John C Morris, and Beau M Ances. Alzheimer disease family history impacts resting state functional connectivity. *Annals of Neurology*, 72(4):571–577, 2012.

Anderson M Winkler, Gerard R Ridgway, Matthew A Webster, Stephen M Smith, and Thomas E Nichols. Permutation inference for the general linear model. *NeuroImage*, 92:381–397, 2014.

Jie Xiang, Hao Guo, Rui Cao, Hong Liang, Junjie Chen, et al. An abnormal resting-state functional brain network indicates progression towards alzheimer's disease. *Neural Regeneration Research*, 8(30):2789, 2013.