



Bayesian longitudinal spectral estimation with application to resting-state fMRI data analysis[☆]

Ning Dai^{a,*}, Galin L. Jones^a, Mark Fiecas^b

^aSchool of Statistics, University of Minnesota, Minneapolis, MN 55455, USA

^bDivision of Biostatistics, University of Minnesota, Minneapolis, MN 55455, USA

ARTICLE INFO

Article history:

Received 12 April 2018

Revised 21 January 2019

Accepted 21 January 2019

Available online 25 February 2019

Keywords:

Resting-state fMRI

Longitudinal data analysis

Spectral estimation

Alzheimer's disease

Amplitude of low-frequency fluctuation

Markov chain Monte Carlo

ABSTRACT

The amplitude of the oscillatory patterns present in spontaneous fluctuations of brain signals obtained from resting-state functional magnetic resonance imaging (fMRI), measured using an index called the fractional amplitude of low-frequency fluctuation (fALFF), is a well-known measure of brain activity with potential to serve as a marker for brain dysfunction. With the rise of longitudinal neuroimaging studies, there is a great need for methodologies that take advantage of the longitudinal design in modeling the impact of aging or disease progression. Motivated by the longitudinal design of the Alzheimer's Disease Neuroimaging Initiative (ADNI), a novel Bayesian longitudinal model is developed in order to estimate the spectra of resting-state fMRI time courses, from which one can extract estimates of fALFF that are potentially associated with aging. The model incorporates within-subject correlation to improve estimates of the spectra, in addition to the variability that naturally arises between subjects. The model is validated using simulated data to show the gains in performance for estimating fALFF by taking advantage of the longitudinal design. Finally, a longitudinal analysis on fALFF from the resting-state fMRI data from ADNI is conducted, where the impact of both Alzheimer's disease and aging on the spontaneous fluctuations of brain activity is shown.

© 2019 EcoSta Econometrics and Statistics. Published by Elsevier B.V. All rights reserved.

1. Introduction

In the past decade, a plethora of studies have investigated the clinical utility of indices derived from spontaneous fluctuations of the functional magnetic resonance imaging (fMRI) signals observed during rest (Fox and Raichle, 2007). While much of the literature has investigated the spatial distribution of the temporal correlations of spontaneous fluctuations at spatially distinct regions of the brain, known as “resting-state functional connectivity”, a number of studies have shown that indices derived from the amplitudes of the spontaneous fluctuations are stable indices of brain activity and have potential to differentiate clinical populations compared to healthy controls (Zhang et al., 2008; Zuo et al., 2010; Han et al., 2011; Hu et al., 2014).

[☆] Data used in preparation of this article were obtained from the Alzheimer's 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.

* Corresponding author.

E-mail address: daixx224@umn.edu (N. Dai).

Specifically, an index called amplitude of low-frequency fluctuation (ALFF) of the resting-state fMRI signal has been suggested to reflect the intensity of regional spontaneous brain activity (Zang et al., 2007). This index is based on prior resting-state fMRI studies that showed low frequency components of the fMRI signals are closely related to spontaneous brain activities in the bilateral visual areas (Lowe et al., 1998), the bilateral auditory areas (Cordes et al., 2000), the language systems (Hampson et al., 2002), and the default mode network (DMN) (Greicius et al., 2003). It has been shown that ALFF is significantly higher within the DMN as well as in non-specific areas, which might be due to higher physiological noise in these areas inflating fluctuations across the entire frequency range and thus inflating ALFF (Zang et al., 2007). Seeking to provide a more stable index, Zou et al. (2008) proposed a fractional ALFF (fALFF) approach which selectively enhances signals from cortical regions associated with brain activity while suppressing artifacts from non-specific areas.

Both ALFF and fALFF are indices derived from the spectrum of the resting-state fMRI time courses. The spectrum of a time series is based on the Fourier transformation of the time series, and it quantifies the strength of the variations in the oscillations in the signal as a function of frequency. Thus, the amplitudes of the oscillations within a frequency band of interest can be obtained from the square root of the spectrum evaluated at that frequency band. Specifically, ALFF is defined as the square root of the spectrum averaged across the low-frequency range (0.01–0.08 Hz), while fALFF is the ratio of the square root of the spectrum summed over the low-frequency range to that of the entire frequency range. In practice, both ALFF and fALFF are calculated using an estimate of the spectrum of an fMRI signal.

As the human brain is a constantly changing and developing organ, longitudinal fMRI studies, in which subjects undergo a scan and are measured repeatedly over time, are essential to gain a better understanding of the processes that impede and affect normal development as well as the progression of disease. Compared to cross-sectional studies, longitudinal designs increase statistical power, help identify predictors of change, and distinguish between- and within-subject variation. Despite these design advantages, longitudinal studies can be methodologically challenging mainly because we must account for the correlation between the repeated measures, which only exacerbate the modeling challenges already inherent in fMRI data (Skup, 2010). While there has been work on analyzing longitudinal neuroimaging data of other modalities such as DTI (Goldsmith et al., 2012; Zipunnikov et al., 2014), PET (Chen and DuBois Bowman, 2011; Hyun et al., 2016), and structural MRI (James et al., 2011; Skup et al., 2012), these models are not applicable to fMRI data because they do not deal with time series data. A recent work on longitudinal fMRI methodology by Hart et al. (2018), however, was developed for functional connectivity analyses, and cannot be generalized to model fALFF in a longitudinal manner.

Our primary goal is to develop a longitudinal model for fALFF. To this end, we begin by developing a longitudinal model for the spectrum of fMRI time series. To our knowledge, existing statistical methodologies for estimating the spectrum of time series data are only appropriate for cross-sectional designs. As a result, existing methodologies for estimating fALFF are inappropriate because they do not account for the longitudinal design, leading to biased or inefficient estimates as well as misleading statistical inferences. To address these issues, we introduce a Bayesian longitudinal model that estimates the spectrum of the fMRI time series obtained from multiple visits while accounting for the longitudinal design. The proposed method utilizes additional longitudinal information and enables the characterization of functional changes that occur over time in the low frequency oscillations in the brain.

The rest of the paper is organized as follows. In Section 2, we describe the novel Bayesian longitudinal model with an accompanying computational strategy for spectral estimation. We further apply the proposed method to the Alzheimer's Disease Neuroimaging Initiative (ADNI) data and conduct a resting-state fMRI analysis and inference using fALFF. The performance of the proposed method is examined in a simulation study. Section 3 presents the results of the ADNI data application and the simulation study. We end with a discussion and conclusions in Sections 4 and 5, respectively.

2. Methods

2.1. Experimental data and analysis overview

Our primary goal is to model aging effects in the low frequency oscillations of spontaneous brain activity, which we quantified using fALFF, and to investigate how Alzheimer's disease (AD) moderates the association between fALFF and aging. We selected 8 regions of interest (ROIs): bilateral posterior cingulate (PCC), precuneus (PCU), supplementary motor area (SMA), and putamen (PU). These ROIs were selected based on previous cross-sectional studies on the association between AD with fALFF and aging with fALFF (Greicius et al., 2004; Hu et al., 2014; Ren et al., 2016). ROI-level time series were obtained by averaging all the time series within an ROI.

We applied the proposed Bayesian longitudinal model to a subset of the ADNI data. The resting-state fMRI images were collected at baseline, 3 months from baseline, 6 months from baseline, 12 months from baseline, and annually thereafter. The data consists of patients with AD and those who were cognitively normal (CN). To better separate the groups, only patients that remained in one group for the entirety of the follow-up were considered in our analysis. We focused our attention on late-onset AD and included only patients that were 65 years of age or older at baseline (van der Flier et al., 2017; Holland et al., 2012). To study the longitudinal trend within subject, only patients having at least two visits were considered. The remaining CN group consists of 106 visits from 25 patients with each patient having between 2 and 6 visits. The AD group consists of 72 visits from 19 patients with each patient having between 2 and 5 visits. The average age is 75.74 for the CN group with a range of 65.2–90.9. The AD group average age is 77.17 with a range of 68.3–88.6.

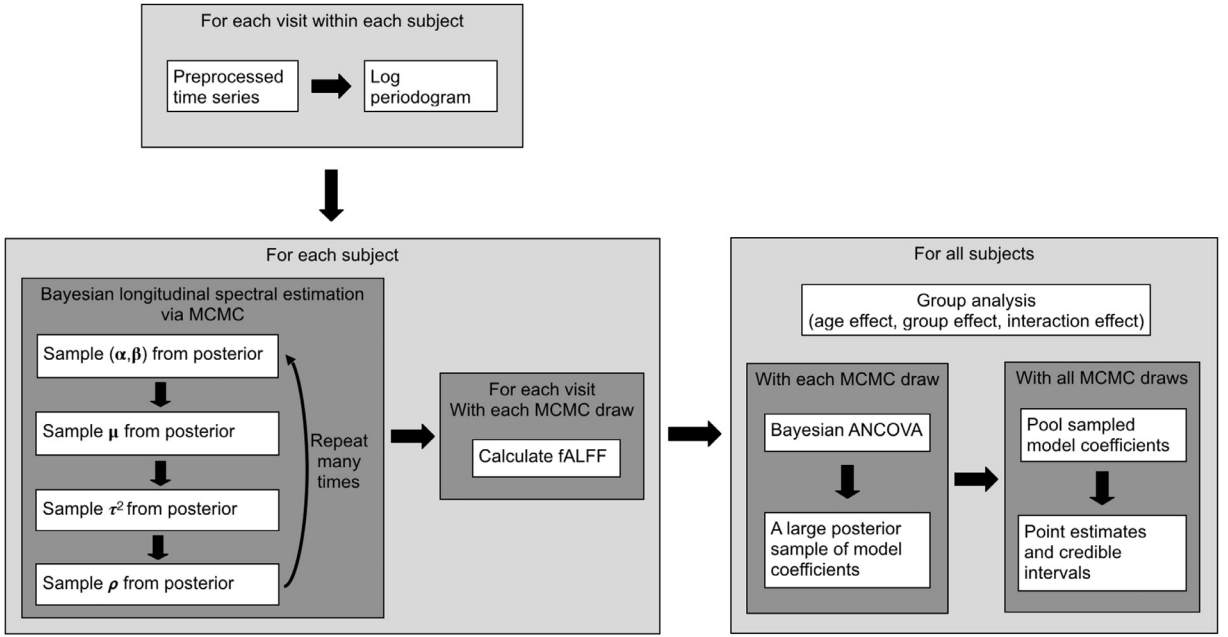


Fig. 1. A workflow chart of the fALFF analysis procedure using the proposed Bayesian longitudinal spectral estimation model.

We discarded the first 6 volumes of each functional time series to allow for magnetization equilibrium and for the adaptation of the subjects to the scanning situation. The remaining 134 volumes were preprocessed using both FSL (version 5.0.9, <https://fsl.fmrib.ox.ac.uk/>) and AFNI (version 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 `mcfliirt` (rigid body transform; cost function normalized correlation; middle volume was used as reference 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.

We provide below a sketch of our data analysis steps as shown in Fig. 1. First, we calculated the log periodogram of each preprocessed time series. For each subject, we obtained a series of log periodograms, one for each visit (Fig. 3). Second, for each subject separately, we used the proposed longitudinal model to jointly estimate the spectra across visits. Based on the estimated spectrum for each visit, we calculated the fALFF as described in Section 1. Third, we carried out group analysis using the estimated values of fALFF for all subjects and all visits to investigate how fALFF is associated with aging and AD.

2.2. Longitudinal spectral estimation model

We describe our longitudinal model for a single subject with V visits, and then we will describe our modeling strategy for all subjects. In the following, we fit our model to each ROI independently. For a given subject, let $\mathbf{x}_v = (x_{v1}, \dots, x_{vK})^\top$ denote the ROI-level time series observed at the v th visit, where K is the number of observed time points, and let $f_v(\cdot)$ denote its spectrum. Define the vector of the Fourier frequencies $\boldsymbol{\omega} = (\omega_0, \omega_1, \dots, \omega_{K-1})^\top$, where $\omega_k = k/K$, for $k = 0, \dots, K-1$. Our modeling approach begins with the visit-specific periodogram $I_v(\boldsymbol{\omega})$ to estimate the spectrum, which we can obtain by transforming the time series to the frequency domain using the fast Fourier transform (FFT) (Shumway and Stoffer, 2011). Our longitudinal spectral estimation method is motivated by the framework developed by Rosen et al. (2012), where the spectrum was modeled on the log scale. To this end, let $\mathbf{g}_v = \log(f_v(\boldsymbol{\omega}))$ and $\mathbf{y}_v = \log(I_v(\boldsymbol{\omega}))$, and consider the Whittle likelihood (Whittle, 1957)

$$p(\mathbf{y}|\mathbf{g}) = \prod_{v=1}^V p(\mathbf{y}_v|\mathbf{g}_v) \propto \exp\left[-\frac{1}{2}\mathbf{1}_{VK}^\top\{\mathbf{g} + \exp(\mathbf{y} - \mathbf{g})\}\right], \quad (1)$$

where $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_V^\top)^\top$, $\mathbf{g} = (\mathbf{g}_1^\top, \dots, \mathbf{g}_V^\top)^\top$, and $\mathbf{1}_{VK}$ is a VK -vector of ones.

We model the log spectrum at each visit using smoothing splines, $\mathbf{g}_v = \alpha_v \mathbf{1}_K + B\boldsymbol{\beta}_v$, where the scalar α_v is the constant component and $B\boldsymbol{\beta}_v$ captures the nonlinear and smooth component. The columns of B are $\sqrt{2}\cos(2j\pi\boldsymbol{\omega})$ for $j = 1, \dots, J$, i.e., the Demmler–Reinsch basis functions (Eubank, 1999; Rosen et al., 2012) and $\boldsymbol{\beta}_v = (\beta_{v1}, \dots, \beta_{vJ})^\top$ is a vector of spline

coefficients. For computational efficiency, we retain only the first $J = 6$ basis functions, corresponding to the $J = 6$ largest eigenvalues. In our analysis, we considered $J = 8$ and $J = 10$ and observed a negligible difference, while using $J = 5$ led to an over-smoothed estimate.

The above described visit-specific models. To model the spectra across visits, we use the Whittle likelihood in Eq. (1) and prior distributions on the spline coefficients to induce correlation in the spectra across visits.

2.3. Priors and posteriors

To jointly estimate the spectra across visits, the prior on the spline coefficients $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_V^\top)^\top$ is set to be dependent across visits. To allow for modeling unevenly spaced visits, we choose the continuous-time AR(1) correlation structure, i.e., the AR decay with respect to the elapsed time, as opposed to visit count. Let η_ν be the elapsed time between the ν th and $(\nu + 1)$ th visits, where $\nu = 1, \dots, V - 1$, and further define $\boldsymbol{\eta} = (\eta_1, \dots, \eta_{V-1})^\top$. Then let $R_\eta(\rho)$ denote the correlation matrix, whose diagonal entries are all ones and the (u, v) th or (v, u) th entry is $\rho^{\eta_u + \dots + \eta_{v-1}}$ for all $1 \leq u < v \leq V$. This term captures the dependence between multiple visits from the same subject. Other correlation structures are possible.

Within a visit, the spline coefficients are assumed independent and centered at a spline-basis-function-specific mean $\boldsymbol{\mu}$, whose j th element corresponds to the j th spline basis function, where $j = 1, \dots, J$. This is different from the prior in Rosen et al. (2012) which is centered at zero. Our modification guarantees that, within each subject, the coefficients across visits will shrink towards the subject's log spectrum averaged across all visits instead of shrinking to zero, which corresponds to the undesirable consequence of the spectrum shrinking to a flat line.

In contrast to the spline coefficients, the prior on the constant components $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_V)^\top$ is independent across visits. This is not only for the sake of simplicity, but also because we are more interested in how the distribution of power across frequencies, which is captured by the nonlinear component $B\boldsymbol{\beta}_\nu$ for each visit ν , varies across visits compared to the overall level, represented by α_ν .

The likelihood (1) and the prior (A.1) (see Appendix A) give rise to a posterior distribution $p(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \tau^2, \rho | \mathbf{y})$ which we want to use for inference. However, the posterior is analytically intractable and hence we are forced to use Markov chain Monte Carlo (MCMC) to estimate posterior quantities of interest. The required conditional posterior distributions are given in Appendix B. Details of our MCMC algorithm are provided in Appendix C.

To reduce the computational cost, we used the Hamiltonian Monte Carlo (HMC) (Neal, 2011) instead of random walk Metropolis-Hastings as the sampling scheme. This is because HMC proposals are less correlated and converge to the posterior distribution more rapidly, and thus it takes fewer iterations to achieve the same level of precision of estimating the posterior mean.

Ideally, the number of leapfrog steps and the stepsize in HMC need to be tuned on a case-by-case basis to achieve good performance, where the optimal acceptance rate is 65% (Neal, 2011). For ease of computation, we tuned the parameters on a subset of the experimental data and then used the same parameter values on all data. In our analysis, the number of steps and the stepsize in sampling $(\boldsymbol{\alpha}^\top, \boldsymbol{\beta}^\top)^\top$ were 15 and 0.065 respectively, with an acceptance rate ranging from 0.43 to 0.86. For simulating ρ , the number of steps and the stepsize were 10 and 0.13 respectively, with an acceptance rate ranging from 0.51 to 0.78. We also assessed performance by examination of trace plots, density plots, and autocorrelation plots, all of which indicated fast mixing of the Markov chain simulation. In addition, multivariate effective sample size (ESS) calculations (Vats et al., 2018; Dai and Jones, 2017) indicated that a burn-in of 100 was sufficient in the data application. See Flegal et al. (2008) and Geyer (2011) for more on this approach.

To determine when to terminate the simulation, we used the relative standard deviation fixed-volume sequential stopping rule (Vats et al., 2015), where the multivariate initial sequence method (Dai and Jones, 2017) was implemented to assess the Monte Carlo error. Analyses were performed in R 3.3.3 (R Core Team, 2017) using the `mcmcse` package (Flegal et al., 2017). Note that implementing the sequential stopping rule itself is time-consuming, especially when the number of visits and the number of spline coefficients are not small. Indeed, in our data application, it was impractical to compute the sufficient number of iterations for all subjects. Therefore, we performed the calculation on a subset of the experimental data and found out that when the number of iterations was 20,000, the estimated multivariate ESS was greater than 5300. We then applied the same setting, 20,000 iterations to all data. The MCMC algorithm took about 50 s for one subject.

2.4. Data analysis

To study how fALFF depends on covariates, population-level modeling is essential. However, the proposed spectral estimation method applies to individual subjects. The computational burden is one of the main barriers to extending the proposed method to multisubject modeling. On the other hand, single-subject modeling is more flexible as it does not assume that different subjects have similar smoothness across visits. Therefore, we decided to keep the spectral estimation model at the single-subject level.

For each subject independently, we estimate the spectra using our Bayesian longitudinal model, and then calculate the fALFF based on each estimated spectrum as described in Section 1. Then to model fALFF at the population level, we combined the longitudinal subject-specific results in a regression analysis. The two-stage formulation of the full procedure is

briefly described below and demonstrated in a workflow chart in Fig. 1. More details are provided in Sections 2.4.1 and 2.4.2.

1. At the subject level,
 - (a) we conducted the proposed Bayesian model for longitudinal spectral estimation via MCMC;
 - (b) we calculated the fALFF value for each visit and with each MCMC draw.
2. At the population level,
 - (a) we fit a Bayesian regression model with each MCMC draw of fALFF;
 - (b) we pooled the posterior samples of regression model coefficients from all MCMC draws and then used the pooled sample for inference.

2.4.1. Subject-level analysis

We implemented the proposed Bayesian model for longitudinal spectral estimation on each subject. The elapsed time used in the correlation structure of the spline coefficients was taken in units of 3 months, which is the shortest and most common time interval in the ADNI data.

For each subject, at each MCMC iteration, we computed the spectra of all visits using the sampled parameters and further obtained the fALFF values. Recall from Section 1 that fALFF is computed as the ratio of the square root of the spectrum summed over the low-frequency range to that of the entire frequency range. Instead of averaging over a sampled Markov chain, we retained all the MCMC draws of fALFF to accurately capture the variance in spectral estimation.

2.4.2. Population-level analysis

To investigate how fALFF depends on covariates, we performed regression analysis, making use of all the MCMC iterations obtained at the subject level. Specifically, we fit a Bayesian regression model, separately using each MCMC draw of the fALFF values of all subjects and visits. Priors were selected to ensure that the posterior of the regression coefficients and the error variance was tractable so that sampling from the posterior was easy and fast.

Let \boldsymbol{y} and σ^2 denote the vector of the regression coefficients and the error variance, respectively. The prior takes the form $p(\boldsymbol{y}, \sigma^2) = p(\sigma^2)p(\boldsymbol{y}|\sigma^2)$, where $p(\sigma^2) \propto \sigma^{-1}I(\sigma^2 \leq c_\sigma^2)$ for some c_σ^2 is induced from a uniform prior on σ (Gelman, 2006). The conditional prior $p(\boldsymbol{y}|\sigma^2)$ is a normal distribution $\boldsymbol{y}|\sigma^2 \sim N(\hat{\boldsymbol{y}}, \sigma^2 I)$, where $\hat{\boldsymbol{y}}$ is a hyper-parameter. A popular data-driven value of $\hat{\boldsymbol{y}}$ is the ordinary least squares estimate. Let $\boldsymbol{\psi}$ and X denote the vector of response values and the design matrix, respectively. Then the posterior can be written as $p(\boldsymbol{y}, \sigma^2 | \boldsymbol{\psi}, X) = p(\boldsymbol{y} | \sigma^2, \boldsymbol{\psi}, X)p(\sigma^2 | \boldsymbol{\psi}, X)$, where the two factors correspond to the densities of $N(\hat{\boldsymbol{y}}, \sigma^2(X^\top X + I)^{-1})$ and an inverse gamma distribution $IG((n-1)/2, (\boldsymbol{\psi}^\top \boldsymbol{\psi} - \hat{\boldsymbol{y}}^\top X^\top X \hat{\boldsymbol{y}})/2)$ truncated at $\sigma^2 \in (0, c_\sigma^2]$, and n is the length of $\boldsymbol{\psi}$.

With each MCMC draw, we obtained a large posterior sample of the regression coefficients. We then pooled the samples from all MCMC draws. The pooled sample reflects both sources of variation in estimating the effects: (1) the error in spectral estimation, reflected in the MCMC draws obtained at the subject level, and (2) the error in estimating the regression coefficients, reflected in the posterior samples obtained at the population level. Therefore, using the pooled sample for inference enables us to properly account for the errors and to assess effects. A Bayesian point estimate for each regression coefficient was computed as the average of the pooled sample. Further a credible interval was constructed by taking the corresponding sample quantiles.

An alternative to using all MCMC draws is to average across the simulated Markov chain. This approach is in accord with the literature on fALFF analysis where the smoothed periodograms are used as the true spectra. However, this approach ignores the variation introduced in spectral estimation, resulting in unreasonably narrow credible intervals for the effects and inflated false positive rates. Using all MCMC draws can address these issues effectively, but poses computational challenges, which we were able to manage by using parallel computing and HMC.

In our application, we considered the ANCOVA model with fALFF as the response, age as a continuous predictor, AD status as a group effect, and an interaction term for age and AD status. To put the main coefficients on the same scale, we centered age at 75 years old and then divided it by 10 years. The interaction coefficient reflects the difference between AD and CN in the change of fALFF per 10 years.

2.5. Simulation experiment

We examined the performance of the proposed method via a simulation experiment. To simulate realistic ROI-level time series, we randomly selected a CN subject “002_S_5230” and an AD subject “006_S_4153” from the ADNI dataset. The time series in the left PCU were detrended and the spectra were estimated using the proposed Bayesian longitudinal method. We used the estimated spectra of the first three visits of both subjects as the true group-level spectra to simulate data. Therefore, the simulated data was ensured to be close to the ADNI data. Then we jittered the ground truth to add subject-specific variability. The data generating process is described below, and demonstrated via a workflow chart in Fig. 2.

1. We jittered the estimated spline coefficients of the first three visits of both subjects “002_S_5230” and “006_S_4153” 20 times by adding 20 different random samples from a normal distribution with mean zero and a small variance to the estimated spline coefficients. Consequently, each subject generated 20 sets of spline coefficients, with each set having 3 visits.






Fig. 2. A workflow chart of the data generating process for the simulation experiment.



(a) Log spectrum estimates of the ROI-level time series in the right PCC of Subject “130_S_4730”.



(b) Log spectrum estimates of the ROI-level time series in the left PU of Subject “002_S_5018”.

Fig. 3. Log spectrum estimates of two randomly selected subjects in the ADNI dataset. Displayed are the periodogram (gray), smoothed periodogram (blue) and Bayesian longitudinal estimate (black) with 95% pointwise credible intervals (dashed). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

2. We treated the obtained 40 sets of spline coefficients as the true spline coefficients of the 40 generated subjects and calculated the corresponding true spectra using the smoothing spline model.
3. For each visit within each subject, we simulated one time series from the true spectrum using the time-frequency toggle (Kirch and Politis, 2011; Fiecas and von Sachs, 2014).

We replicated the data generating process to obtain 100 datasets that share a common group-level spectrum, with each individual-level spectrum being the group-level spectrum with an individual-specific and a visit-specific jitter. Each simulated dataset contains 20 CN subjects and 20 AD subjects, with each subject having 3 visits and each visit within each subject having one time series.

We implemented the proposed Bayesian longitudinal method on each simulated subject to estimate the spectra of the time series of the three visits simultaneously. At each MCMC iteration, we recorded the estimated spline coefficients and the estimated log spectra. A point estimate of the log spectrum value at each frequency grid was obtained by averaging over the simulated Markov chain. A nominal 95% Bayesian credible interval of the log spectrum value at each frequency grid was constructed using the 2.5% and 97.5% sample quantiles. The point estimate and pointwise credible intervals for log spectrum were then exponentiated to obtain the corresponding point estimate and pointwise credible interval for the spectrum.

To assess the quality of the spectral estimation, we considered two error criteria: the mean squared error (MSE) for estimating the log spectrum evaluated at each frequency, and the total variation distance (TVD) between a normalized spectrum estimator and the true normalized spectrum. TVD measures the dissimilarity between two spectra in terms of the distribution of the power across frequencies rather than the magnitude of the oscillations (Alvarez-Esteban et al., 2016; Euán et al., 2018).

We examined the performance of three methods: the smoothed periodogram, the single-visit Bayesian model, where the multiple visits of each subject were treated as independent single-visit subjects, and the proposed longitudinal Bayesian




Fig. 4. Boxplot of the estimated within-subject correlation coefficients.

method that models the correlation across visits. The first method is a standard approach (Shumway and Stoffer, 2011). Comparing the second and third methods allows us to see the gains in performance by accounting for the correlation across visits. We compared the methods using the percent improvement in TVD or MSE relative to the TVD or MSE in using the smoothed periodogram.

We paid special attention to how well we estimated fALFF, the variable of interest in our ADNI data analysis. Three frequency bands were considered for fALFF calculation: the low-frequency band (0.01–0.08 Hz), slow-4 (0.027–0.073 Hz), and slow-5 (0.01–0.027 Hz). For each simulated time series, we computed a spectrum estimate, based on which the fALFF was approximated and a squared error was obtained. Then by averaging across all subjects within a simulated dataset, we obtained an estimate for MSE.

3. Results

3.1. ADNI data analysis

Fig. 3 shows the log spectrum estimates of two randomly selected subjects obtained using the proposed method and the smoothed periodogram approach. As observed in Fig. 3(a) visit 1–3 and (b) visit 3, the smoothed periodogram estimates using the generalized cross validation (GCV) of span selection (Ombao et al., 2001) appeared too flat. In comparison, the estimates obtained using the proposed model were in better shape. On the other hand, the log raw periodogram could be very low at frequency 0, as presented in Fig. 3(a) visit 3 and (b) visit 3, thus pulling down the smoothed periodogram. By borrowing information from the other visits of the same subject, the Bayesian longitudinal method overcame this problem.

Fig. 4 shows for each ROI the estimated within-subject correlations of all subjects. Recall that the window of the correlation coefficient is 3 months. The degree of within-subject correlation varied across subjects and across regions. In general, the correlations appeared well above zero. For most subjects, we observed moderate to high correlations in all the ROIs.

Table 1 shows the point estimates and the 95% credible intervals for the age effect, the group effect I_{AD} , and the interaction between age and I_{AD} . Recall that we rescaled age in the unit of 10 years, as described in Section 2.4.2. For the CN subjects, the age effect appeared significant in the bilateral PCU and PU in the low-frequency band, and in the bilateral PU in the slow-4 band. Specifically, fALFF in the low-frequency band decreased between approximately 0.01 and 0.04 every 10 years in the left PCU and the bilateral PU, while in the right PCU, the decrease was between approximately 0.00 and 0.03. The AD patients had significantly lower fALFF than the CN subjects in all the eight ROIs in both the low-frequency band and the slow-4 band, and in the bilateral PCU in the slow-5 band. The largest difference between AD and CN was in the left PCC and the bilateral PCU in the low-frequency band, in the left PCC in the slow-4 band, and in the bilateral PCU in the slow-5 band. The interaction between age and I_{AD} showed no difference between AD and CN except in the left PCU in the low-frequency band, where fALFF decreased with age in the CN subjects but not in the AD subjects.

3.2. Simulation experiment

Fig. 5 compares the log spectrum estimates of two randomly selected subjects obtained using the proposed Bayesian longitudinal method, the single-visit Bayesian method and the smoothed periodogram approach with GCV for automatic

Table 1
ADNI fMRI data analysis results.

ROI	Effect	Low-freq. (0.01–0.08 Hz)			Slow-4 (0.027–0.073 Hz)			Slow-5 (0.01–0.027 Hz)		
		Estimate	95% CI		Estimate	95% CI		Estimate	95% CI	
PCC L	Age	-0.0125	-0.0311	0.0061	-0.0111	-0.0281	0.0060	-0.0012	-0.0095	0.0072
	AD	-0.0319	-0.0485	-0.0152	-0.0260	-0.0411	-0.0108	-0.0060	-0.0135	0.0015
	Age:AD	0.0026	-0.0282	0.0335	0.0101	-0.0178	0.0381	-0.0081	-0.0219	0.0058
PCC R	Age	-0.0024	-0.0206	0.0159	-0.0052	-0.0214	0.0111	0.0025	-0.0056	0.0107
	AD	-0.0190	-0.0353	-0.0028	-0.0199	-0.0344	-0.0054	0.0016	-0.0057	0.0090
	Age:AD	-0.0105	-0.0405	0.0195	0.0034	-0.0232	0.0300	-0.0146	-0.0282	-0.0009
PCU L	Age	-0.0222	-0.0384	-0.0060	-0.0147	-0.0299	0.0006	-0.0066	-0.0149	0.0017
	AD	-0.0313	-0.0459	-0.0167	-0.0181	-0.0317	-0.0045	-0.0143	-0.0216	-0.0069
	Age:AD	0.0281	0.0011	0.0550	0.0215	-0.0036	0.0465	0.0064	-0.0072	0.0199
PCU R	Age	-0.0164	-0.0317	-0.0011	-0.0083	-0.0234	0.0068	-0.0063	-0.0143	0.0016
	AD	-0.0312	-0.0448	-0.0176	-0.0156	-0.0290	-0.0022	-0.0188	-0.0258	-0.0118
	Age:AD	0.0188	-0.0065	0.0440	0.0125	-0.0121	0.0372	0.0079	-0.0051	0.0209
SMA L	Age	-0.0046	-0.0228	0.0137	-0.0046	-0.0213	0.0121	0.0008	-0.0078	0.0094
	AD	-0.0182	-0.0347	-0.0018	-0.0188	-0.0338	-0.0038	0.0031	-0.0046	0.0107
	Age:AD	0.0071	-0.0235	0.0375	0.0127	-0.0151	0.0404	-0.0077	-0.0219	0.0065
SMA R	Age	-0.0076	-0.0276	0.0124	-0.0086	-0.0260	0.0088	0.0044	-0.0040	0.0127
	AD	-0.0179	-0.0358	-0.0000	-0.0196	-0.0351	-0.0040	0.0039	-0.0034	0.0113
	Age:AD	0.0195	-0.0137	0.0526	0.0201	-0.0087	0.0489	-0.0033	-0.0170	0.0104
PU L	Age	-0.0243	-0.0413	-0.0074	-0.0174	-0.0338	-0.0009	-0.0073	-0.0147	0.0002
	AD	-0.0224	-0.0376	-0.0071	-0.0232	-0.0381	-0.0083	-0.0004	-0.0071	0.0063
	Age:AD	0.0143	-0.0137	0.0422	0.0098	-0.0173	0.0370	0.0044	-0.0079	0.0167
PU R	Age	-0.0264	-0.0430	-0.0099	-0.0220	-0.0382	-0.0059	-0.0025	-0.0099	0.0049
	AD	-0.0150	-0.0298	-0.0001	-0.0159	-0.0302	-0.0015	0.0007	-0.0058	0.0073
	Age:AD	0.0139	-0.0135	0.0412	0.0167	-0.0098	0.0430	-0.0067	-0.0189	0.0054

Table 2
Percent improvement measured in MSE and TVD, using the Bayesian models relative to the smoothed periodogram approach. “l-Bayes” and “Bayes” are the longitudinal and single-visit Bayesian methods, respectively. Standard errors are given in parentheses.

Error type	Method	Overall	Visit 1	Visit 2	Visit 3
MSE	l-Bayes	0.240 (0.004)	0.230 (0.008)	0.253 (0.007)	0.237 (0.008)
	Bayes	0.039 (0.006)	0.018 (0.010)	0.051 (0.011)	0.048 (0.010)
TVD	l-Bayes	0.175 (0.003)	0.184 (0.005)	0.170 (0.005)	0.171 (0.005)
	Bayes	0.074 (0.003)	0.084 (0.006)	0.067 (0.006)	0.071 (0.006)

bandwidth selection. We observed that GCV tended to over smooth. Compared to the smoothed periodogram approach, Bayesian methods better captured the shape of the curve. The longitudinal estimate was generally closer to the truth than the single-visit estimate, especially at the ends and in regions with sharp curvatures. By forcing similar spectrum estimates across visits, the longitudinal estimate was smoother and more accurate, while the single-visit estimate appeared too variable.

Fig. 6 compares the three spectral estimation methods using two error criteria: MSE integrated across all frequencies, and TVD from the true spectrum. With each criterion, we observed that both Bayesian methods had generally smaller estimation errors than the smoothed periodogram approach. Further, the longitudinal Bayesian model showed smaller errors than the single-visit model, which shows the gains in accounting for the within-subject dependence. Table 2 shows how much the spectral estimation quality was improved by using the Bayesian models instead of the smoothed periodogram approach. Errors were computed for different visits and also integrated as “overall error”. The proposed Bayesian longitudinal model showed substantial improvement, generally about 24% and 17.5%, measured in MSE and TVD, respectively. Compared to the single-visit Bayesian model, the longitudinal model was at least 5 times better when using MSE, and 2 times better when using TVD.

Fig. 7 shows the MSE for estimating fALFF using different spectral estimation methods. Three frequency bands were considered: the low-frequency band (0.01–0.08 Hz), slow-4 (0.027–0.073 Hz), and slow-5 (0.01–0.027 Hz). On all three frequency bands, the Bayesian longitudinal model estimated fALFF with substantially smaller errors compared to the smoothed periodogram approach. The single-visit Bayesian method also showed better performance than the smoothed periodogram approach, especially on the slow-5 band, but in general inferior to the longitudinal model.

The take-home message was that the proposed longitudinal method allows for modeling developmental trends across aging, which cross-sectional models cannot do. In addition, we were able to obtain more precise estimates of the spectra by borrowing information across visits within an individual.

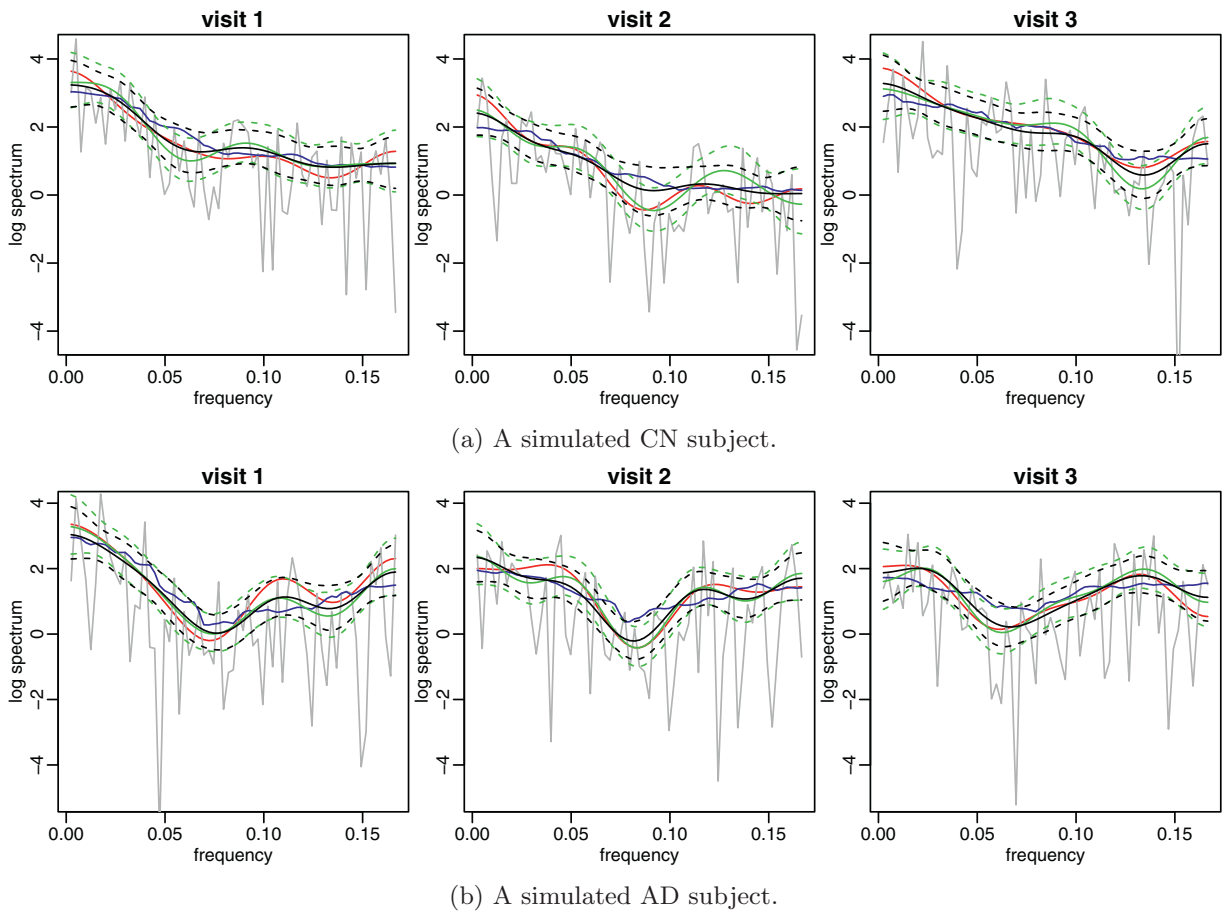


Fig. 5. True and estimated log spectra of two simulated subjects. Red is the truth. Gray is the raw periodogram. Blue is the smoothed periodogram. Black is the longitudinal Bayesian estimate (solid) with 95% pointwise credible interval (dashed). Green is the single-visit Bayesian estimate (solid) with 95% pointwise credible interval (dashed). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

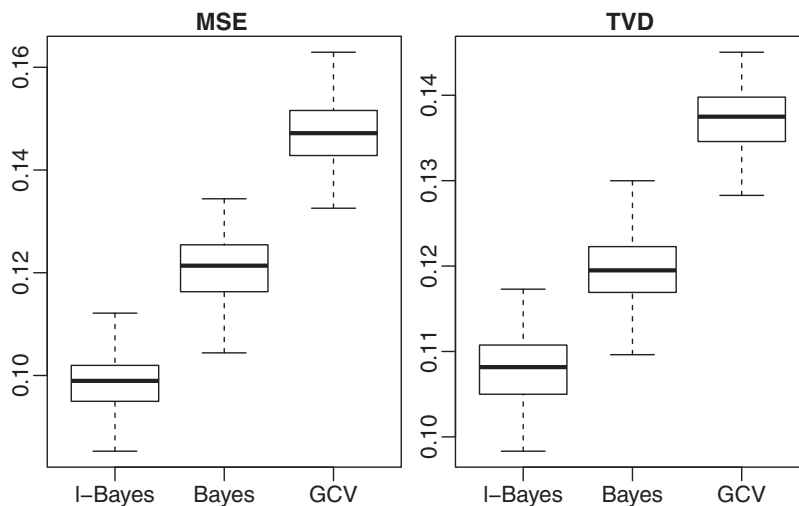


Fig. 6. Boxplots of MSE and TVD for spectral estimation. “I-Bayes” is the Bayesian longitudinal method; “Bayes” is the single-visit Bayesian method; “GCV” is the smoothed periodogram with GCV.

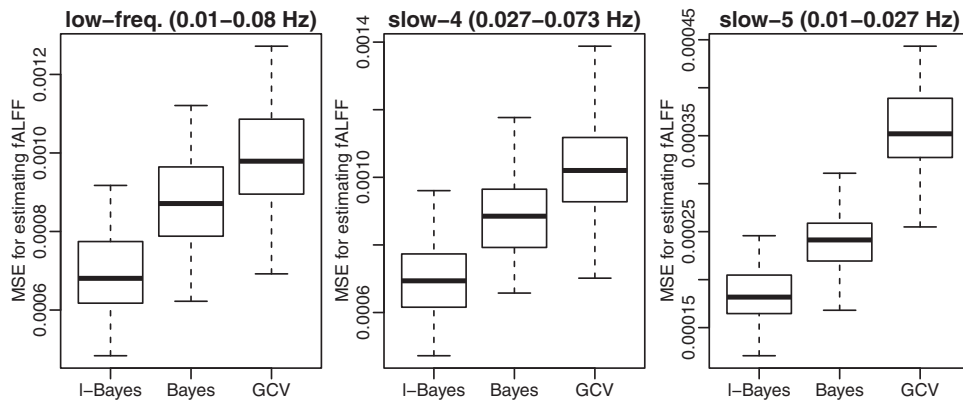


Fig. 7. Boxplots of MSE for estimating fALFF on three frequency bands: the low-frequency band (0.01–0.08 Hz), slow-4 (0.027–0.073 Hz), and slow-5 (0.01–0.027 Hz). “l-Bayes” is the Bayesian longitudinal method; “Bayes” is the single-visit Bayesian method; “GCV” is the smoothed periodogram with GCV.

4. Discussion

We developed a Bayesian longitudinal model for the spectra across visits to model and estimate longitudinal trend of fALFF. Using our model, we found that (i) many regions showed a decrease in fALFF between AD and CN, (ii) fALFF from many regions decreased with aging, and (iii) the differences between AD and CN in fALFF and their longitudinal trends were frequency-dependent.

When looking at the entire low-frequency band (0.01–0.08 Hz), we saw that the bilateral PCC and the bilateral PCU had lower fALFF values for those with AD. The disruption in the low frequency oscillations from these regions has been observed in cross-sectional studies on those with amnesic mild cognitive impairment or AD (Han et al., 2011; Liu et al., 2014; Cha et al., 2015; Long et al., 2016). Indeed, the PCC and PCU make up a part of the DMN, and the DMN has been shown to distinguish AD from healthy aging (Greicius et al., 2004). Furthermore, previous cross-sectional studies have provided evidence of disruption in the activity and in the functional networks that involve the SMA and the PU (Wang et al., 2007; Binnewijzend et al., 2012; Vidoni et al., 2012).

Healthy aging is associated with cognitive decline that affects speed, episodic memory, and reasoning, with effects most pronounced after the age of 60 (Hedden and Gabrieli, 2004). Our results showed that aging is associated with a decrease in low frequency activity, which agrees with previous studies on aging effects. For instance, previous studies have shown decreased activity in the DMN and disruptions in connectivity networks (Damoiseaux et al., 2007; Onoda et al., 2012). While results about aging effects on the DMN are fairly consistent, there have been contrasting results on aging effects that involve motor and subcortical networks (Sala-Llonch et al., 2015). Activity in the low frequencies within the SMA and the DMN have been shown to correlate negatively with age (Hu et al., 2014). Though we found aging effects in the PCU and the PU in the low-frequency band, we did not find aging effects in the PCC and the SMA. However, we did not find aging effects in fALFF in the slow-4 and slow-5 frequency bands except in the slow-4 frequency band for the PU, contrasting a previous study (Ren et al., 2016). We point out that previous resting-state fMRI studies on normal aging have compared cohorts comprised of “young” and “old” participants or were conducted using a sample of participants that covered a wide range in age (Damoiseaux et al., 2007), which contrasts our analysis that leverages the longitudinal design of ADNI to study aging effects. Therefore, our results on the longitudinal changes in fALFF enrich the existing literature on the aging effects on low frequency oscillations.

MCMC often raises computational challenges when data is high-dimensional, as is the case in longitudinal fMRI studies. Indeed, a voxel-level analysis was impractical, hence, we worked with ROI-level fMRI signals, substantially reducing the computation time. As spectral estimation was independent between ROIs and between subjects, we were able to run jobs in parallel on several machines simultaneously. Additionally, using HMC instead of random walk Metropolis-Hastings effectively lowered the computational cost.

One limitation of our analysis is that we did not account for missing data. Missing data is not uncommon in longitudinal fMRI studies. For example, situations may arise where data must be discarded due to excessive movement or scanner artifact. Data might also be missing due to missed appointments or attrition. Note that in spectral estimation, missing data means missing a whole spectrum curve instead of only one observation. The problem of missing data makes analysis even more challenging. In non-imaging analyses, there is a rich Bayesian literature on handling missing data (Daniels and Hogan, 2008; Diggle and Kenward, 1994; Little and Rubin, 2014; Molenberghs and Kenward, 2007). We hope to adapt the existing approaches to the longitudinal imaging setting in future research.

5. Conclusion

We introduced a novel Bayesian model for longitudinal spectral estimation of resting-state fMRI data. The model properly accounts for the correlation inherent in repeated measures data. We also provided a Bayesian framework for subsequent

data analysis and inference that accounts for the variation in spectral estimation. We used our method to study the longitudinal changes in the low frequency components of resting-state fMRI signals from the ADNI database. We generally saw lower fALFF in AD and a decreasing trend over time, consistent with previous studies on mild cognitive impairment and Alzheimer's disease. Our results highlight abnormalities of low frequency oscillations in AD and their association with aging.

The code for implementing the proposed method and data analysis is available at <https://github.com/mfiecas/longBayes>.

Acknowledgment

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI; Principal Investigator Michael W. Weiner, MD) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for NeuroImaging at the University of Southern California.

The authors would like to thank Timothy D. Johnson, James S. Hodges, and Brian B. Hart for helpful conversations.

Appendix A. Prior specifications

Let \otimes denote Kronecker product. We describe the prior specifications as follows:

$$\begin{aligned}\boldsymbol{\alpha} &\sim N(0, \sigma_{\alpha}^2 I_V), \\ \boldsymbol{\beta} &\sim N(\mathbf{1}_V \otimes \boldsymbol{\mu}, \tau^2 R_{\eta}(\rho) \otimes I_J), \\ \boldsymbol{\mu} &\sim N(0, \sigma_{\mu}^2 I_J), \\ \tau &\sim U(0, c_{\tau}), \text{ and} \\ \rho &\sim U(0, 1),\end{aligned}\tag{A.1}$$

where σ_{α}^2 , σ_{μ}^2 , and c_{τ} are fixed positive numbers. The uniform prior on τ , recommended by Gelman (2006), induces a prior on τ^2 , $p(\tau^2) \propto \tau^{-1} I(\tau^2 \leq c_{\tau}^2)$.

Appendix B. Conditional posterior distributions

The conditional posterior distributions are given by

$$p(\boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{y}, \boldsymbol{\mu}, \tau^2, \rho) \propto \exp \left[-\frac{1}{2} \mathbf{1}_{VK}^{\top} \{ \mathbf{g} + \exp(\mathbf{y} - \mathbf{g}) \} - \frac{\boldsymbol{\alpha}^{\top} \boldsymbol{\alpha}}{2\sigma_{\alpha}^2} - \frac{c}{2\tau^2} \right],\tag{B.1}$$

$$\boldsymbol{\mu} | \boldsymbol{\beta}, \tau^2, \rho \sim N \left(\frac{d}{\tau^2} [\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_V] R_{\eta}^{-1}(\rho) \mathbf{1}_V, d I_J \right),\tag{B.2}$$

$$\tau^2 | \boldsymbol{\beta}, \boldsymbol{\mu}, \rho \sim IG \left(\frac{VJ-1}{2}, \frac{c}{2} \right), \text{ for } \tau^2 \in (0, c_{\tau}^2], \text{ and}\tag{B.3}$$

$$p(\rho | \boldsymbol{\beta}, \boldsymbol{\mu}, \tau^2) \propto \exp \left(-\frac{c}{2\tau^2} \right) \prod_{v=1}^{V-1} (1 - \rho^{2\eta_v})^{-\frac{1}{2}} I(0 < \rho < 1),\tag{B.4}$$

where $c = (\boldsymbol{\beta} - \mathbf{1}_V \otimes \boldsymbol{\mu})^{\top} (R_{\eta}^{-1}(\rho) \otimes I_J) (\boldsymbol{\beta} - \mathbf{1}_V \otimes \boldsymbol{\mu})$ and $d = (\tau^{-2} \mathbf{1}^{\top} R_{\eta}^{-1}(\rho) \mathbf{1} + \sigma_{\mu}^{-2})^{-1}$.

Appendix C. MCMC algorithm

Our MCMC algorithm proceeds as follows:

1. α and β are sampled jointly from $p(\alpha, \beta | \mathbf{y}, \mu, \tau^2, \rho)$ (B.1).
2. μ is sampled from the normal distribution (B.2).
3. τ^2 is sampled from the truncated inverse gamma distribution (B.3).
4. ρ is sampled from $p(\rho | \beta, \mu, \tau^2)$ (B.4).

We used HMC in steps 1 and 4. In both HMC algorithms, the potential energy function of position was the negative log posterior density, and the kinetic energy function of momentum was assumed to be Gaussian.

References

- Alvarez-Esteban, P.C., Euán, C., Ortega, J., 2016. Time series clustering using the total variation distance with applications in oceanography. *Environmetrics* 27 (6), 355–369. doi:10.1002/env.2398. Env.2398.
- Binnewijzend, M.A.A., Schoonheim, M.M., Sanz-Arigitá, E., Wink, A.M., van der Flier, W.M., Tolboom, N., Adriaanse, S.M., Damoiseaux, J.S., Scheltens, P., van Berckel, B.N.M., Barkhof, F., 2012. Resting-state fMRI changes in Alzheimer's disease and mild cognitive impairment. *Neurobiol. Aging* 33 (9), 2018–2028. doi:10.1016/j.neurobiolaging.2011.07.003.
- Cha, J., Hwang, J.-M., Jo, H.J., Seo, S.W., Na, D.L., Lee, J.-M., 2015. Assessment of functional characteristics of amnesic mild cognitive impairment and Alzheimer's disease using various methods of resting-state fMRI analysis. *BioMed Res. Int.* 2015.
- Chen, S., DuBois Bowman, F., 2011. A novel support vector classifier for longitudinal high-dimensional data and its application to neuroimaging data. *Stat. Anal. Data Min.* 4 (6), 604–611. doi:10.1002/sam.10141.
- Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. *Am. J. Neuroradiol.* 21 (9), 1636–1644.
- Dai, N., Jones, G., 2017. Multivariate initial sequence estimators in Markov chain Monte Carlo. *J. Multivar. Anal.* 159, 184–199. doi:10.1016/j.jmva.2017.05.009.
- Damoiseaux, J., Beckmann, C., Arigita, E.S., Barkhof, F., Scheltens, P., Stam, C., Smith, S., Rombouts, S., 2007. Reduced resting-state brain activity in the “default network” in normal aging. *Cereb. Cortex* 18 (8), 1856–1864.
- Daniels, M., Hogan, J., 2008. Missing data in longitudinal studies: strategies for Bayesian modeling and sensitivity analysis. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.
- Diggle, P., Kenward, M.G., 1994. Informative drop-out in longitudinal data analysis. *J. R. Stat. Soc. Ser. C (Appl. Stat.)* 43 (1), 49–93.
- Euán, C., Ombao, H., Ortega, J., 2018. The Hierarchical Spectral Merger algorithm: a new time series clustering procedure. *J. Classif.* 35, 71–99.
- Eubank, R.L., 1999. Nonparametric regression and spline smoothing. *Statistics: A Series of Textbooks and Monographs*, 2, revised CRC Press.
- Fiecas, M., von Sachs, R., 2014. Data-driven shrinkage of the spectral density matrix of a high-dimensional time series. *Electron. J. Stat.* 8 (2), 2975–3003. doi:10.1214/14-EJS977.
- Flegal, J.M., Haran, M., Jones, G.L., 2008. Markov chain Monte Carlo: can we trust the third significant figure? *Stat. Sci.* 23 (2), 250–260. doi:10.1214/08-STS257.
- Flegal, J.M., Hughes, J., Vats, D., Dai, N., 2017. mcmcse: Monte Carlo standard errors for MCMC. Riverside, CA, Denver, CO, Coventry, UK, and Minneapolis, MN. R package version 1.3-2.
- van der Flier, W.M., Pijnenburg, Y.A., Fox, N.C., Scheltens, P., 2017. Early-onset versus late-onset Alzheimer's disease: the case of the missing APOE ϵ 4 allele. *Lancet Neurol.* 10 (3), 280–288. doi:10.1016/S1474-4422(10)70306-9.
- Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* 8, 700–711. doi:10.1038/nrn2201.
- Gelman, A., 2006. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Anal.* 1 (3), 515–534. doi:10.1214/06-BA117A.
- Geyer, C.J., 2011. Introduction to Markov chain Monte Carlo. In: Brooks, S., Gelman, A., Meng, X.-L., Jones, G.L. (Eds.), *Handbook of Markov Chain Monte Carlo*. CRC Press.
- Goldsmith, J., Crainiceanu, C.M., Caffo, B., Reich, D., 2012. Longitudinal penalized functional regression for cognitive outcomes on neuronal tract measurements. *J. R. Stat. Soc. Ser. C (Appl. Stat.)* 61 (3), 453–469. doi:10.1111/j.1467-9876.2011.01031.x.
- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc. Natl. Acad. Sci.* 100 (1), 253–258. doi:10.1073/pnas.0135058100. <http://www.pnas.org/content/100/1/253.full.pdf>.
- Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc. Natl. Acad. Sci. USA* 101 (13), 4637–4642.
- Hampson, M., Peterson, B.S., Skudlarski, P., Gatenby, J.C., Gore, J.C., 2002. Detection of functional connectivity using temporal correlations in MRI images. *Hum. Brain Map.* 15 (4), 247–262. doi:10.1002/hbm.10022.
- Han, Y., Wang, J., Zhao, Z., Min, B., Lu, J., Li, K., He, Y., Jia, J., 2011. Frequency-dependent changes in the amplitude of low-frequency fluctuations in amnesic mild cognitive impairment: a resting-state fMRI study. *NeuroImage* 55 (1), 287–295.
- Hart, B., Cribben, I., Fiecas, M., 2018. A longitudinal model for functional connectivity networks using resting-state fMRI. *NeuroImage* 178, 687–701. doi:10.1016/j.neuroimage.2018.05.071.
- Hedden, T., Gabrieli, J.D.E., 2004. Insights into the ageing mind: a view from cognitive neuroscience. *Nat. Rev. Neurosci.* 5 (2), 87–96.
- Holland, D., Desikan, R.S., Dale, A.M., McEvoy, L.K., for the Alzheimers Disease Neuroimaging Initiative, 2012. Rates of decline in Alzheimer's disease decrease with age. *PLOS ONE* 7 (8), 1–12. doi:10.1371/journal.pone.0042325.
- Hu, S., Chao, H.H., Zhang, S., Ide, J.S., Li, C.-S. R., 2014. Changes in cerebral morphometry and amplitude of low-frequency fluctuations of BOLD signals during healthy aging: correlation with inhibitory control. *Brain Struct. Funct.* 219 (3), 983–994.
- Hyun, J.W., Li, Y., Huang, C., Styner, M., Lin, W., Zhu, H., 2016. STGP: Spatio-temporal Gaussian process models for longitudinal neuroimaging data. *NeuroImage* 134, 550–562. doi:10.1016/j.neuroimage.2016.04.023.
- James, B.D., Caffo, B., Stewart, W.F., Yousem, D., Davatzikos, C., Schwartz, B.S., 2011. Genetic risk factors for longitudinal changes in structural MRI in former organolead workers. *J. Aging Res.* 2011. doi:10.4061/2011/362189.
- Kirch, C., Politis, D.N., 2011. TFT-bootstrap: resampling time series in the frequency domain to obtain replicates in the time domain. *Ann. Stat.* 39 (3), 1427–1470. doi:10.1214/10-AOS868.
- Little, R., Rubin, D., 2014. *Statistical analysis with missing data*. Wiley Series in Probability and Statistics. Wiley.
- Liu, X., Wang, S., Zhang, X., Wang, Z., Tian, X., He, Y., 2014. Abnormal amplitude of low-frequency fluctuations of intrinsic brain activity in Alzheimer's disease. *J. Alzheimer's Dis.* 40 (2), 387–397. doi:10.3233/JAD-131322.
- Long, Z., Jing, B., Yan, H., Dong, J., Liu, H., Mo, X., Han, Y., Li, H., 2016. A support vector machine-based method to identify mild cognitive impairment with multi-level characteristics of magnetic resonance imaging. *Neuroscience* 331, 169–176.
- Lowe, M., Mock, B., Sorenson, J., 1998. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *NeuroImage* 7 (2), 119–132.

- Molenberghs, G., Kenward, M., 2007. *Missing data in clinical studies*. *Statistics in Practice*. Wiley.
- Neal, R.M., 2011. MCMC using Hamiltonian dynamics. In: Brooks, S., Gelman, A., Meng, X.-L., Jones, G.L. (Eds.), *Handbook of Markov chain Monte Carlo*. CRC Press.
- Ombao, H.C., Raz, J.A., Strawderman, R.L., von Sachs, R., 2001. A simple generalised crossvalidation method of span selection for periodogram smoothing. *Biometrika* 88 (4), 1186–1192.
- Onoda, K., Ishihara, M., Yamaguchi, S., 2012. Decreased functional connectivity by aging is associated with cognitive decline. *J. Cognit. Neurosci.* 24 (11), 2186–2198.
- R Core Team, 2017. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ren, P., Lo, R.Y., Chapman, B.P., Mapstone, M., Porsteinsson, A., Lin, F., 2016. Longitudinal alteration of intrinsic brain activity in the striatum in mild cognitive impairment. *J. Alzheimer's Dis.* 54 (1), 69–78. doi:10.3233/JAD-160368.
- Rosen, O., Wood, S., Stoffer, D.S., 2012. AdaptSPEC: adaptive spectral estimation for nonstationary time series. *J. Am. Stat. Assoc.* 107 (500), 1575–1589. doi:10.1080/01621459.2012.716340.
- Sala-Llonch, R., Bartrés-Faz, D., Junqué, C., 2015. Reorganization of brain networks in aging: a review of functional connectivity studies. *Front. Psychol.* 6, 663. doi:10.3389/fpsyg.2015.00663.
- Shumway, R., Stoffer, D., 2011. *Time Series Analysis and Its Applications*. Springer.
- Skup, M., 2010. Longitudinal fMRI analysis: a review of methods. *Stat. Interface* 3 (2), 232–252.
- Skup, M., Zhu, H., Zhang, H., 2012. Multiscale adaptive marginal analysis of longitudinal neuroimaging data with time-varying covariates. *Biometrics* 68 (4), 1083–1092. doi:10.1111/j.1541-0420.2012.01767.x.
- Vats, D., Flegal, J.M., Jones, G.L., 2018. Strong consistency of multivariate spectral variance estimators in Markov chain Monte Carlo. *Bernoulli* 24 (3), 1860–1909. doi:10.3150/16-BEJ914.
- Vats, D., M. Flegal, J., Jones, G., 2015. *Multivariate output analysis for Markov chain Monte Carlo*.
- Vidoni, E.D., Thomas, G.P., Honea, R.A., Loskutova, N., Burns, J.M., 2012. Evidence of altered corticomotor system connectivity in early-stage Alzheimer's disease. *J. Neurol. Phys. Ther.* 36 (1), 8.
- Wang, K., Liang, M., Wang, L., Tian, L., Zhang, X., Li, K., Jiang, T., 2007. Altered functional connectivity in early Alzheimer's disease: a resting-state fMRI study. *Hum. Brain Map.* 28 (10), 967–978. doi:10.1002/hbm.20324.
- Whittle, P., 1957. Curve and periodogram smoothing. *J. R. Stat. Soc. Ser. B (Methodol.)* 19 (1), 38–63.
- Zang, Y.-F., He, Y., Zhu, C.-Z., Cao, Q.-J., Sui, M.-Q., Liang, M., Tian, L.-X., Jiang, T.-Z., Wang, Y.-F., 2007. Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI. *Brain Dev.* 29 (2), 83–91. doi:10.1016/j.braindev.2006.07.002.
- Zhang, Z.-Q., Lu, G.-M., Zhong, Y., Tan, Q.-F., Zhu, J.-G., Jiang, L., Chen, Z.-L., Wang, Z.-Q., Shi, J.-X., Zang, Y.-F., Liu, Y.-J., 2008. Application of amplitude of low-frequency fluctuation to the temporal lobe epilepsy with bilateral hippocampal sclerosis: an fMRI study. *Zhonghua Yi Xue Za Zhi* 88 (23), 1594–1598.
- Zipunnikov, V., Greven, S., Shou, H., Caffo, B.S., Reich, D.S., Crainiceanu, C.M., 2014. Longitudinal high-dimensional principal components analysis with application to diffusion tensor imaging of multiple sclerosis. *Ann. Appl. Stat.* 8 (4), 2175–2202. doi:10.1214/14-AOAS748.
- Zou, Q.-H., Zhu, C.-Z., Yang, Y., Zuo, X.-N., Long, X.-Y., Cao, Q.-J., Wang, Y.-F., Zang, Y.-F., 2008. An improved approach to detection of amplitude of low-frequency fluctuation (ALFF) for resting-state fMRI: fractional ALFF. *J. Neurosci. Methods* 172 (1), 137–141.
- Zuo, X.-N., Di Martino, A., Kelly, C., Shehzad, Z.E., Gee, D.G., Klein, D.F., Castellanos, F.X., Biswal, B.B., Milham, M.P., 2010. The oscillating brain: complex and reliable. *NeuroImage* 49 (2), 1432–1445.