

A Bayesian Spatial Model for Imaging Genetics

Yin Song¹, Shufei Ge², Jiguo Cao³, Liangliang Wang³, Farouk S. Nathoo^{1,*}

¹Mathematics and Statistics, University of Victoria, British Columbia, PO BOX 1700 STN CSC, Canada.

²Institute of Mathematical Sciences, ShanghaiTech University, Shanghai, 201210, China.

³Statistics and Actuarial Science, Simon Fraser University, British Columbia, V5A 1S6, Canada.

* to whom correspondence should be addressed.

**email*: nathoo@uvic.ca

SUMMARY: We develop a Bayesian bivariate spatial model for multivariate regression analysis applicable to studies examining the influence of genetic variation on brain structure. Our model is motivated by an imaging genetics study of the Alzheimer's Disease Neuroimaging Initiative (ADNI), where the objective is to examine the association between images of volumetric and cortical thickness values summarizing the structure of the brain as measured by magnetic resonance imaging (MRI) and a set of 486 SNPs from 33 Alzheimer's Disease (AD) candidate genes obtained from 632 subjects. A bivariate spatial process model is developed to accommodate the correlation structures typically seen in structural brain imaging data. First, we allow for spatial correlation on a graph structure in the imaging phenotypes obtained from a neighbourhood matrix for measures on the same hemisphere of the brain. Second, we allow for correlation in the same measures obtained from different hemispheres (left/right) of the brain. We develop a mean-field variational Bayes algorithm and a Gibbs sampling algorithm to fit the model. We also incorporate Bayesian false discovery rate (FDR) procedures to select SNPs. We implement the methodology in a new release of the R package *bgsmt*. We show that the new spatial model demonstrates superior performance over a standard model in our application. Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (<https://adni.loni.usc.edu>).

KEY WORDS: Bayesian Model, Gibbs Sampling, Imaging Genetics, Spatial Model, Variational Bayes

1. Introduction

We consider multivariate multiple regression modeling within the context of imaging genetics where interest lies in uncovering the associations between genetic variations and neuroimaging measures as quantitative traits (QTs). This problem has received a great deal of attention recently and is challenging because it combines the analysis of neuroimaging data with genetic data (see e.g., Vounou et al., 2010; Stein et al., 2010; Silver et al., 2010; Inkster et al., 2010; Hibar et al., 2011; Ge et al., 2012; Thompson et al., 2013; Stingo et al., 2013; Zhu et al., 2014; Hibar et al., 2015; Huang et al., 2015; Huang et al., 2017; Lu et al., 2017). Recent reviews of statistical issues in this area are discussed in Liu and Calhoun (2014) and Nathoo et al. (2018).

In our work, we consider the setting where interest lies in assessing the association between a moderate number of brain imaging phenotypes (e.g., 111 ROIs in Vounou et al., 2010; 12 ROIs in Wang et al., 2012; 93 ROIs in Zhu et al., 2014; 56 ROIs in Greenlaw et al., 2017) and with the number of SNPs ranging from between a few hundred to a few thousand. Within this setting a multivariate model with regression matrix jointly characterizing the associations between all ROIs and genetic markers is feasible.

Greenlaw et al. (2017) propose a Bayesian group sparse multi-task regression model where the primary focus is the use of a shrinkage prior based on a product of multivariate Laplace kernels developed following the ideas of Park and Casella (2008) and Kyung et al. (2010). The specific prior developed is motivated by the penalized multi-task regression estimator proposed by Wang et al. (2012). This development is an effort to move from point estimation to Bayesian credible intervals and fully Bayesian inference.

While these authors demonstrate the advantage of characterizing posterior uncertainty in their imaging genetics application to the ADNI study, their model makes a simplifying assumption for the covariance matrix of the imaging phenotypes, where the first level of the

model assumes:

$$\mathbf{y}_\ell | \mathbf{W}, \sigma^2 \stackrel{ind}{\sim} MVN_c(\mathbf{W}^T \mathbf{x}_\ell, \sigma^2 I_c) \quad \ell = 1, \dots, n, \quad (1)$$

where $\mathbf{y}_\ell = (\mathbf{y}_{\ell 1}, \dots, \mathbf{y}_{\ell c})^T$ denotes the vector of imaging phenotypes for subject ℓ and c is the dimension of the imaging phenotype, where $\ell = 1, \dots, n$; \mathbf{A}^T denotes transpose of matrix \mathbf{A} , \mathbf{W} is the regression matrix; $\mathbf{x}_\ell = (\mathbf{x}_{\ell 1}, \dots, \mathbf{x}_{\ell d})^T$, where \mathbf{x}_ℓ denotes the vector of genetic markers for subject ℓ and d is the number of such markers. The assumed covariance structure ignores spatial correlation as well as bilateral correlation across brain hemispheres.

We develop a new model that allows for this type of correlation by adopting a proper bivariate conditional autoregressive process (BCAR; see, e.g., Gelfand and Vounatsou, 2003; Jin et al., 2005) for the errors in the regression model. While spatial models for functional magnetic resonance imaging (fMRI) and other neuroimaging modalities have been developed to a large extent (see, e.g., Penny et al., 2005; Bowman, 2005; Bowman et al., 2008; Derado et al., 2013; Teng et al., 2018a; Teng et al., 2018b), to our knowledge there has been very little development of explicitly spatial models for imaging genetics. One exception is the mixture model developed by Stingo et al. (2013) where an Ising prior, a binary Markov random field, is used for Bayesian variable selection. Our model is rather different in both its aims and structure as it is based on a continuous bivariate Markov random field that is specified at the first level of the model for the imaging phenotype directly.

In Figure 1 we show several summaries of the data from our motivating application demonstrating the need to account for correlation across brain hemispheres. For example, the sample correlation between the volume of the right cerebral cortex and the volume of the left cerebral cortex is 0.90, and across all 28 pairs of measurements, the median left/right correlation between corresponding phenotypes is approximately 0.8. The bivariate CAR structure allows us to account for this between-hemisphere correlation while also allowing

us to account for within-hemisphere spatial correlation using a graph structure based on a neighbourhood matrix.

[Figure 1 about here.]

Typically, models incorporating multivariate CAR specifications are used for modelling observations (in the case of a proper CAR model) or spatially-varying parameters when multiple observations or parameters appear at each spatial site. For our application the use of this process is non-standard in the sense that we do not model multiple observations at each site, but rather, we pair corresponding observations on opposite hemispheres of the brain and use the bivariate spatial process to model a combination of the bilateral correlation across the left and right brain hemispheres as well as potential correlation within each hemisphere. As a matter of fact for the MRI data considered in our application the bilateral correlation is a very strong signal in the observed data and so this feature of the model is practically important.

For the bivariate spatial model we use a separable BCAR process as it is reasonable in our application to assume (as it might be in other neuroimaging studies) that the spatial covariance on the two hemispheres of the brain is similar. Non-separable multivariate spatial models (see, e.g., Gelfand and Banerjee, 2010; MacNab 2016) could be adopted for more flexibility allowing the spatial structure on the two hemispheres to be different; however, we do not expect that this additional flexibility would be useful in the current context. This spatial process is combined with a group Lasso prior for the regression coefficients, where each group corresponds to a single row of \mathbf{W} . Each row in this case represents the associations between a given SNP and all of the phenotypes.

To compute the posterior distribution we develop two algorithms, both of which are implemented in our R package *bgsmttr* for imaging genetics regression modelling. The package is available for download on the Comprehensive R Archive Network (CRAN). The first

algorithm is a Gibbs sampling algorithm and the second is a faster mean-field variational Bayes (VB) approximation to the posterior distribution (see e.g., Ormerod and Wand, 2010). Within the context of hierarchical models for spatial data, mean-field VB inference has been considered by Ren et al. (2011) who make comparisons with inference from MCMC within the context of spatial process models. In addition to the computation of the posterior distribution, the *bgsmttr* package now incorporates Bayesian FDR procedures (Morris et al., 2008) for SNP selection. This can be used alongside or as an alternative to SNP selection based on credible intervals.

The overall contribution of our work is four-fold. First, we develop an explicitly spatial model for imaging genetics based on the BCAR process. Second, we develop both an MCMC algorithm and a mean-field VB algorithm for approximating the posterior distribution. Third, we incorporate Bayesian FDR procedures for SNP selection within the new spatial model. Fourth, our new developments are implemented in the latest version of the *bgsmttr* R package.

The remainder of this paper is structured as follows. In Section 2, we present our new spatial model for imaging genetics. Section 3 presents a simulation study evaluating the performance of the spatial model relative to a non-spatial model and inference based on MCMC relative to that from VB. In Section 4 we apply our new model to our motivating application examining data from the ADNI-1 study, examining 56 structural brain imaging phenotypes, 486 SNPs from 33 genes, and 632 subjects. The paper concludes with a discussion in Section 5.

2. Bayesian Spatial Regression Model

Let $\mathbf{y}_\ell = (\mathbf{y}_{\ell,1}, \dots, \mathbf{y}_{\ell,c})^T$ and $\mathbf{x}_\ell = (\mathbf{x}_{\ell,1}, \dots, \mathbf{x}_{\ell,d})^T$ denote the imaging measures at c ROIs and the genetic data respectively for subject ℓ , $\ell = 1, \dots, n$, where $\mathbf{x}_{\ell,j} \in \{0, 1, 2\}$ represents the number of minor alleles of the j_{th} SNP for subject ℓ . The regression model takes the form $E(\mathbf{y}_\ell) = \mathbf{W}^T \mathbf{x}_\ell$, $\ell = 1, \dots, n$, where \mathbf{W} has dimensions $d \times c$ and $W_{i,j}$ represents the association between the i_{th} SNP and the j_{th} imaging phenotype.

Our model is developed for settings where the imaging data are symmetric with the same measures collected on each hemisphere of the brain. This is true when the neuroimaging data are considered at the voxel level and it is also the case for the study considered here where we analyze MRI data from the ADNI-1 database preprocessed using the FreeSurfer V4 software (Fischl, 2012). We conduct automated parcellation to define volumetric and cortical thickness values from the 28 ROIs considered in Shen et al. (2010), Szefer et al. (2017), and Greenlaw et al. (2017) on each hemisphere leading to $c = 56$ brain measures in total.

As described in Szefer et al. (2017), potential confounders in the analysis are population stratification and APOE genotype. Since true population structure is not observed, a set of principal coordinates from multidimensional scaling are used to derive proxy variables for population stratification in the data. We also adjust for APOE genotype, since it can account for the population stratification in the data, over and above the principal components or principal coordinates (Lucotte et al. 1997).

The response imaging measures at each brain ROI are first adjusted for the ten principal coordinates, as well as for dummy variables representing APOE genotype, using weighted ordinary least squares regression. The residuals from each regression are then used as the adjusted neuroimaging phenotypes (Szefer et al., 2017).

Let $\mathbf{y}_{\ell,i} = (y_{i,i}^{(L)}, y_{i,i}^{(R)})'$ be the brain summary measures obtained at the i_{th} ROI in the left hemisphere (L) and the right hemisphere (R). Then $\mathbf{y}_{\ell} = (\mathbf{y}'_{\ell,1}, \dots, \mathbf{y}'_{\ell,c/2})'$ is the imaging data ordered so that left-right imaging phenotype pairs are adjacent in the response vector. There are thus $c/2$ ROIs on each hemisphere and we let \mathbf{A} denote a $c/2 \times c/2$ symmetric neighborhood matrix which in the simplest case can have binary elements, where $A_{i,j} = 1$ indicates that ROI i and j are neighbors $i \neq j$, or more generally $A_{i,j} \geq 0$ and $A_{i,i} = 0$, $i = 1, \dots, c/2$. A user input is thus a neighborhood matrix \mathbf{A} or in the absence of user input our software implementation takes $A_{i,j}$ to be the average of the absolute value of the sample

correlation between phenotype/ROI i and phenotype/ROI j , where the average is taken over left/right hemisphere. The regression model then takes the form

$$\mathbf{y}_\ell = \mathbf{W}^T \mathbf{x}_\ell + \boldsymbol{\epsilon}_\ell \quad (2)$$

and the model for the errors $\boldsymbol{\epsilon}_\ell$ is a mean-zero multivariate normal distribution of dimension c , which can be specified through a set of $c/2$ compatible bivariate conditional distributions for $\boldsymbol{\epsilon}_{l,i} = (\epsilon_{l,i}^{(L)}, \epsilon_{l,i}^{(R)})'$, specified as follows:

$$\boldsymbol{\epsilon}_{l,i} | \boldsymbol{\epsilon}_{l(-i)}, \rho, \boldsymbol{\Sigma} \sim \text{BVN}\left(\frac{\rho}{A_i} \sum_{j=1}^{c/2} A_{i,j} \boldsymbol{\epsilon}_{l,j}, \frac{1}{A_i} \boldsymbol{\Sigma}\right)$$

where, $\boldsymbol{\epsilon}_{l(-i)}$ denotes the rest after removing $\boldsymbol{\epsilon}_{l,i}$ from $\boldsymbol{\epsilon}_l$, $A_i = \sum_{j=1}^{c/2} A_{i,j}$, $\rho \in [0, 1)$ characterizes spatial dependence with $\rho = 0$ corresponding to independence across all ROI pairs and $\boldsymbol{\Sigma}$ is a 2×2 matrix where $\kappa = \Sigma_{12} / \sqrt{\Sigma_{11}\Sigma_{22}} \in (-1, 1)$ quantifies within pair dependence, and with $\kappa = 0$ corresponding to independence across hemispheres.

As far as we are aware, this spatial model for neuroimaging data is one of the first to explicitly model dependence across brain hemispheres in addition to accounting for local dependence. Often, this left/right bilateral dependence is ignored with neuroimaging data. In the ADNI-1 data it is a very clear and strong signal as is evident in Figure 1. As the parameter $\boldsymbol{\Sigma}$ is free in the model it can be informed by the data. Therefore, we expect that the posterior will reflect some degree of between hemisphere correlation when it is present to a sufficient degree, and will remain at roughly a diagonal form when it is not, in diseases that have major differences across hemispheres, for example. The prior for $\boldsymbol{\Sigma}$ is chosen so that it is centered on a diagonal matrix.

Under this new specification the first level of the regression model takes the following form:

$$\mathbf{y}_\ell | \mathbf{W}, \boldsymbol{\Sigma} \stackrel{ind}{\sim} \text{MVN}_c(\mathbf{W}^T \mathbf{x}_\ell, (\mathbf{D}_A - \rho \mathbf{A})^{-1} \otimes \boldsymbol{\Sigma}), \quad l = 1, \dots, n, \quad (3)$$

where MVN_c denotes a c -dimensional multivariate normal distribution, \otimes is the kronecker product, $\mathbf{D}_A = \text{diag}\{A_i, i = 1, \dots, c/2\}$ and as before $A_i = \sum_{j=1}^{c/2} A_{i,j}$. For the regression

coefficients, we let $\tilde{W}_{i,j^*} = (W_{i,j}, W_{i,j+1})$, $j = 2j^* - 1$, $j^* = 1, \dots, \frac{c}{2}$, and we adopt a shrinkage prior based on a bivariate Gaussian scale mixture $\tilde{W}_{i,j^*} | \omega_i^2, \Sigma \stackrel{ind}{\sim} \text{BVN}(\mathbf{0}, \omega_i^2 \Sigma)$, $\omega_i^2 | \lambda^2 \stackrel{iid}{\sim} \text{Gamma}(\frac{c+1}{2}, \lambda^2/2)$, $\Sigma \sim \text{Inv-Wishart}(v, \mathbf{S})$, where ρ and λ^2 are tuning parameters controlling spatial dependence and regression sparsity respectively. Through the bivariate mixture, this prior also encourages shrinkage across brain hemispheres.

Posterior computation can be implemented using Gibbs sampling. The derivation and update steps for this algorithm are listed in Algorithm 1 in the Supplementary Material (Web Appendix A). As a faster alternative approach to computing the posterior distribution, we also develop a mean-field VB algorithm that is based on maximizing an evidence lower bound (ELBO). The derivations and VB algorithm are presented as Algorithm 2 in the Supplementary Material (Web Appendix A). To initialize the variational Bayes algorithm, we use a ridge regression estimator obtained separately for each column of \mathbf{W} obtained by fitting ridge regression with individual scalar-valued phenotypes as the response. The output of VB is then used to initialize the MCMC sampling algorithm. After the algorithms have converged, we base SNP selection on the Bayesian FDR procedure developed in Morris et al. (2008) based on a posterior probability cut-off value of $c^* > 0$. Adapted to our model, the procedure is described in the Supplementary Material (Web Appendix B). To compare the spatial and non-spatial models and to choose values for the tuning parameters, one option is the use of the WAIC (Vehtari et al., 2017) which can be computed from posterior samples, and with lower values being preferred. An alternative approach to tuning the model is to use a simple moment estimator based on the ridge regression estimator. This estimator implicitly uses cross-validation (CV) in the sense that CV is used to tune the ridge estimator and is much faster than applying CV directly to our Bayesian model. The details for model tuning and selection are provided in Supplementary Material (Web Appendix C).

3. Simulation Studies

We conduct a simulation study where the dimension of the data is the same as that in our motivating application with 56 structural brain imaging phenotypes, 486 SNPs from 33 genes, and 632 subjects. The data are simulated from the spatial model with parameter values set at the estimates obtained from the real data.

We use 100 simulation replicates and run the MCMC algorithms for both the spatial model and the non-spatial model for 10000 iterations with 5000 iterations used as burn-in for each replicate, and we run the variational Bayes algorithm to convergence of the ELBO, where convergence is declared when the absolute relative change in the ELBO is smaller than $\epsilon = 10^{-4}$ for $K = 2$ consecutive iterations. The tuning parameters are chosen using the WAIC over a coarse grid for the spatial model implemented via MCMC; set using the default approach in the R package for the model of Greenlaw et al. (based on five-fold cross-validation over a grid of possible values with out-of-sample prediction based on an approximate posterior mode, the estimator of Wang et al. (2012), that can be computed quickly) and using our moment based approach to obtain λ^2 for VB, while the spatial parameter is held fixed at the default value of $\rho = 0.95$ for VB (the true value set when simulating the data is $\rho = 0.8$).

Five measures are used for comparison in this study and are listed in Table 1. For the best overall estimation performance, our spatial model implemented using MCMC has an average mean-squared error (MSE) of 0.0055, followed by the non-spatial model with an average MSE of 0.012 and finally the spatial model implemented with VB with an average MSE of 0.043. Here, the average is taken over the 27,216 regression coefficients in the model. Next, we examine the correlation between the posterior mean estimates of $\text{vec}(\mathbf{W})$ and the true values averaged over simulation replicates. The spatial model implemented using MCMC has the best performance (0.69), followed by the non-spatial model implemented via MCMC (0.67) and then the spatial model implemented using VB (0.65). The average-squared bias

is lowest for the non-spatial model (3.29×10^{-4}), followed by the spatial model implemented with MCMC (4.70×10^{-3}) and finally the spatial model implemented using VB (9.10×10^{-3}) has the highest average-squared bias. Both the non-spatial and spatial models appear to exhibit adequate coverage probability overall when implemented with MCMC while the variational Bayes implementation of the spatial model exhibits coverage of 95% equal-tail credible intervals that is slightly lower at 0.91. The average posterior standard deviation obtained from the spatial model (MCMC) is slightly larger than that obtained from the non-spatial model perhaps indicating that the spatial model is adequately accounting for the dependence. As is typical with VB, the posterior standard deviation is underestimated. This arises from the Kullback-Leibler objective function that VB is based on which under-penalizes approximations that are under-dispersed.

To evaluate the empirical FDR of the Bayesian FDR approach when the expected Bayesian FDR is controlled with $\alpha = 0.05$, leading to the threshold for posterior probabilities $\phi_{\alpha=0.05}$, we conduct an additional 100 simulation replicates where the error structure underlying the simulated data is again set to that obtained from the real data application. All but 100 SNPs have their 56 coefficients set to exactly zero, while 50 rows of \mathbf{W} have coefficients set to $W_{ij} = 1$, 25 rows have coefficients set to $W_{ij} = 2$ and the remaining 25 rows have coefficients set to $W_{ij} = 3$. Here we evaluate the empirical FDR for the spatial model and both of its implementations. Both the VB and MCMC algorithms use the moment estimate for λ^2 . The empirical proportion of false discoveries averaged over simulation replicates is reported in the Supplementary Material (Web Appendix D) for both VB and MCMC as a function of the threshold c^* . We note that these c^* values are based on each simulated response matrix \mathbf{Y} having its columns centered and scaled. The empirical FDR appears to drop to 0 with increasing c^* faster with the MCMC implementation which is consistent with the VB implementation being more liberal in selecting SNPs.

Finally, we have conducted a simulation study investigating the performance of the spatial model under a misspecified model setting where data are generated without spatial correlation or bilateral correlation. The results of this study are reported in the Supplementary Material (Web Appendix E) and show that the performance of our model is adequate in this case, though as expected the non-spatial model outperforms our proposed model in this setting both with respect to estimation and also with respect to SNP selection and true and false positive rates summarized with ROC curves in the Supplementary Material (Figure 1).

4. ADNI-1 Study of MRI and Genetics

We apply our spatial model as well as the group sparse multi-task regression model of Greenlaw et al. (2017) to MRI and genetic data collected from $n = 632$ subjects from the ADNI-1 database. The genetic data comprise SNPs belonging to the top 40 Alzheimer's Disease (AD) candidate genes listed on the AlzGene database as of June 10, 2010. The data presented here are queried from the genome build as of December 2014, from the ADNI-1 data. After quality control and imputation steps, the genetic data used for this study include 486 SNPs from the 33 targeted genes discussed in Szefer et al. (2017). The freely available software package PLINK (Purcell et.al., 2007) is used for genomic quality control. Subjects are included if their genotyping data is available, they have a baseline MRI scan, and they have at least one additional follow-up baseline scan. Among all SNPs, only SNPs belonging to the top 40 AD candidate genes listed on the AlzGene database (<https://www.alzgene.org>) as of June 10, 2010, are selected after the standard quality control (QC) and imputation steps. The QC criteria for the SNP data include (i) call rate check per subject and per SNP marker, (ii) gender check, (iii) sibling pair identification, (iv) the Hardy-Weinberg equilibrium test, (v) marker removal by the minor allele frequency and (vi) population stratification. Our thresholds for SNP and subject exclusion are the same as in Wang et al. (2012) with three exceptions. In the data quality control step, we used a stricter minimum call rate of 95% on

SNPs vs. Wang et al.’s call rate of 90%. To assign SNPs to genes, we use a genome build (Build GRCh38.p2) from December 2014 whereas these authors use a genome build (Build 36.2) from September 2006, and use all subjects with a baseline measurement whereas we choose subjects with a baseline MRI scan and a scan at at least one additional time point in the longitudinal study.

The response measures are obtained by preprocessing the MRI data using the FreeSurfer V4 software which conducts automated parcellation to define volumetric and cortical thickness values from the 28 ROIs considered in Szefer et al. (2017) and Greenlaw et al. (2017) on each hemisphere of the brain, leading to $c = 56$ brain measures in total. These ROIs are chosen based on prior knowledge that they are related to Alzheimer’s Disease. Each of the response variables are adjusted for age, gender, education, handedness, baseline total intracranial volume (ICV), potential population stratification and APOE genotype and centered to have zero-sample-mean and unit-sample-variance.

We fit our spatial model to these data using both Gibbs sampling and VB. In addition, we fit the non-spatial model of Greenlaw et al. (2017) using the MCMC sampler derived therein, with the tuning parameters for the non-spatial model set at $\lambda_1^2 = 1000$ and $\lambda_2^2 = 1000$ based on the values selected in Greenlaw et al. (2017). In all cases, MCMC sampling is run for 10,000 iterations with the initial 5,000 iterations discarded. The required computation time for the spatial model (MCMC) is 50 hours on a single core (2.66-GHz Xeon x5650) with 20GB of RAM, while the computation for the non-spatial model is 5hrs. Some trace plots and MCMC convergence diagnostics are presented in the Supplementary Material (Web Appendix F) and these demonstrate rapid convergence and good mixing of the MCMC sampling chains. The VB algorithm is run to convergence and requires 45 minutes with the ELBO converging in approximately 16 iterations. The tuning parameter λ^2 is set based on the moment estimator

and we set $\rho = 0.95$ for VB. The convergence of VB based on successive values of the ELBO is depicted in Figure 2.

[Figure 2 about here.]

The Supplementary Material (Figure 4) presents the WAIC computed for a number of different choices of the tuning parameters ρ and λ^2 . For the values of $\rho = 0.8$ and $\lambda^2 = 10,000$, the value of the WAIC is 83,170. While the WAIC is able to identify a value for ρ , it is monotone decreasing as a function of λ^2 . We will thus summarize the results over several values of λ^2 . The WAIC obtained for the non-spatial model with value of the tuning parameters used in Greenlaw et al. (2017) ($\lambda_1^2 = 1000$ and $\lambda_2^2 = 1000$) is 108,745.

Figure 3 presents the number of SNPs chosen by the spatial model for each ROI using Bayesian FDR (based on a critical value of $c^* = 0.044$) as a function of the tuning parameter λ^2 for both Gibbs sampling and VB. As expected, the curves are monotone decreasing but we note that their shapes differ when comparing the algorithms. In particular, VB selects a larger number of SNPs at all values of λ^2 . We suggest that the VB algorithm be used for obtaining starting values to initialize the MCMC as well as a tool to gain some initial insight into the data while the MCMC sampler runs to completion. This is useful because the MCMC sampler requires a relatively long run time, and the VB algorithm can be used initially (requiring 45 minutes in our study) while the MCMC sampler runs (requiring 50 hrs in our study).

[Figure 3 about here.]

For the values of the tuning parameters $\rho = 0.8$, $\lambda^2 = 10,000$, the average number of SNPs selected per ROI is 2, while more than half of the ROIs have no SNPs selected. In total, 75 SNPs across all 56 ROIs are selected and these are listed in the Supplementary Material (Web Appendix G) along with the corresponding phenotypes that they are associated with. With the VB approximation, 150 SNPs are selected, and the set of SNP-ROI pairs selected

by MCMC is a proper subset of the set selected by VB. In addition, the subset of SNPs and phenotypes also selected by the approach of Greenlaw et al. (2017) where the marginal posterior 95% credible interval is used for SNP selection are also highlighted in bold in Table 3 of the Supplementary Material.

Considering all three approaches, a consistent signal is found at the APOE gene, where all three methods select SNP rs405509 and the Bayesian FDR procedure selects this SNP for four phenotypes, namely, the middle temporal gyrus thickness on the right hemisphere, supramarginal gyrus thickness on the right hemisphere, mean thickness of the caudal mid-frontal, rostral midfrontal, superior frontal, lateral orbitofrontal, and medial orbitofrontal gyri and frontal pole on the right hemisphere, and the mean thickness of the inferior temporal, middle temporal, and superior temporal gyri on the right hemisphere. We note that the selected associations for this SNP all correspond to ROIs in the right brain hemisphere. The associations between the genetic signal represented in our analysis by APOE SNP rs405509 with phenotypes on the right hemisphere of the brain may be of potential interest for further investigation.

The associated point estimates and 95% equal-tail credible intervals for all 56 phenotypes and APOE SNP rs405509 are presented in the Supplementary Material (Web Appendix H) and a subset of these results for 26 phenotypes is presented in Table 2. Overall, there is broad agreement based on the overlap of the interval estimates, though, we also see a number of examples where the 95% interval estimate obtained from the spatial model includes the value 0 while the corresponding interval estimate for the non-spatial model does not include zero. The higher bias arising from VB in the simulation studies may also be apparent in the results presented in Table 2. The interval estimates arising from VB appear to be slightly more narrow than those obtained from MCMC, but not to a large degree, and this may also be reflected in the selection rates depicted in Figure 3 which are slightly higher.

Another consistent signal is found at the ACE gene with SNP rs4311, which is found associated with 12 ROIs. We note that all but one of these ROIs is in the right hemisphere, and three of these ROIs (all of which are in the right hemisphere) are in common with the ROIs selected for this SNP by Greenlaw et al. (2017). In Figure 4 we indicate the SNPs chosen for each ROI, where the SNPs are grouped on the x-axis by gene and the ROIs are grouped in left/right pairs on the y-axis. The selected SNPs for each ROI are shown for tuning parameter values $\lambda^2 = 1000$ and $\lambda^2 = 10,000$. In both cases the value of the spatial tuning parameter for the CAR model is set at $\rho = 0.8$ as suggested by Supplementary Material Figure 4.

[Figure 4 about here.]

Examining Figure 4, two ROIs stand out as having a relatively broad genetic signal that persists even as the tuning parameter increases from $\lambda^2 = 1000$ to $\lambda^2 = 10,000$. These are *Left-Supramarg* (thickness of the left supramarginal gyrus) and *Left-SupTemporal* (thickness of the left superior temporal gyrus). For the case where $\lambda^2 = 1000$, phenotype *Left-Supramarg* is associated with 188 SNPs (top panel of Figure 4) and this decreases to 24 SNPs (bottom panel of Figure 4) when $\lambda^2 = 10,000$. When $\lambda^2 = 1000$ phenotype *Left-SupTemporal* is associated with 188 SNPs and this decreases to 23 SNPs when $\lambda^2 = 10,000$. This is to be compared with the average number of SNPs selected over all ROIs when $\lambda^2 = 10,000$ which is just 2. The regularization paths for these two regions are shown in the Supplementary Material (Figure 5) where the SNPs having the most persistent signal across values of λ^2 are highlighted. These include SNP rs10868609 for the thickness of the left supramarginal gyrus and rs10501426 for the thickness of the left superior temporal gyrus.

5. Conclusion

We have developed a spatial multi-task regression model for relating genetic data to imaging phenotypes. The error structure for the imaging phenotype is based on a computationally convenient proper bivariate conditional autoregressive model, which allows us to account for bilateral correlation across brain hemispheres while also allowing us to account for within-hemisphere correlation using a graph structure. In our case it is clear that a model accounting for correlation is indeed required (see, e.g., Figure 1) for bilateral correlation. Examination of alternative spatial models and comparisons with respect to estimation, inference and robustness relative to that currently implemented under different settings will be part of future work. Future work investigate more general settings under which the addition of spatial correlation will prove useful in comparison to fitting a non-spatial model.

Future developments may also incorporate the Bayesian false discovery probability proposed by Wakefield et al. (2007) which allows the user to account explicitly for the cost of false discoveries and the cost of false non-discovery. With regards computation algorithms, we recommend that the approximate VB procedures be used to initialize the MCMC algorithm and also to obtain an initial insight into the data while the MCMC sampler runs. It appears that VB combined with Bayesian FDR tends to be more liberal in the selection of SNPs, and in our application the SNP-ROI pairs selected by MCMC + Bayesian FDR are a proper subset of that selected by VB + Bayesian FDR. Nevertheless, Figure 3 suggests that the approximation can be reasonable when good initializations are used.

ACKNOWLEDGEMENTS

Research is supported by funding from the Natural Sciences and Engineering Research Council of Canada, Canadian Statistical Sciences Institute and the Canada Research Chairs Program. Research was enabled in part by support provided by WestGrid and Compute Canada. Data collection and sharing for this project was funded by the Alzheimer's Disease

Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). The authors wish it to be known that the first two authors should be regarded as joint first authors.

REFERENCES

- Alzgene. The Alzgene Database. <http://www.alzgene.org/>. Accessed: 2016-07.
- CRAN. The Comprehensive R Archive Network. <https://cran.r-project.org/>. Accessed: 2020-07.
- Bowman, F. D. (2005). Spatio-temporal modeling of localized brain activity. *Biostatistics* **6**, 558–575.
- Bowman, F. D., Caffo, B., Bassett, S. S., and Kilts, C. (2008). A bayesian hierarchical framework for spatial modeling of fMRI data. *NeuroImage* **39**, 146–156.
- Corder, E. H., Saunders, A. M., Strittmatter, W. J., Schmechel, D. E., Gaskell, P. C., Small, G., Roses, A. D., Haines, J., and Pericak-Vance, M. A. (1993). Gene dose of apolipoprotein e type 4 allele and the risk of alzheimer’s disease in late onset families. *Science* **261**, 921–923.
- Derado, G., Bowman, F. D., Zhang, L., and Initiative, A. D. N. (2013). Predicting brain activity using a bayesian spatial model. *Statistical methods in medical research* **22**, 382–397.
- Fischl, B. (2012). Freesurfer. *Neuroimage* **62**, 774–781.
- Ge, T., Feng, J., Hibar, D. P., Thompson, P. M., and Nichols, T. E. (2012). Increasing power for voxel-wise genome-wide association studies: the random field theory, least square kernel machines and fast permutation procedures. *Neuroimage* **63**, 858–873.
- Gelfand, A. E. and Banerjee, S. (2010). Multivariate spatial process models. *Handbook of Spatial Statistics* pages 495–515.
- Gelfand, A. E. and Vounatsou, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* **4**, 11–15.

- Greenlaw, K., Szefer, E., Graham, J., Lesperance, M., Nathoo, F. S., and Initiative, A. D. N. (2017). A Bayesian group sparse multi-task regression model for imaging genetics. *Bioinformatics* **33**, 2513–2522.
- Hibar, D. P., Stein, J. L., Kohannim, O., Jahanshad, N., Saykin, A. J., Shen, L., Kim, S., Pankratz, N., Foroud, T., Huentelman, M. J., et al. (2011). Voxelwise gene-wide association study (vgenewas): multivariate gene-based association testing in 731 elderly subjects. *Neuroimage* **56**, 1875–1891.
- Hibar, D. P., Stein, J. L., Renteria, M. E., Arias-Vasquez, A., Desrivières, S., Jahanshad, N., Toro, R., Wittfeld, K., Abramovic, L., Andersson, M., et al. (2015). Common genetic variants influence human subcortical brain structures. *Nature* **520**, 224.
- Huang, C., Thompson, P., Wang, Y., Yu, Y., Zhang, J., Kong, D., Colen, R. R., Knickmeyer, R. C., Zhu, H., Initiative, A. D. N., et al. (2017). Fgwas: Functional genome wide association analysis. *Neuroimage* **159**, 107–121.
- Huang, M., Nichols, T., Huang, C., Yu, Y., Lu, Z., Knickmeyer, R. C., Feng, Q., Zhu, H., Initiative, A. D. N., et al. (2015). Fvgwas: Fast voxelwise genome wide association analysis of large-scale imaging genetic data. *Neuroimage* **118**, 613–627.
- Inkster, B., Nichols, T. E., Saemann, P. G., Auer, D. P., Holsboer, F., Muglia, P., and Matthews, P. M. (2010). Pathway-based approaches to imaging genetics association studies: Wnt signaling, gsk3beta substrates and major depression. *Neuroimage* **53**, 908–917.
- Jin, X., Carlin, B. P., and Banerjee, S. (2005). Generalized hierarchical multivariate car models for areal data. *Biometrics* **61**, 950–961.
- Kyung, M., Gill, J., Ghosh, M., Casella, G., et al. (2010). Penalized regression, standard errors, and bayesian lassos. *Bayesian Analysis* **5**, 369–411.
- Liu, J. and Calhoun, V. D. (2014). A review of multivariate analyses in imaging genetics.

- Frontiers in neuroinformatics* **8**, 29.
- Lu, Z.-H., Khondker, Z., Ibrahim, J. G., Wang, Y., Zhu, H., Initiative, A. D. N., et al. (2017). Bayesian longitudinal low-rank regression models for imaging genetic data from longitudinal studies. *Neuroimage* **149**, 305–322.
- Lucotte, G., Loirat, F., and Hazout, S. (1997). Pattern of gradient of apolipoprotein e allele* 4 frequencies in western europe. *Human biology* pages 253–262.
- MacNab, Y. C. (2016). Linear models of coregionalization for multivariate lattice data: a general framework for coregionalized multivariate car models. *Statistics in medicine* **35**, 3827–3850.
- Morris, J. S., Brown, P. J., Herrick, R. C., Baggerly, K. A., and Coombes, K. R. (2008). Bayesian analysis of mass spectrometry proteomic data using wavelet-based functional mixed models. *Biometrics* **64**, 479–489.
- Nathoo, F. S., Kong, L., and Zhu, H. (2018). A review of statistical methods in imaging genetics. *Canadian Journal of Statistics* page DOI: 10.1002/cjs.11487.
- Ormerod, J. T. and Wand, M. P. (2010). Explaining variational approximations. *The American Statistician* **64**, 140–153.
- Park, T. and Casella, G. (2008). The bayesian lasso. *Journal of the American Statistical Association* **103**, 681–686.
- Penny, W. D., Trujillo-Barreto, N. J., and Friston, K. J. (2005). Bayesian fmri time series analysis with spatial priors. *NeuroImage* **24**, 350–362.
- Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M. A., Bender, D., Maller, J., Sklar, P., De Bakker, P. I., Daly, M. J., et al. (2007). Plink: a tool set for whole-genome association and population-based linkage analyses. *The American Journal of Human Genetics* **81**, 559–575.
- Ren, Q., Banerjee, S., Finley, A. O., and Hodges, J. S. (2011). Variational Bayesian methods

- for spatial data analysis. *Computational Statistics & Data Analysis* **55**, 3197–3217.
- Scott, S. L., Blocker, A. W., Bonassi, F. V., Chipman, H. A., George, E. I., and McCulloch, R. E. (2016). Bayes and big data: The consensus monte carlo algorithm. *International Journal of Management Science and Engineering Management* **11**, 78–88.
- Shen, X., Papademetris, X., and Constable, R. T. (2010). Graph-theory based parcellation of functional subunits in the brain from resting-state fmri data. *Neuroimage* **50**, 1027–1035.
- Silver, M., Montana, G., Nichols, T. E., Initiative, A. D. N., et al. (2011). False positives in neuroimaging genetics using voxel-based morphometry data. *Neuroimage* **54**, 992–1000.
- Stein, J. L., Hua, X., Lee, S., Ho, A. J., Leow, A. D., Toga, A. W., Saykin, A. J., Shen, L., Foroud, T., Pankratz, N., et al. (2010). Voxelwise genome-wide association study (vgwas). *Neuroimage* **53**, 1160–1174.
- Stein, J. L., Medland, S. E., Vasquez, A. A., Hibar, D. P., Senstad, R. E., Winkler, A. M., Toro, R., Appel, K., Bartecek, R., Bergmann, Ø., et al. (2012). Identification of common variants associated with human hippocampal and intracranial volumes. *Nature genetics* **44**, 552.
- Stingo, F. C., Guindani, M., Vannucci, M., and Calhoun, V. D. (2013). An integrative bayesian modeling approach to imaging genetics. *Journal of the American Statistical Association* **108**, 876–891.
- Szefer, E., Lu, D., Nathoo, F., Beg, M. F., Graham, J., et al. (2017). Multivariate association between single-nucleotide polymorphisms in alzgene linkage regions and structural changes in the brain: discovery, refinement and validation. *Statistical applications in genetics and molecular biology* **16**, 367–386.
- Teng, M., Johnson, T. D., and Nathoo, F. S. (2018a). Time series analysis of fmri data: Spatial modelling and Bayesian computation. *Statistics in medicine* .
- Teng, M., Nathoo, F. S., and Johnson, T. D. (2018b). Bayesian analysis of functional

- magnetic resonance imaging data with spatially varying auto-regressive orders. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* .
- Thompson, P. M., Ge, T., Glahn, D. C., Jahanshad, N., and Nichols, T. E. (2013). Genetics of the connectome. *Neuroimage* **80**, 475–488.
- Vehtari, A., Gelman, A., and Gabry, J. (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing* **27**, 1413–1432.
- Vounou, M., Janousova, E., Wolz, R., Stein, J. L., Thompson, P. M., Rueckert, D., Montana, G., Initiative, A. D. N., et al. (2012). Sparse reduced-rank regression detects genetic associations with voxel-wise longitudinal phenotypes in alzheimer’s disease. *Neuroimage* **60**, 700–716.
- Vounou, M., Nichols, T. E., Montana, G., Initiative, A. D. N., et al. (2010). Discovering genetic associations with high-dimensional neuroimaging phenotypes: A sparse reduced-rank regression approach. *Neuroimage* **53**, 1147–1159.
- Wakefield, J. (2007). A Bayesian measure of the probability of false discovery in genetic epidemiology studies. *The American Journal of Human Genetics* **81**, 208–227.
- Wang, H., Nie, F., Huang, H., Risacher, S. L., Saykin, A. J., Shen, L., and Initiative, A. D. N. (2012). Identifying disease sensitive and quantitative trait-relevant biomarkers from multidimensional heterogeneous imaging genetics data via sparse multimodal multitask learning. *Bioinformatics* **28**, 127–136.
- Zhu, H., Khondker, Z., Lu, Z., and Ibrahim, J. G. (2014). Bayesian generalized low rank regression models for neuroimaging phenotypes and genetic markers. *Journal of the American Statistical Association* **109**, 977–990.

SUPPORTING INFORMATION

Web Appendices, Tables, and Figures referenced in Sections 2, 3 and 4 are available with this paper at the Biometrics website on Wiley Online Library. The R software package *bgsmttr*

implementing the methodology is available for download on The Comprehensive R Archive Network (CRAN) <https://cran.r-project.org/>.

Accepted Article

Table 1: Summaries from Simulation Study I. The table entries are average mean-squared error (MSE), average correlation (Corr.) between the posterior mean and the true value, average squared-bias (Bias²), average coverage probability (Cov. Prob.) of 95% equal-tail credible intervals, average posterior standard deviation (P. SD). The averages are taken over the 27,216 regression coefficients in all but 'Corr.' where the average of $Corr(\text{vec}(\hat{\mathbf{W}}), \text{vec}(\mathbf{W}_{true}))$ is taken over 100 simulation replicates and $\hat{\mathbf{W}} = E[\mathbf{W}|\mathbf{Y}]$.

Model	MSE	Corr.	Bias ²	Cov. Prob.	P. SD.
Non-Spatial MCMC	0.012	0.67	3.29×10^{-4}	0.96	0.087
Spatial Model MCMC	0.0055	0.69	4.70×10^{-3}	0.97	0.10
Spatial Model VB	0.043	0.65	9.10×10^{-3}	0.91	0.037

Table 2: ADNI-1 Study: posterior means and 95% equal-tail credible intervals for a subset of the ROIs and their association with APOE SNP rs405509.

Region	Spatial Model (MCMC)		Spatial Model (MFVB)		Non-Spatial Model (MCMC)	
	Mean	95% CI	Mean	95% CI	Mean	95% CI
Amygdala volume (L)	0.09	[-0.03,0.22]	0.17	[0.08,0.28]	0.12	[0.02,0.23]
Cerebral white matter volume (L)	0.09	[-0.04,0.22]	0.19	[0.10,0.28]	0.13	[0.03,0.23]
Inferior lateral ventricle volume (L)	-0.13	[-0.24,-0.01]	-0.14	[-0.24,-0.05]	-0.08	[-0.18,0.02]
Inferior parietal gyrus thickness (R)	0.11	[0.01,0.21]	0.20	[0.11,0.29]	0.13	[0.03,0.23]
Fusiform, parahippocampal, and lingual gyri, temporal pole and transverse temporal pole mean thickness (L)	0.12	[0.01,0.24]	0.22	[0.13,0.32]	0.14	[0.04,0.24]

Continued on next page

Table 2 – *Continued from previous page*

Region	Spatial Model (MCMC)		Spatial Model (MFVB)		Non-Spatial Model (MCMC)	
	Mean	95% CI	Mean	95% CI	Mean	95% CI
Inferior temporal, middle temporal, superior temporal, fusiform, parahippocampal, and lingual gyri, temporal pole and transverse temporal pole mean thickness (L)	0.11	[0.01,0.22]	0.20	[0.11,0.28]	0.13	[0.03,0.23]
Postcentral gyrus thickness (L)	0.13	[0.03,0.24]	0.23	[0.13,0.33]	0.14	[0.03,0.23]
Superior frontal gyrus thickness (R)	0.12	[0.02,0.22]	0.19	[0.11,0.28]	0.14	[0.04,0.24]
Supramarginal gyrus thickness (L)	0.13	[0.02,0.23]	0.21	[0.13,0.29]	0.14	[0.04,0.25]
Supramarginal gyrus thickness (R)	0.14	[0.04,0.24]	0.24	[0.17,0.33]	0.15	[0.05,0.25]
Fusiform gyrus thickness (L)	0.11	[-0.02,0.23]	0.19	[0.09,0.28]	0.13	[0.03,0.23]
Hippocampus volume (L)	0.12	[-0.02,0.26]	0.22	[0.14,0.34]	0.15	[0.05,0.26]
Caudal midfrontal, rostral midfrontal, superior frontal, lateral orbitofrontal, and medial orbitofrontal gyri and frontal pole mean thickness (L)	0.10	[-0.02,0.22]	0.18	[0.10,0.27]	0.13	[0.03,0.23]
Inferior temporal, middle temporal, and superior temporal gyri mean thickness (L)	0.10	[-0.01,0.21]	0.17	[0.08,0.26]	0.12	[0.02,0.22]
Inferior temporal, middle temporal, and superior temporal gyri mean thickness (R)	0.10	[-0.01,0.21]	0.20	[0.11,0.30]	0.13	[0.02,0.22]

Continued on next page

Table 2 – Continued from previous page

Region	Spatial Model (MCMC)		Spatial Model (MFVB)		Non-Spatial Model (MCMC)	
	Mean	95% CI	Mean	95% CI	Mean	95% CI
Fusiform, parahippocampal, and lingual gyri, temporal pole and transverse temporal pole mean thickness (R)	0.10	[-0.01,0.21]	0.20	[0.11,0.29]	0.12	[0.02,0.22]
Inferior and superior parietal gyri, supramarginal gyrus, and precuneus mean thickness (L)	0.09	[-0.02,0.21]	0.16	[0.07,0.25]	0.12	[0.02,0.22]
Precentral and postcentral gyri mean thickness (R)	0.09	[-0.02,0.20]	0.17	[0.09,0.26]	0.12	[0.01,0.22]
Precentral and postcentral gyri mean thickness (L)	0.10	[-0.01,0.20]	0.16	[0.08,0.25]	0.12	[0.02,0.22]
Middle temporal gyrus thickness (L)	0.08	[-0.03,0.19]	0.15	[0.06,0.23]	0.11	[0.01,0.21]
Middle temporal gyrus thickness (R)	0.09	[-0.01,0.20]	0.18	[0.11,0.27]	0.12	[0.02,0.22]
Postcentral gyrus thickness (R)	0.09	[-0.01,0.20]	0.18	[0.08,0.28]	0.11	[0.01,0.21]
Precentral gyrus thickness (R)	0.08	[-0.03,0.19]	0.15	[0.06,0.25]	0.11	[0.01,0.21]
Precuneus thickness (R)	0.09	[-0.01,0.19]	0.17	[0.08,0.25]	0.11	[0.01,0.21]
Superior parietal gyrus thickness (R)	0.08	[-0.02,0.18]	0.14	[0.06,0.22]	0.11	[0.01,0.21]
Superior temporal gyrus thickness (R)	0.10	[-0.01,0.22]	0.22	[0.14,0.31]	0.13	[0.03,0.23]

Received October xxx. Revised February xxx.

Accepted March xxx.

Figure 1: The correlation plot for left and right brain measures showing the bilateral correlation between 28 brain measurements across brain hemispheres for 632 subjects. Panel (a) presents the boxplot of all of the pairwise bilateral correlations for the 28 brain measurements across left and right hemispheres. Panel (b) presents a scatter plot showing mean thickness of left/right frontal with correlation 0.91. Panel (c) presents a scatter plot showing the volume of left/right cerebral cortex with correlation 0.90. Panel (d) presents a scatter plot between volume of left/right cerebral white matter with correlation 0.90.

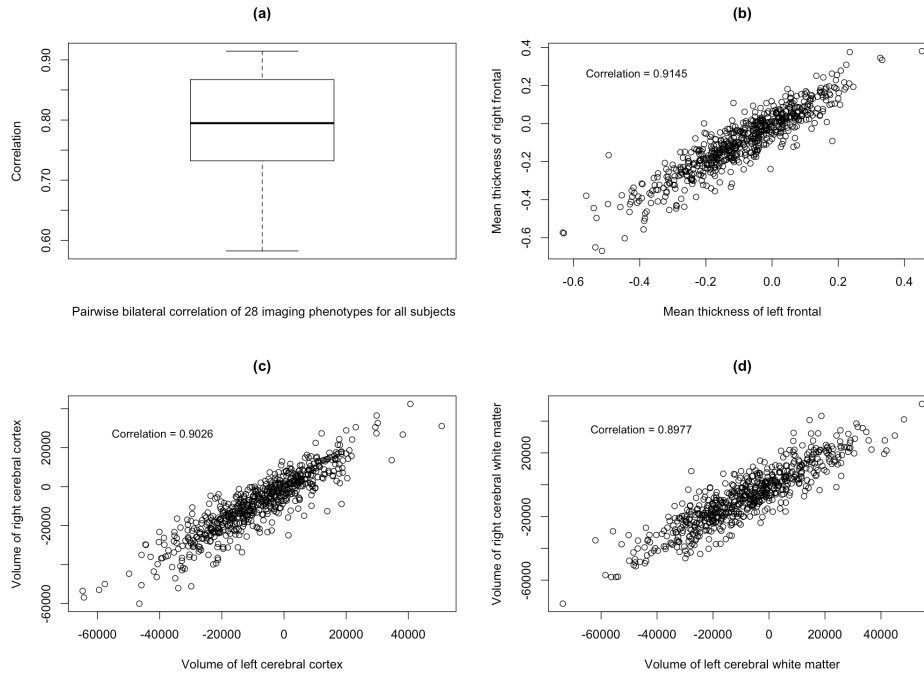


Figure 2: Variational Bayes - convergence of the evidence lower bound (ELBO) for the ADNI MRI and genetic data considered in the application.




Figure 3: ADNI-1 Data - Relationship between the number of selected SNPs in each region and λ^2 . Each region is represented with a curve in each panel of the figure. The left panel shows this relationship for MCMC combined with Bayesian FDR ($\alpha = 0.05$) while the right panel shows the same relationship for VB with Bayesian FDR ($\alpha = 0.05$).




Figure 4: ADNI-1 Data: SNPs chosen with the spatial model fit using Gibbs sampling and Bayesian FDR ($\alpha = 0.05$) are highlighted in red for each phenotype. The black ticks on y-axis indicate the phenotypes from the left/right hemisphere, and the SNPs from same gene are indicated by the ticks on x-axis. The top panel corresponds to the case $\lambda^2 = 1000$ while the bottom panel corresponds to the case $\lambda^2 = 10,000$.

