The Canadian Journal of Statistics
Vol. 00, No. 00, 2021, Pages 000–000
La revue canadienne de statistique

A Bayesian latent spatial model for mapping the cortical signature of progression to Alzheimer's disease

Ning DAI[1], Hakmook KANG[2], Galin L. JONES[1], and Mark B. FIECAS[3*] for the Alzheimer's Disease Neuroimaging Initiative[†]

[1]*School of Statistics, University of Minnesota, Minneapolis, MI 55455, U.S.A.*
[2]*Department of Biostatistics, Vanderbilt University, Nashville, TN 37203, U.S.A.*
[3]*Division of Biostatistics, University of Minnesota, Minneapolis, MI 55455, U.S.A.*

Key words and phrases: Alzheimer's disease; Bayesian modelling; longitudinal studies; spatial statistics; survival analysis.

MSC 2010: Primary 62N02; secondary 62P10.

Abstract: Prior studies have shown that atrophy in vulnerable cortical regions is associated with an increased risk of progression to clinical dementia. In this work, we utilize the longitudinal structural magnetic resonance imaging (MRI) data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) to investigate the relationship between the temporally changing spatial topography of cortical thickness and conversion from mild cognitive impairment to Alzheimer's disease (AD). We develop a novel Bayesian latent spatial model that employs the spatial information underlying the thickness effects across the cortical surface. The proposed method facilitates the development of imaging markers by reliably quantifying and mapping the regional vulnerability to AD progression across the cortical surface. Simulation results showed substantial gains in statistical power and estimation performance by accounting for the spatial structure of the association. Using MRI data from ADNI, we examined the topographic patterns of anatomic regions where cortical thinning is associated with an increased risk of developing AD. *The Canadian Journal of Statistics* 00: 000–000; 2021 © 2021 Statistical Society of Canada

Résumé: Des études précédentes ont montré que l'atrophie de régions corticales vulnérables est associée à une augmentation du risque de la progression de la démence clinique. Les auteurs utilisent les données d'imagerie par résonnance magnétique (IRM) longitudinales de l'initiative d'imagerie médicale pour la maladie d'Alzheimer (IIMMA) afin d'étudier la relation entre l'évolution dans le temps de la topographie spatiale de l'épaisseur corticale et la conversion de déficiences cognitives modérées (DCM) vers la maladie d'Alzheimer (MA). Ils développent un nouveau modèle latent spatial bayésien qui utilise l'information spatiale par rapport au cortex dans la mesure de l'effet lié à l'épaisseur corticale. La méthode proposée facilite le développement de marqueurs d'imagerie en quantifiant et en localisant de façon fiable des vulnérabilités régionales qui mènent à une progression vers la MA. Les auteurs présentent des résultats de simulations montrant des gains substantiels de puissance statistique et de performance pour l'estimation

---

Additional Supporting Information may be found in the online version of this article at the publisher's website.

**Author to whom correspondence may be addressed.*
E-mail: mfiecas@umn.edu

[†]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.

lorsque la structure spatiale de l'association est prise en compte. Avec les données d'IRM de l'IIMMA, ils examinent les motifs topographiques des régions anatomiques où l'amincissement cortical est associé à un risque accru de développer la MA. *La revue canadienne de statistique* 00: 000–000; 2021 © 2021 Société statistique du Canada

1. INTRODUCTION

Research on aging and dementia has identified a transitional state between normal cognition and Alzheimer's disease (AD), known as mild cognitive impairment (MCI) (Petersen et al., 2001). Individuals classified with MCI progress to AD at a considerably accelerated rate compared with healthy individuals (Petersen et al., 1999). Studying the risk factors associated with conversion to AD is of immense importance as therapeutic interventions for the delay or prevention of AD onset are developed (Galvin, 2017; Crous-Bou et al., 2017), and treatment trials are under way.

Neuroimaging has accelerated the study of AD-related clinical conditions as it enables researchers to investigate the temporal dynamics of the underlying biological mechanisms through serial imaging measurements. For example, structural magnetic resonance imaging (MRI) can visualise the degenerative changes caused by AD that occur before AD can be clinically diagnosed, and therefore can be used for predicting AD onset (Gomez-Isla, Price & Hyman, 1996). Prior studies suggest that analysis of the cortical thickness in key brain regions can be used to identify individuals at greatest risk for developing AD (Buckner et al., 2005; Dickerson et al., 2009). Cortical thickness analysis can quantify the atrophy attributable to AD along the cortical surface. The specific pattern of cortical thinning that appears to be associated with AD comprises the medial, inferior, and pole of the temporal lobe; the angular gyrus; the superior and inferior frontal lobe; the superior parietal lobule; the supramarginal gyrus; and the precuneus (Brun & Gustafson, 1976; Braak & Braak, 1991; Lerch et al., 2005; Dickerson et al., 2009, 2011).

Previous analyses of the relationship between cortical measurements and AD often used a case–control approach (Jack et al., 2012; Sabuncu et al., 2012) which may involve linear mixed modelling (Bernal-Rusiel et al., 2013). The case–control approach treats those who have developed AD and those who have not developed AD as distinct groups and assumes a certain amount of within-group homogeneity, an assumption which is often violated in Alzheimer's studies because MCI subjects who have not yet converted to AD are likely to develop AD in the not-too-distant future. Treating these high-risk subjects as controls can introduce bias into the statistical analysis. Also, available information in those who have yet to convert after the fixed follow-up time cannot be used in this type of analysis (Vemuri et al., 2011).

Instead of dichotomizing the MCI subjects into converters and non-converters, one can treat them as being in different stages of disease progression. Under this setup, survival analysis can be used to directly model the timing of AD onset, called survival time, while accounting for varying follow-up times and incorporating partial information in the form of censoring (Klein & Moeschberger, 2013). Various survival models, including the Cox proportional hazards model (Cox, 1972), proportional odds model, additive hazards model, accelerated failure time model, etc., have been employed in prior neuroimaging studies to demonstrate the association between cortical thinning and conversion to AD (Stoub et al., 2005; Devanand et al., 2007; Desikan et al., 2010; Vemuri et al., 2011; Lee et al., 2015). These prior studies relied on standard survival models that assumed time-invariant explanatory variables, and therefore were restricted to using MRI data obtained from a single visit and could not take into account the dynamic changes of the imaging measurements whenever the study had a longitudinal design. To use serial MRIs, we propose to use survival models that incorporate time-dependent covariates. The interested reader is referred to Fisher & Lin (1999), Ibrahim, Chen & Sinha (2001), Klein & Moeschberger (2013), and Klein et al. (2016) for examples of these methods.

Existing approaches for the spatial analysis of survival data related to human healthcare are not applicable to our particular dataset. The literature has been focused on modeling and analysis for geographically referenced survival data (Rushton, 2003; Banerjee, 2016; Zhou & Hanson, 2018), whereas each subject's time-to-conversion to AD does not vary across the cortical surface despite geographically referenced longitudinal imaging measurements. Recently, Sabuncu et al. (2014) developed an approach to investigating the association between space-invariant survival time and space-varying longitudinal imaging measurements based on survival analysis using each spatial location in the serial MRIs as a separate predictor. This type of analysis, which we call a massively univariate analysis (MUA), suffers from inefficient estimation that reduces statistical power due to the separate modeling of the covariance components at each spatial location.

The goal of this study is to develop a novel statistical method for identifying the topographic patterns of anatomic regions where cortical thinning is associated with an increased risk of developing AD. We propose a Bayesian latent spatial model that employs the spatial information underlying the thickness effects across the cortical surface. Specifically, we build on the Weibull model (Carroll, 2003; Zhou & Hanson, 2018) for survival analysis with time-dependent covariates to quantify the effects at each spatial location while borrowing information from nearby locations, effectively improving statistical power and stability.

The rest of the paper is organised as follows. In Section 2 we describe the novel Bayesian latent spatial model. Then, in Section 3 we carry out a simulation study to validate our approach. We apply our Bayesian latent spatial model to the Alzheimer's Disease Neuroimaging Initiative (ADNI) data in Section 4, and end with a discussion and conclusions in Section 5.

2. METHODS

2.1. Experimental Data and Pre-processing

We use the longitudinal clinical and structural MRI data from ADNI, consisting of 300 subjects with MCI at study entry. Specifically, our data comes from ADNIMERGE, which is a merged dataset containing ADNI 1/GO/2 clinical data. We use all the subjects in ADNIMERGE that meet the following three conditions: (i) The subject's MRI scan is available in the ADNI database; (ii) the subject was MCI, early MCI, or late MCI at study entry; (iii) the subject did not go back and forth during follow-up, that is, if a subject converted to AD and then went back to MCI at a later visit, the subject would be excluded from data analysis. Only 300 subjects meet these conditions.

The subjects are on average 74.7 years old at study entry with a standard deviation of 7.03 years. About 37.7% of them are female. The subjects fall into two categories: converters, that is, those who converted to clinical AD during the follow-up period; and non-converters, that is, those who did not convert to AD during follow-up. There are 181 converters, whose first diagnosis of AD is on average 2.51 years with a standard deviation of 1.95 years. For the remaining 119 MCI subjects, the average censoring time, that is, the final follow-up visit, is on average 4.32 years with a standard deviation of 2.86 years. Table 1 summarizes the demographic, clinical, and imaging characteristics of the participants. To illustrate the time-to-conversion from MCI to AD, Figure 1 presents the Kaplan–Meier estimator of the survival function.

We processed the MRI scans using FreeSurfer (version 6.0, Fischl, 2012), specifically its longitudinal stream (Reuter et al., 2012) to extract reliable cortical thickness estimates. FreeSurfer computes subject-specific thickness measurements across the cortical surface of each hemisphere. Each subject's data were resampled onto a standard surface-based template, called fsaverage. Then we applied minimal spatial smoothing using a Gaussian kernel with a standard deviation of 0.85 mm, which corresponds to a full-width at half-maximum (FWHM) of 2 mm. For computational efficiency, we downsampled the cortical thickness measurements by mapping the data onto the fsaverage4 template, which is a lower resolution version of fsaverage and has 2,562 vertices and 5,120 triangular faces in each hemisphere's surface mesh. Finally, we

TABLE 1: Demographic, clinical, and imaging characteristics of the participants at study entry. Survival time or censoring time, age, MMSE score, global cortical thickness and education are in mean ± standard deviation. The range of the number of visits is listed in square brackets.

	MCI subjects	Converters	Non-converters
Number of subjects	300	181	119
Number of imaging visits	4.52 [1–7]	4.66 [1–7]	4.32 [1–6]
Survival time or censoring time (years)	3.23 ± 2.51	2.51 ± 1.95	4.32 ± 2.86
Global cortical thickness (mm)	2.08 ± 0.15	2.05 ± 0.15	2.12 ± 0.13
Age (years)	74.70 ± 7.03	74.46 ± 6.96	75.06 ± 7.15
Gender (female/male)	113/187	66/115	47/72
MMSE score	27.04 ± 1.78	26.82 ± 1.75	27.37 ± 1.79
ApoE4 (carriers/non-carriers)	164/136	115/66	49/70
Education (years)	15.67 ± 3.04	15.80 ± 2.89	15.49 ± 3.25

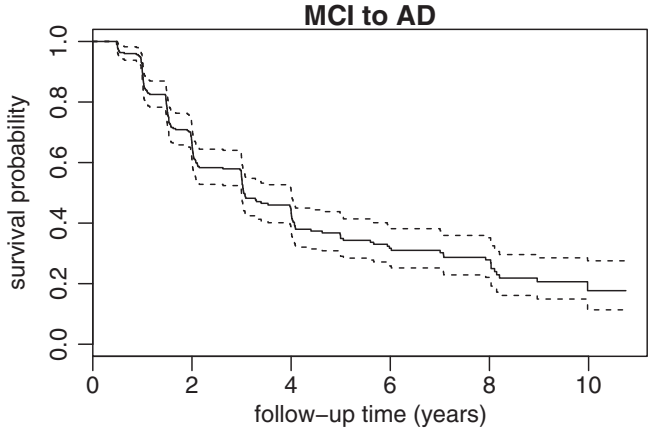

FIGURE 1: Kaplan–Meier estimated survival curve of the MCI-to-AD conversion. The dashed lines show pointwise 95% confidence intervals.

temporally smoothed the measurements using a linear mixed-effects (LME) model where cortical thickness was assumed to be a linear function of time with a random intercept and random slope for each subject. LME was used to find the trend over time due to the limited availability of repeated measures (no more than 3 for 21.67% of subjects) and the small value of cortical thickness (0.4–3.3 mm for 90% of subjects and vertices). Subject-specific random intercept and slope were introduced to account for the large variation between subjects as shown in Figure 2. The model was fitted in `R` using the `lme4` package (version 1.1-19, Bates et al., 2018). Due to the inter-subject variability of the timing of the serial imaging measurements and the lack of coincidence between clinical visits and image acquisitions, we estimated image data at all observed survival times via interpolation and extrapolation using the fitted LME model.

We used the Desikan–Killiany cortical atlas having 36 regions of interest (ROIs) in each hemisphere (Desikan et al., 2006). The "unknown" ROI corresponding to the medial wall and the corpus callosum are non-cortical and therefore were discarded from further analysis. Considering

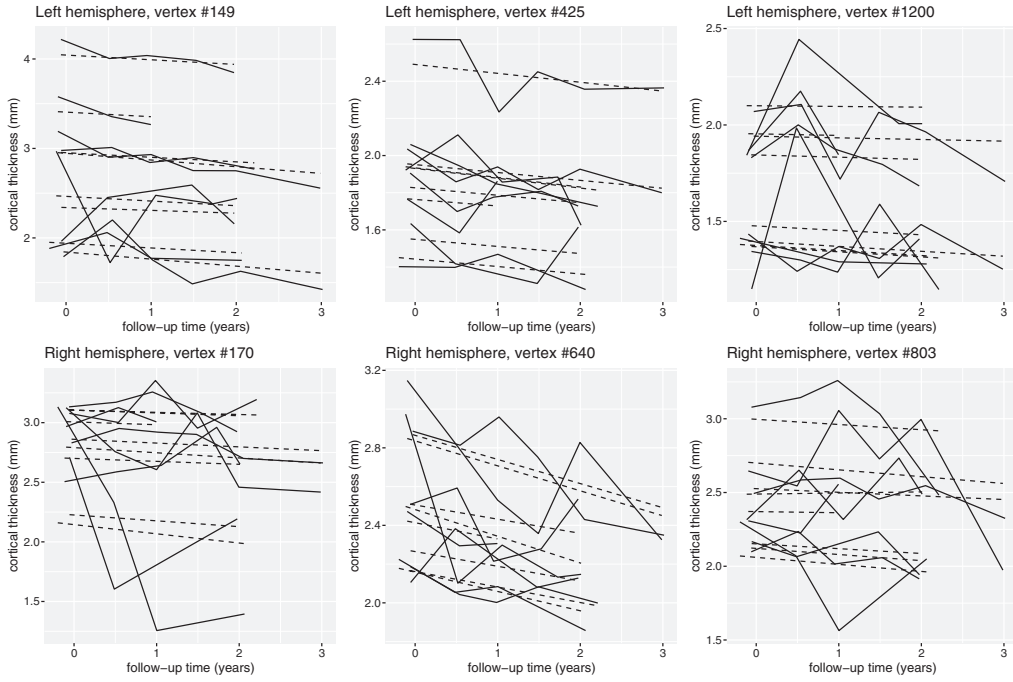

FIGURE 2: Spaghetti plots of the cortical thickness of 10 randomly selected subjects at 3 randomly selected vertices in each hemisphere. The raw data is shown as solid lines while the fitted result is shown as dashed lines.

the cortical folding along the cortical surface, we used the geodesic distance to compute the spatial correlation matrix in data analysis. Specifically, within each ROI, we used the C++ library `geodesic` (Kirsanov, 2008) to compute the exact geodesic distance on a surface mesh, which is determined by the number of edges in a shortest path connecting two vertices.

2.2. Vertex-specific Survival Models

Let Y_i denote the time-to-AD conversion for subject i ($i = 1, \ldots, n$), where we set the time origin at study entry. Define δ_i as an indicator that subject i converted to AD during the follow-up period. For non-converters, that is, $\delta_i = 0$, let y_i denote the time of the last visit, and thus $Y_i > y_i$. For converters, that is, $\delta_i = 1$, let y_{i0} denote the time of the last visit identified with MCI and let y_{i1} denote the time of the first visit identified with AD; thus, the survival time Y_i satisfies $y_{i0} < Y_i \leq y_{i1}$. Let $S_i(\cdot)$ denote the survival function for subject i. Assuming independent censoring, we write the likelihood as (Klein & Moeschberger, 2013)

$$\prod_{i=1}^{n} \{S_i(y_{i0}) - S_i(y_{i1})\}^{\delta_i} S_i(y_i)^{1-\delta_i}. \tag{1}$$

For subjects who converted to AD, that is, $\delta_i = 1$, we let $y_i := (y_{i0} + y_{i1})/2$ denote the midpoint of the interval and approximate the likelihood component $S_i(y_{i0}) - S_i(y_{i1})$ by $(y_{i1} - y_{i0})S_i(y_i)h_i(y_i)$ for ease of computation, where $h_i(\cdot)$ is the hazard function, defined as the instantaneous rate of occurrence of conversion to AD given that it has not previously occurred.

We let $x_{iv}(t)$ denote subject i's cortical thickness at the location of vertex v and time t. Additionally, we consider the following time- and space-invariant covariates: age at study entry, gender, and ApoE genotype status (1 if an e4 carrier and 0 otherwise), which has been found to

greatly increase the odds that an individual will develop Alzheimer's disease (Liu et al., 2013). Let Z denote the matrix of time- and space-invariant covariates, together with an additional column of ones and let Z_i be the ith row of Z, corresponding to subject i. Then the hazard function for subject i at vertex v is

$$h_{iv}(t) = h_{0v}(t) \exp\left(x_{iv}(t)\beta_v + Z_i^\top \boldsymbol{\gamma}_v\right), \tag{2}$$

where $h_{0v}(t)$ is the baseline hazard function, β_v is the coefficient for the cortical thickness effect at vertex v, and $\boldsymbol{\gamma}_v$ is the vector of coefficients for the time- and space-invariant covariates. Then the corresponding cumulative hazard is

$$H_{iv}(t) = \int_0^t h_{iv}(u)du, \tag{3}$$

and the associated survival function is $S_{iv}(t) = \exp\{-H_{iv}(t)\}$.

We assume a Weibull baseline hazard, $h_{0v}(t) = \alpha_v t^{\alpha_v - 1}$ with the shape parameter $\alpha_v > 0$. While other baseline hazard functions are possible, we used a Weibull baseline hazard because it simplifies the likelihood, leads to a tractable posterior distribution, and is computationally faster. Define $\boldsymbol{y} = (y_1, \ldots, y_n)^\top$ and $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_n)^\top$. Combining Equations (1)–(3), we rewrite the likelihood at vertex v up to a normalizing constant as

$$\begin{aligned} p(\boldsymbol{y}, \boldsymbol{\delta} | \beta_v, \boldsymbol{\gamma}_v, \alpha_v) \propto & \prod_{i=1}^n \left\{ \alpha_v y_i^{\alpha_v - 1} \exp\left(x_{iv}(y_i)\beta_v + Z_i^\top \boldsymbol{\gamma}_v\right) \right\}^{\delta_i} \\ & \times \exp\left\{ -\exp(Z_i^\top \boldsymbol{\gamma}_v) \int_0^{y_i} \alpha_v t^{\alpha_v - 1} \exp\left(x_{iv}(t)\beta_v\right) dt \right\}. \end{aligned} \tag{4}$$

We take advantage of the fitted LME model from the temporal smoothing in pre-processing and obtain subject i's cortical thickness at vertex v as a linear function of time t, $x_{iv}(t) = x_{iv}(0) + k_{iv}t$, where $x_{iv}(0)$ and k_{iv} are the estimated intercept and slope, respectively. We compute the integral $\int_0^{y_i} \alpha_v t^{\alpha_v - 1} \exp\left(x_{iv}(t)\beta_v\right) dt$ using the lower incomplete gamma function when $k_{iv}\beta_v < 0$, or using numerical integration when $k_{iv}\beta_v \geq 0$.

2.3. Spatial Model for the Cortical Thickness Effects

We conducted a preliminary ROI-level analysis to explore the spatially varying association between conversion to AD and time-varying cortical thickness averaged within each ROI. We analyzed each ROI separately using a Cox proportional hazard model with extension to incorporate time-dependent covariates. Table 2 shows the z-score from the preliminary analysis of each ROI-level cortical thickness effect reflecting the significance of the association. We observed appreciable association in most ROIs and large spatial variation in the strength of the association across ROIs. However, the results can be misleading as the imaging measurements are averaged within each ROI. The effect is most pronounced with large ROIs. An ROI-level analysis cannot address whether the strength of the association is evenly distributed within an ROI, or due to a small sub-region strongly associated with conversion to AD. To overcome this limitation, we need to investigate the association at each vertex.

We examined the spatially varying association between vertices via MUA, in which an extended Cox proportional hazard model was fitted at each vertex separately. Prior to MUA, the MRI data was spatially smoothed using a Gaussian kernel with a standard deviation of 3.4 mm, which corresponds to a FWHM of 8 mm. The amount of spatial smoothing was selected according to the recommendation by Bernal-Rusiel et al. (2013). In general, the optimal extent

TABLE 2: ROI-level ADNI MCI data analysis results: z-scores per hemisphere reflecting the significance of the association between time-to-conversion to AD and time-varying cortical thickness.

ROI	Left	Right	ROI	Left	Right
bankssts	−4.86	−2.60	parahippocampal	−3.09	−3.49
caudalanteriorcingulate	0.40	−0.81	parsopercularis	−4.17	−2.90
caudalmiddlefrontal	−2.51	−1.76	parsorbitalis	−3.64	−3.07
cuneus	−1.18	−0.76	parstriangularis	−1.61	−2.01
entorhinal	−5.61	−5.69	pericalcarine	−2.72	−3.25
frontalpole	−3.26	−2.60	postcentral	−1.83	−0.78
fusiform	−5.19	−4.63	posteriorcingulate	−3.29	−1.96
inferiorparietal	−4.66	−4.81	precentral	−1.81	−1.40
inferiortemporal	−5.26	−4.69	precuneus	−4.23	−4.31
insula	−3.67	−3.87	rostralanteriorcingulate	−2.50	−1.11
isthmuscingulate	−5.13	−4.16	rostralmiddlefrontal	−3.47	−2.20
lateraloccipital	−2.88	−2.79	superiorfrontal	−2.26	−1.87
lateralorbitofrontal	−3.02	−2.91	superiorparietal	−2.37	−2.60
lingual	−2.66	−2.81	superiortemporal	−3.70	−3.67
medialorbitofrontal	−2.35	−2.68	supramarginal	−4.20	−4.02
middletemporal	−5.76	−4.80	temporalpole	−5.83	−3.46
paracentral	−2.08	−2.18	transversetemporal	−2.46	−3.46

of smoothing depends on the sample size, the effect size, the spatial extent of the effect, and the type of multiple comparison correction. The estimated coefficients obtained using MUA at the vertices having uncorrected $P < 0.05$ are plotted in Figure 3. We observe large spatial variation in the thickness effects across the cortical surface. There seems to be a strong association between conversion to AD and time-varying cortical thickness in the temporal lobe and the parietal lobe, while many other areas show little effect. Furthermore, the empirical semivariogram of the estimated coefficients within each ROI had distinct shapes that varied across ROIs. Altogether, the MUA suggested an investigation of the spatial structure of the estimated coefficients within each ROI and a spatially varying association between time-to-conversion to AD and cortical thickness. The distinct shapes of the empirical semivariograms suggested a spatial structure in the estimated effects that is non-stationary across the cortical surface.

Hereafter, we present our method for one arbitrary ROI and assume that cortical thickness is observed at V vertices. We model the spatially varying coefficient for cortical thickness $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_V)^\top$ as a zero-mean Gaussian process (GP) capturing spatial association and measurement error or fine-scale variation. We assume the spatial correlation is isotropic to avoid the difficulty of accounting for directions on the topologically complex cortical surface. The covariance function of the GP that we use is given by $C_{\phi,\sigma^2}(d(v,v')) + \tau^2 I(d(v,v') = 0)$, where $C_{\phi,\sigma^2}(d(v,v')) = \sigma^2 \exp(-\phi d(v,v'))$, and $d(v,v')$ denotes the geodesic distance between vertex v and vertex v'. Other correlation structures are possible. The partial sill σ^2, nugget τ^2, and the decay or inverse range ϕ are unknown parameters that we need to estimate. For ease of notation, we define $C_{\phi,\sigma^2}(v,v') = C_{\phi,\sigma^2}(d(v,v'))$. We let $\mathcal{V} = (1, \ldots, V)^\top$ and write the spatial model as

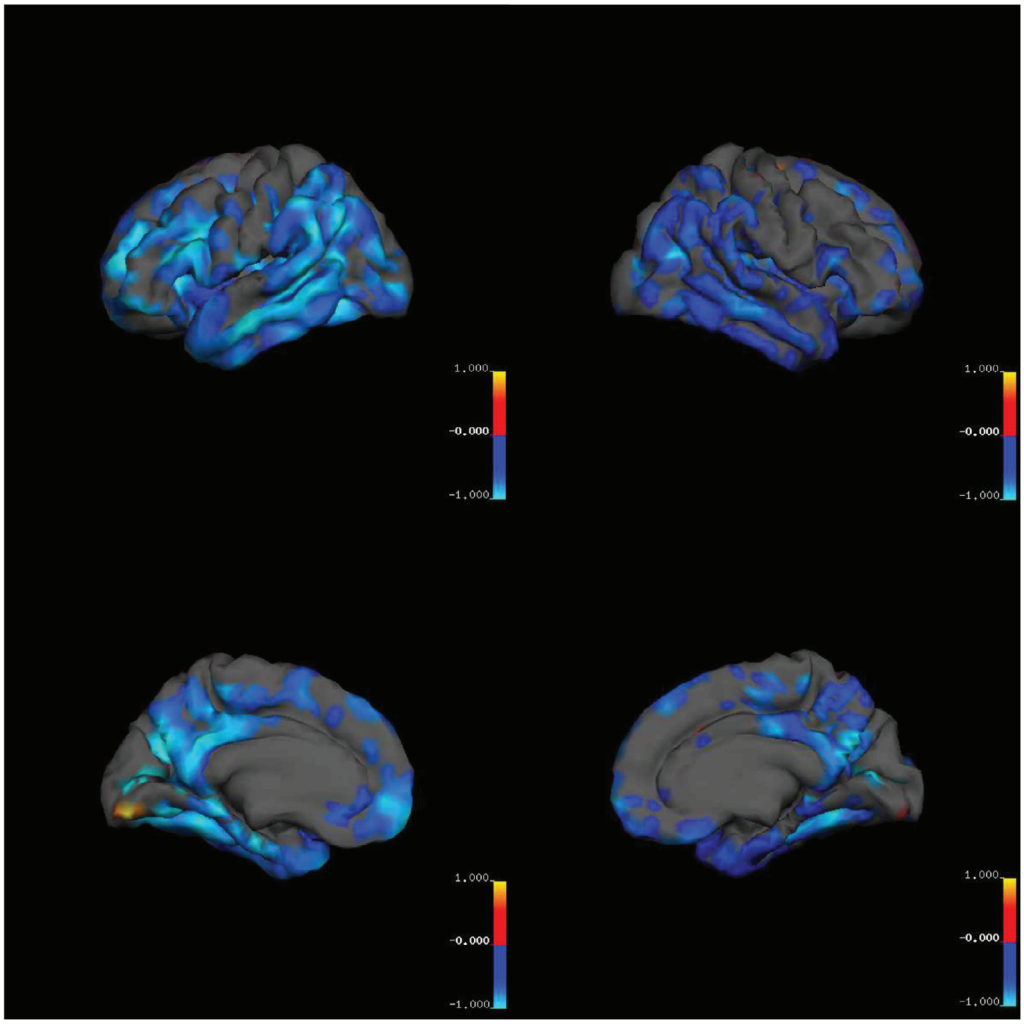

FIGURE 3: Results of the preliminary MUA of the ADNI data: Estimated coefficient for the cortical thickness effect at each vertex having uncorrected $P < 0.05$. The top and bottom rows show views of the lateral and medial surfaces, respectively. The left and right panels correspond to the left and right hemispheres, respectively.

$$\boldsymbol{\beta}|\phi, \sigma^2, \tau^2 \sim N(0, C_{\phi,\sigma^2}(\mathcal{V}, \mathcal{V}) + \tau^2 I_V)\,. \tag{5}$$

Traditional spatial-statistical techniques for likelihood-based inference involve the inverse and determinant of the dense $V \times V$ covariance matrices, where V, the number of cortical thickness measurements in an ROI, can be large. Incorporating time-dependent covariates in modeling survival data poses additional challenges to computing and storage. The GP is therefore computationally prohibitive and we need a highly scalable alternative to tackle the difficulties. Specifically, we use the nearest-neighbour Gaussian process (NNGP) proposed by Datta et al. (2016). In NNGP, the precision matrix of the GP $(C_{\phi,\sigma^2}(\mathcal{V}, \mathcal{V}) + \tau^2 I_V)^{-1}$ is approximated by $(I - A_{\phi,\sigma^2,\tau^2})^\top (D_{\phi,\sigma^2,\tau^2})^{-1}(I - A_{\phi,\sigma^2,\tau^2})$, so that the spatial model specified in Equation (5) becomes

$$\boldsymbol{\beta}|\phi,\sigma^2,\tau^2 \sim N(0,((I-A_{\phi,\sigma^2,\tau^2})^\top(D_{\phi,\sigma^2,\tau^2})^{-1}(I-A_{\phi,\sigma^2,\tau^2}))^{-1}), \tag{6}$$

where A_{ϕ,σ^2,τ^2} is a sparse and strictly lower triangular matrix with at most M non-zero entries in each row, where M is much less than V, and D_{ϕ,σ^2,τ^2} is a diagonal matrix (Finley et al., 2018). The likelihood that arises using the NNGP spatial model identified in (6) is a scalable alternative compared to the corresponding likelihood based on using the GP identified in Equation (5) using the GP, but the storage and computational burden of the former is linear in V (Datta et al., 2016). Datta et al. (2016) showed that the NNGP performs well and sometimes better than the GP when estimating the true underlying spatial process at a fraction of the computational cost. In fact, Datta et al. (2016) illustrated that usually a small value of M between 10 and 15 for the NNGP produces predictive performance on par with the GP when $V = 2,000$. In our data application, V is the number of cortical thickness measurements in an ROI which is less than 200 and therefore $M = 5$ suffices.

We now describe how to construct the matrices A and D, where we omit subscripts for ease of notation. Matrix A is strictly lower triangular with at most M non-zero entries in each row, where M is much less than V. Let $N(v)$ be the set of indices of at most M closest points to vertex v among the vertices indexed less than v. The vth row ($v > 1$) of A has non-zero entries only at positions indexed by $N(v)$, calculated via

$$A(v,N(v)) = C_{\phi,\sigma^2}(v,N(v))\{C_{\phi,\sigma^2}(N(v),N(v)) + \tau^2 I\}^{-1}. \tag{7}$$

Matrix D is diagonal with the vth ($1 \le v \le V$) diagonal entry given by

$$D(v,v) = \sigma^2 + \tau^2 - C_{\phi,\sigma^2}(v,N(v))\{C_{\phi,\sigma^2}(N(v),N(v)) + \tau^2 I\}^{-1}C_{\phi,\sigma^2}(N(v),v). \tag{8}$$

2.4. Bayesian Latent Spatial Modelling and Analysis

We now formally introduce our Bayesian latent spatial model, which combines the survival models that predict the timing of MCI-to-AD conversion at all vertices by putting a spatial prior on the coefficients for the time-varying cortical thickness. The spatial effects of the imaging measurements are not directly observed and hence are called latent. A graphical representation of the Bayesian latent spatial model is shown in Figure 4.

We adopt a Bayesian approach to parameter estimation. For parameters in the survival model specified in Equation (4), we follow Ibrahim, Chen & Sinha (2001) and assume the prior of the shape parameter of the Weibull distribution α_v's is independent and identically distributed (i.i.d.) Gamma(a,b) with $a = 1$ and $b = 0.02$. The prior of the coefficients for the time- and space-invariant covariates $\boldsymbol{\gamma}_v$'s is assumed to be i.i.d. $N(0, c_{\boldsymbol{\gamma}}^2 I)$ with $c_{\boldsymbol{\gamma}}^2 = 100$. For the spatial model specified in Equation (6), we follow Banerjee, Carlin & Gelfand (2014) and assume $\sigma^2 \sim$ Inv-Gamma$(\lambda_{\sigma^2},\kappa_{\sigma^2})$, $\tau^2 \sim$ Inv-Gamma$(\lambda_{\tau^2},\kappa_{\tau^2})$ with $\lambda_{\sigma^2} = \lambda_{\tau^2} = 2$ and $\kappa_{\sigma^2} = \kappa_{\tau^2} = 0.02$. Using the empirical semivariograms from the MUA, we saw that the range parameter varied between 0.1 and 1 across ROIs, and thus we choose $\phi \sim$ Uniform(θ_1,θ_2) with $\theta_1 = 1$ and $\theta_2 = 10$.

Define $\boldsymbol{\gamma} = (\boldsymbol{\gamma}_1^\top,\ldots,\boldsymbol{\gamma}_V^\top)^\top$ and $\boldsymbol{\alpha} = (\alpha_1,\ldots,\alpha_V)^\top$. Combining the survival models specified in Equation (4), the spatial model identified in Equation (6), and the prior specifications, the Bayesian latent spatial model is as follows:

$$p(\mathbf{y},\boldsymbol{\delta}|\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\alpha}) \propto \prod_{v=1}^{V}\prod_{i=1}^{n}\left\{\alpha_v y_i^{\alpha_v-1}\exp\left(x_{iv}(y_i)\beta_v + Z_i^\top\boldsymbol{\gamma}_v\right)\right\}^{\delta_i}$$
$$\times\exp\left\{-\exp(Z_i^\top\boldsymbol{\gamma}_v)\int_0^{y_i}\alpha_v t^{\alpha_v-1}\exp\left(x_{iv}(t)\beta_v\right)dt\right\},$$

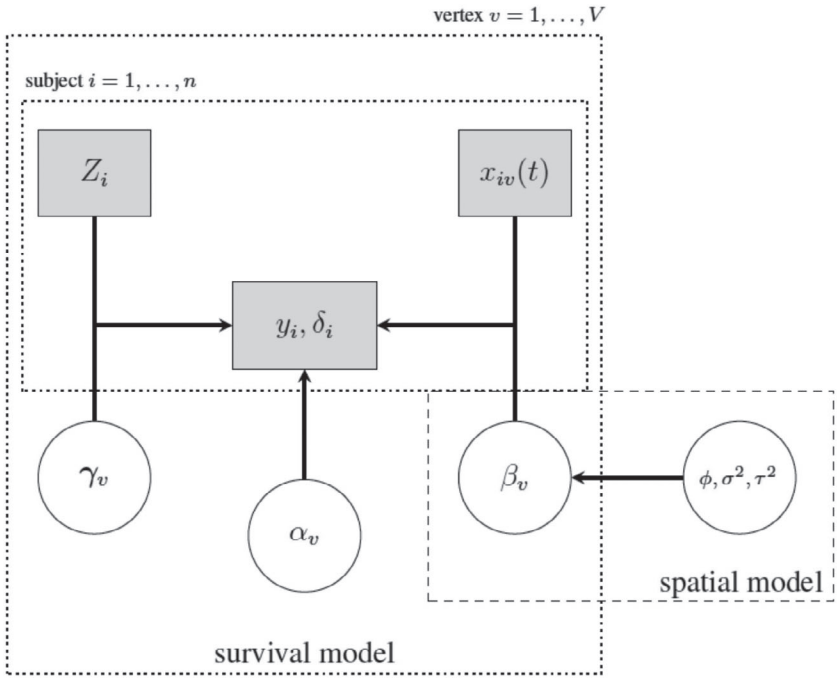

FIGURE 4: Graphical representation of the Bayesian latent spatial model. Shaded squares represent observed quantities and circles represent unknowns.

$$
\begin{aligned}
\boldsymbol{\beta}|\phi, \sigma^2, \tau^2 &\sim N(0, ((I - A_{\phi,\sigma^2,\tau^2})^\top D^{-1}_{\phi,\sigma^2,\tau^2}(I - A_{\phi,\sigma^2,\tau^2}))^{-1}),\\
\boldsymbol{\gamma}_v &\overset{i.i.d.}{\sim} N(0, c^2_{\boldsymbol{\gamma}} I),\\
\alpha_v &\overset{i.i.d.}{\sim} \text{Gamma}(a, b),\\
\phi &\sim \text{Uniform}(\theta_1, \theta_2),\\
\sigma^2 &\sim \text{Inv-Gamma}(\lambda_{\sigma^2}, \kappa_{\sigma^2}),\\
\tau^2 &\sim \text{Inv-Gamma}(\lambda_{\tau^2}, \kappa_{\tau^2}).
\end{aligned}
$$

Bayesian inference can proceed by sampling from the posterior $p(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\alpha}, \phi, \sigma^2, \tau^2|\boldsymbol{y}, \boldsymbol{\delta})$ using Markov chain Monte Carlo (MCMC) methods. We estimated the coefficients of the time-varying cortical thickness at all vertices using the posterior means and the highest posterior density (HPD) intervals were calculated using the MCMC simulations. Association at a vertex was declared to be significant if the 95% HPD interval excluded zero.

2.5. Computational Details

We used the Hamiltonian Monte Carlo (HMC) algorithm (Neal, 2011) with the No-U-Turn sampler (NUTS) with dual averaging (Hoffman & Gelman, 2014), which is an extension to HMC that automatically adapts the tuning parameters. Compared to the random walk Metropolis-Hastings, the HMC substantially reduces the computational cost because HMC iterations are less correlated, and thus it takes fewer iterations to achieve the same level of precision when estimating the posterior mean (Neal, 2011). HMC and NUTS are implemented

in Stan (Stan Development Team, 2018b) and its `R` interface as the `rstan` package (version 2.17.3 Stan Development Team, 2018a) which we used for model fitting and data analysis.

We assessed performance of the MCMC algorithm by examination of trace plots, density plots, and autocorrelation plots, all of which indicated sufficiently fast mixing of the simulation. For convergence diagnostics, we followed Geyer (2011) and Vats, Flegal & Jones (2019) by assessing the Monte Carlo error using the multivariate initial sequence method (Dai & Jones, 2017). To select the length of the warmup phase used to adapt any tunable MCMC parameters and determine when to terminate the simulation, we used the relative standard deviation fixed-volume stopping rule and selected the simulation effort so that the multivariate effective sample size (ESS) exceeded 3,000 (Dai & Jones, 2017; Vats, Flegal & Jones, 2018, 2019). Analyses of MCMC output were performed in `R` using the `mcmcse` package (version 1.3-2, Flegal et al., 2017). The results varied across ROIs depending mainly on the number of vertices in each ROI. For a medium-sized ROI, an ESS of 3,000 required 10,000 iterations with a warmup length of 500, and 96 h of computing time.

3. SIMULATION STUDY

In a two-dimensional Euclidean space, we considered a domain of a grid of size 32×32, where each location was indexed by (v_x, v_y) with $v_x, v_y = 1, 2, \ldots, 32$. To generate the coefficient β_v, $v = 1, 2, \ldots, V$ at all $V = 1,024$ locations, we used a zero-mean Gaussian random field, where the covariance function was given by an exponential semivariogram model with a sill of 0.2, a range of 8 and a nugget of 0. The simulated values having a magnitude less than 0.1 were reset to 0. Figure 5 shows the resulting coefficients.

We generated the covariates at all locations as follows. For each subject independently, we generated i.i.d. normal errors with mean 0 and standard deviation 0.5 at all locations, followed by Gaussian kernel smoothing with a standard deviation of 1. Then we generated a random effect for each subject independently from a normal distribution with mean 3 and standard deviation 1 and added it to the simulated values at all locations.

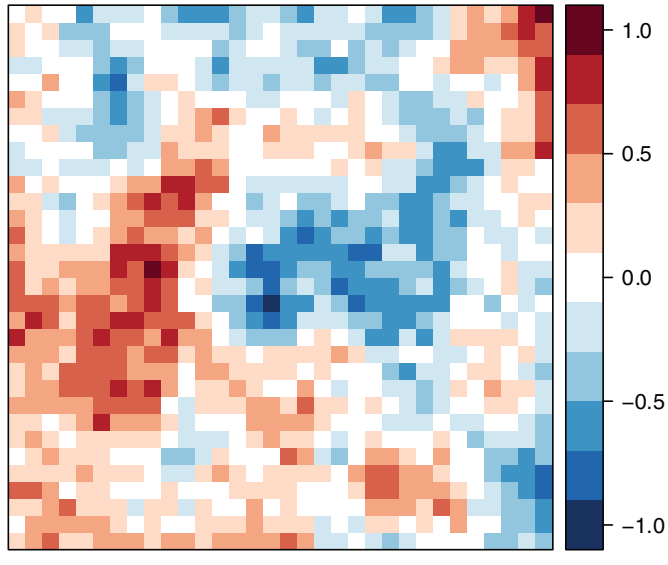

FIGURE 5: Simulated coefficients. The simulated coefficients are non-zero at 794 locations shown in colour, and zero at 230 locations shown in white.

TABLE 3: Simulation results obtained using the Bayesian latent spatial model and MUA. The FNR and FPR within the spatial domain were estimated using the 100 simulated data sets with standard errors given in parentheses. The MSE of the estimated coefficient at each location was calculated using the 100 simulated data sets and is presented as a median with 25 and 75% quantiles in square brackets across the non-null locations and the null locations, respectively.

Performance measure	Bayesian latent spatial model	MUA
FNR	0.020 (0.000)	0.415 (0.001)
FPR	0.038 (0.001)	0.052 (0.001)
MSE at non-null locations	0.002 [0.002, 0.003]	0.018 [0.014, 0.032]
MSE at null locations	0.002 [0.001, 0.002]	0.017 [0.015, 0.018]

We simulated each subject's survival time from a Weibull model with a shape of 1 and a scale equal to the exponential of the negative of the simulated covariate times the simulated coefficient at each location. Censoring times were independently simulated from a uniform distribution from 0 to c, where c was chosen to achieve a desired censoring rate of approximately 40%, which was similar to what we observed in the ADNI data.

We chose a sample size of $n = 100$ and simulated 100 data sets of survival time and censoring time using the same set of simulated covariates and coefficients. The Bayesian latent spatial model was implemented to estimate the coefficients. The computation time for one repetition was approximately 3 h to achieve an ESS of 3,000. The effect at a location was declared significant if the 95% HPD interval excluded zero. We also conducted a MUA for benchmarking. To examine the performance of each method, we calculated the false negative rate (FNR) and the false positive rate (FPR) within the spatial domain, and the mean squared error (MSE) at each location. We first obtained an FNR and an FPR within the spatial domain for each repetition, and then calculated the averages and standard errors from the 100 simulated data sets. The MSE at each location was calculated using the 100 simulated data sets. We reported the median, 25 and 75% quantiles of the MSE among the non-null locations where the coefficients were non-zero, and among the null locations, respectively.

Table 3 presents the results where the performance of the Bayesian latent spatial model is compared to MUA. The Bayesian latent spatial model shows substantial improvement in the FNR, which corresponds to greater statistical power. We also observed a small but consistent decrease in the FPR, which indicates that the Bayesian latent spatial model is sharp in detecting true association but not oversensitive. The MSE obtained from the Bayesian latent spatial model appears small at each location. In addition to reduced estimation error, the Bayesian latent spatial model also shows smaller variance, which led to more stable inference. Altogether, this simulation study shows the substantial improvements in the error rates and in the estimation of the coefficients of the Bayesian latent spatial model compared to the MUA, which resulted from the latent spatial model's use of information from spatial neighbours, which the MUA ignored.

Additionally, we conducted a simulation study that uses the ADNI study data from cortical thickness images as covariates and the Gaussian random field to generate the true outcome on the cortical surface, and a simulation study where the true baseline hazard is not Weibull as assumed in the proposed method. The results in terms of the FNR, FPR and MSE obtained from these new scenarios led to similar conclusions concerning improved performance from using the proposed model. A detailed description and the results of the additional simulations are provided in the Supplementary Material.

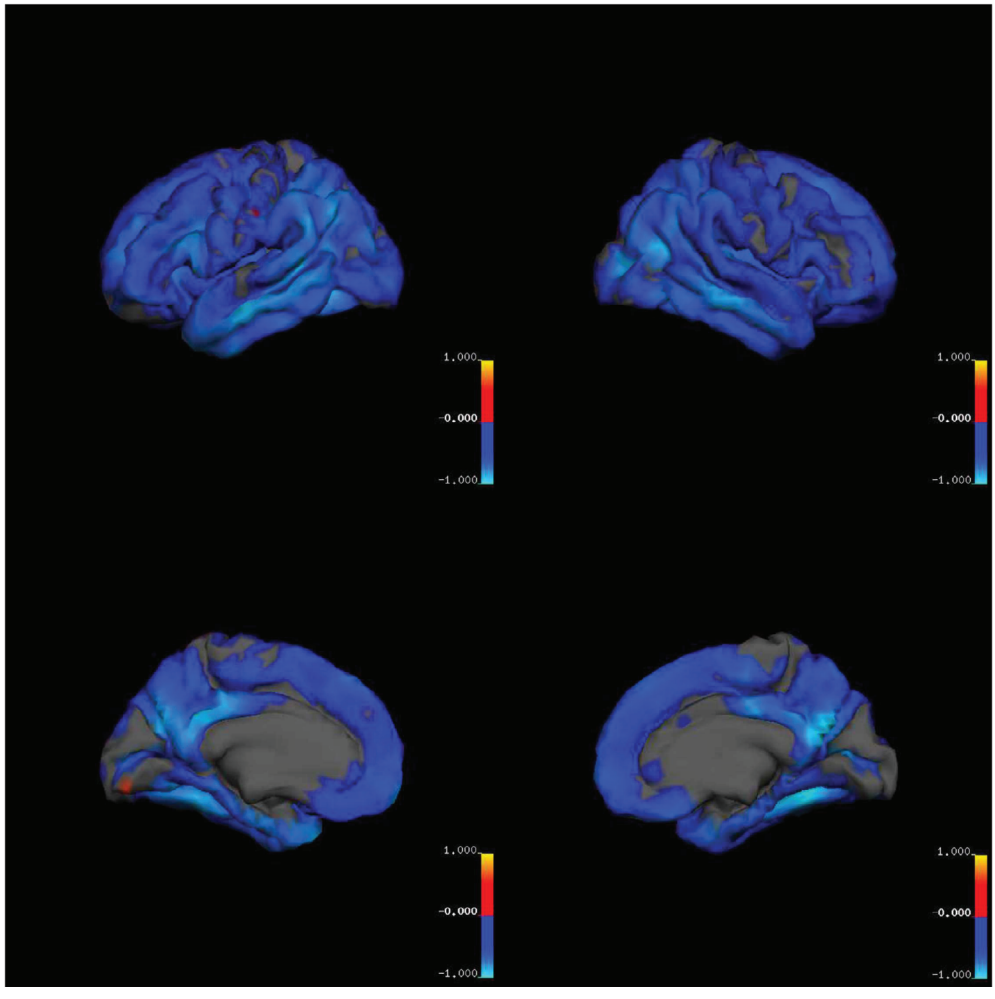

FIGURE 6: ADNI data analysis results obtained using our Bayesian latent spatial model: Estimated coefficient for the cortical thickness effect at each vertex where significant association is detected at the 95% credible level. The top and bottom rows show views of the lateral and medial surfaces, respectively. The left and right panels correspond to the left and right hemispheres, respectively.

4. DATA ANALYSIS

We analyzed each ROI separately using the Bayesian latent spatial model that accounts for spatial dependence between vertices. To visualise the result, we plotted the estimated coefficients at the vertices showing significant association on a whole-brain map; see Figure 6. At the 95% credible level, significant association was detected at 75.5 and 74.6% of all the vertices in the left hemisphere and the right hemisphere, respectively. Almost all significant effects were negative (99.8% in the left hemisphere and 99.9% in the right hemisphere), which indicates that a thinner cortex is associated with a higher risk of developing AD. In the left hemisphere, the estimated coefficient for the cortical thickness effect ranged from -0.94 to -0.11 with a median of -0.46, a 25% quantile of -0.57 and a 75% quantile of -0.37; in the right hemisphere, the estimated coefficient for the cortical thickness effect ranged from -0.97 to -0.15 with a median of -0.44,

a 25% quantile of −0.54 and a 75% quantile of −0.35. An estimated coefficient of −0.5 means that the hazard of developing AD at any time instant will increase by a factor of 5.13% if the cortical thickness at that time instant decreases by 0.1 mm.

Compared to the ROI-level analyses, our method can quantify the cortical thickness effect at each precise spatial location and detect association that is sparsely distributed in an ROI. The 1% largest effect sizes in the left hemisphere, namely, estimated coefficient smaller than −0.815, appeared at 11 vertices in the temporal lobe (bankssts, fusiform, middletemporal, inferiortemporal, and temporalpole), 13 vertices in the parietal lobe (inferiorparietal, precuneus, and isthmuscingulate), and 2 vertices in the frontal lobe (parsopercularis and frontalpole). The 1% largest effect sizes in the right hemisphere, namely, estimated coefficient smaller than −0.798, appeared at 8 vertices in the temporal lobe (fusiform), 12 vertices in the parietal lobe (inferiorparietal and precuneus), 4 vertices in the occipital lobe (lateraloccipital and pericalcarine), and 2 vertices in the frontal lobe (precentral). The results obtained using our Bayesian latent spatial model generally agree with our preliminary MUA and prior studies (Buckner et al., 2005; Dickerson et al., 2009). However, importantly, we were able to detect a larger spatial extent of association between conversion to AD and cortical thickness using the Bayesian latent spatial model. The significant vertices detected by MUA as shown in Figure 3 generally lay within the area detected by the Bayesian latent spatial model as shown in Figure 6. Specifically, the Bayesian latent spatial model detected 1.62 times more significant vertices than were detected by MUA in the left hemisphere, and 2.54 times more in the right hemisphere. These results are consistent with the conclusions resulting from our simulation study, namely, that our model substantially enhances statistical power since it accounts for the spatial correlation between vertices.

5. DISCUSSION

We have developed a novel Bayesian latent spatial model for analyzing the spatially varying association between MCI-to-AD conversion and the longitudinal measurements of cortical thickness across the surface of the brain. The Bayesian latent spatial model was employed in an analysis of the longitudinal MRI data from ADNI to investigate the topographic patterns of cortical thinning associated with an increased risk of developing AD. Though our results yielded significant effects in areas previously reported in the literature, we were able to obtain larger effect sizes and spatial extent of the associations than those previously reported in the literature by incorporating the spatial structure of the cortical thickness effects.

The Bayesian latent spatial model offers excellent statistical power and stability for detecting associations between clinical events and geographically referenced longitudinal imaging measurements. The data analysis results show a substantial boost in statistical power with respect to MUA, which ignores the spatial dependence in the cortical thickness effects. Using a simulation study, we validated that the Bayesian latent spatial model is powerful in detecting true association without being oversensitive. Additionally, compared to previous work that analyzed imaging data obtained from a single visit, our approach opens up the possibility of examining the relationship between the temporal dynamics of imaging data and clinical events by utilizing serial measurements that can vary substantially over time.

While the proposed model assumes a Weibull baseline hazard for ease of computation, other baseline hazards are possible and the Weibull assumption is subject to change based on the dictates of the data. We observed a higher false positive rate when the Weibull baseline assumption was violated in a simulation study. Therefore, we suggest selecting the appropriate baseline hazard on a case-by-case basis when using our proposed method in applications. Another limitation of our analysis is that we did not account for missing data. Missing data is not uncommon in MRI studies due to scanner artifact, missed appointments or attrition. We used approximations to handle issues due to the lack of coincidence between clinical visits and image

acquisitions and a great amount and unknown pattern of variation and missing information when imaging data was acquired repeatedly for each subject. In turn, this lack of information makes it difficult to evaluate the loss attributable to the use of these various approximations. Further investigations are needed to address these problems.

The code for implementing the Bayesian latent spatial model and data analysis is available at https://github.com/mfiecas/ADcorticalthickness.

ACKNOWLEDGEMENTS

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI; Principal Investigator Michael W. Weiner, MD) and Department of Defense ADNI. 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 James S. Hodges for helpful conversations.

BIBLIOGRAPHY

Banerjee, S. (2016). Spatial data analysis. *Annual Review of Public Health*, 37, 47–60.

Banerjee, S., Carlin, B. P., & Gelfand, A. E. (2014). *Hierarchical Modeling and Analysis for Spatial Data*, 2nd ed., Chapman & Hall/CRC, Boca Raton.

Bates, D., Mächler, M., Bolker, B., Walker, S., Orcid, R. H. B. C., Singmann, H., Dai, B., Scheipl, F., Grothendieck, G., Green, P., & Fox, J. (2018). lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1-19, https://cran.r-project.org/web/packages/lme4/index.html.

Bernal-Rusiel, J. L., Reuter, M., Greve, D. N., Fischl, B., & Sabuncu, M. R. (2013). Spatiotemporal linear mixed effects modeling for the mass-univariate analysis of longitudinal neuroimage data. *NeuroImage*, 81, 358–370.

Braak, H. & Braak, E. (1991). Neuropathological stageing of Alzheimer-related changes. *Acta Neuropathologica*, 82, 239–259.

Brun, A. & Gustafson, L. L. (1976). Distribution of cerebral degeneration in Alzheimer's disease. *Archiv für Psychiatrie und Nervenkrankheiten*, 223, 15–33.

Buckner, R. L., Snyder, A. Z., Shannon, B. J., LaRossa, G., Sachs, R., Fotenos, A. F., Sheline, Y. I., Klunk, W. E., Mathis, C. A., Morris, J. C., & Mintun, M. A. (2005). Molecular, structural, and functional characterization of Alzheimer's disease: Evidence for a relationship between default activity, amyloid, and memory. *Journal of Neuroscience*, 25, 7709–7717.

Carroll, K. J. (2003). On the use and utility of the Weibull model in the analysis of survival data. *Controlled Clinical Trials*, 24, 682–701.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society Series B*, 34, 187–220.

Crous-Bou, M., Minguillón, C., Gramunt, N., & Molinuevo, J. L. (2017). Alzheimer's disease prevention: From risk factors to early intervention. *Alzheimer's Research & Therapy*, 9, 71.

Dai, N. & Jones, G. L. (2017). Multivariate initial sequence estimators in Markov chain Monte Carlo. *Journal of Multivariate Analysis*, 159, 184–199.

Datta, A., Banerjee, S., Finley, A. O., & Gelfand, A. E. (2016). Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111, 800–812.

Desikan, R. S., Cabral, H. J., Settecase, F., Hess, C. P., Dillon, W. P., Glastonbury, C. M., Weiner, M. W., Schmansky, N. J., Salat, D. H., & Fischl, B. (2010). Automated MRI measures predict progression to Alzheimer's disease. *Neurobiology of Aging*, 31, 1364–1374.

Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., & Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31, 968–980.

Devanand, D. P., Pradhaban, G., Liu, X., Khandji, A., De Santi, S., Segal, S., Rusinek, H., Pelton, G. H., Honig, L. S., Mayeux, R., Stern, Y., Tabert, M. H., & de Leon, M. J. (2007). Hippocampal and entorhinal atrophy in mild cognitive impairment. *Neurology*, 68, 828–836.

Dickerson, B. C., Bakkour, A., Salat, D. H., Feczko, E., Pacheco, J., Greve, D. N., Grodstein, F., Wright, C. I., Blacker, D., Rosas, H. D., Sperling, R. A., Atri, A., Growdon, J. H., Hyman, B. T., Morris, J. C., Fischl, B., & Buckner, R. L. (2009). The cortical signature of Alzheimer's disease: Regionally specific cortical thinning relates to symptom severity in very mild to mild AD dementia and is detectable in asymptomatic amyloid-positive individuals. *Cerebral Cortex*, 19, 497–510.

Dickerson, B. C., Stoub, T. R., Shah, R. C., Sperling, R. A., Killiany, R. J., Albert, M. S., Hyman, B. T., Blacker, D., & deToledo Morrell, L. (2011). Alzheimer-signature MRI biomarker predicts AD dementia in cognitively normal adults. *Neurology*, 76, 1395–1402.

Finley, A. O., Datta, A., Cook, B. C., Morton, D. C., Andersen, H. E., & Banerjee, S. (2018). Efficient algorithms for Bayesian nearest-neighbor Gaussian processes. *Journal of Computational and Graphical Statistics*, 1–37.

Fischl, B. (2012). Freesurfer. *NeuroImage*, 62, 774–781.

Fisher, L. D. & Lin, D. Y. (1999). Time-dependent covariates in the Cox proportional-hazards regression model. *Annual Review of Public Health*, 20, 145–157.

Flegal, J. M., Hughes, J., Vats, D., & Dai, N. (2017). mcmcse: Monte Carlo standard errors for MCMC. R package version 1.3-2, https://cran.r-project.org/web/packages/mcmcse/index.html.

Galvin, J. E. (2017). Prevention of Alzheimer's disease: Lessons learned and applied. *Journal of the American Geriatrics Society*, 65, 2128–2133.

Geyer, C. J. (2011). Introduction to Markov chain Monte Carlo. In S. Brooks, A. Gelman, G. L. Jones, & X.-L. Meng (Eds.) *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC, Boca Raton, 3–48.

Gomez-Isla, T., Price, J. L., & Hyman, B. T. (1996). Profound loss of layer II entorhinal cortex neurons occurs in very mild Alzheimer's disease. *Journal of Neuroscience*, 16, 4491–4500.

Hoffman, M. D. & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623.

Ibrahim, J. G., Chen, M. H., & Sinha, D. (2001). *Bayesian Survival Analysis*. Springer, New York.

Jack, C. R., Vemuri, P., Wiste, H. J., Weigand, S. D., Lesnick, T. G., Lowe, V., Kantarci, K., Bernstein, M. A., Senjem, M. L., Gunter, J. L., Boeve, B. F., Trojanowski, J. Q., Shaw, L., Aisen, P., Weiner, M., Petersen, R., & Knopman, D. (2012). Shapes of the trajectories of 5 major biomarkers of Alzheimer disease. *Archives of Neurology*, 69, 856–867.

Kirsanov, D. (2008). geodesic: Multiple source/target exact geodesic (shortest path) algorithm for triangular mesh (triangulated 2D surface in 3D), https://code.google.com/archive/p/geodesic/.

Klein, J. P. & Moeschberger, M. L. (2013). *Survival Analysis: Techniques for Censored and Truncated Data*. Springer, New York.

Klein, J. P., van Houwelingen, H. C., Ibrahim, J. G., & Scheike, T. H. (2016). *Handbook of Survival Analysis*. Chapman & Hall/CRC, Boca Raton.

Lee, E., Zhu, H., Kong, D., Wang, Y., Sullivan Giovanello, K., Ibrahim, J. G., & for the Alzheimer's Disease Neuroimaging Initiative. (2015). BFLCRM: A Bayesian functional linear Cox regression model for predicting time to conversion to Alzheimer's disease. *The Annals of Applied Statistics*, 9, 2153–2178.

Lerch, J. P., Pruessner, J. C., Zijdenbos, A., Hampel, H., Teipel, S. J., & Evans, A. C. (2005). Focal decline of cortical thickness in Alzheimer's disease identified by computational neuroanatomy. *Cerebral Cortex*, 15, 995–1001.

Liu, C. -C., Kanekiyo, T., Xu, H., & Bu, G. (2013). Apolipoprotein E and Alzheimer disease: Risk, mechanisms and therapy. *Nature Reviews Neurology*, 9, 106–118.

Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, & X.-L. Meng (Eds.) *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC, Boca Raton, 113–162.

Petersen, R., Smith, G., Waring, S., Ivnik, R., Tangalos, E., & Kokmen, E. (1999). Mild cognitive impairment: Clinical characterization and outcome. *Archives of Neurology*, 56, 303–308.

Petersen, R. C., Doody, R., Kurz, A., Mohs, R. C., Morris, J. C., Rabins, P. V., Ritchie, K., Rossor, M., Thal, L., & Winblad, B. (2001). Current concepts in mild cognitive impairment. *Archives of Neurology*, 58, 1985–1992.

Reuter, M., Schmansky, N. J., Rosas, H. D., & Fischl, B. (2012). Within-subject template estimation for unbiased longitudinal image analysis. *NeuroImage*, 61, 1402–1418.

Rushton, G. (2003). Public health, GIS, and spatial analytic tools. *Annual Review of Public Health*, 24, 43–56.

Sabuncu, M. R., Bernal-Rusiel, J. L., Reuter, M., Greve, D. N., & Fischl, B. (2014). Event time analysis of longitudinal neuroimage data. *NeuroImage*, 97, 9–18.

Sabuncu, M. R., Buckner, R. L., Smoller, J. W., Lee, P. H., Fischl, B., Sperling, R. A., & for the Alzheimer's Disease Neuroimaging Initiative. (2012). The association between a polygenic Alzheimer score and cortical thickness in clinically normal subjects. *Cerebral Cortex*, 22, 2653–2661.

Stan Development Team. (2018a). RStan: The R interface to Stan. R package version 2.17.3, http://mc-stan.org/.

Stan Development Team. (2018b). Stan modeling language: User's guide and reference manual. Stan Version 2.18.0. http://mc-stan.org.

Stoub, T. R., Bulgakova, M., Leurgans, S., Bennett, D. A., Fleischman, D., Turner, D. A., & de Toledo-Morrell, L. (2005). MRI predictors of risk of incident Alzheimer disease. *Neurology*, 64, 1520–1524.

Vats, D., Flegal, J. M., & Jones, G. L. (2018). Strong consistency of multivariate spectral variance estimators in Markov chain Monte Carlo. *Bernoulli*, 24, 1860–1909.

Vats, D., Flegal, J. M., & Jones, G. L. (2019). Multivariate output analysis for Markov chain Monte Carlo. *Biometrika*, 106, 321–337.

Vemuri, P., Weigand, S. D., Knopman, D. S., Kantarci, K., Boeve, B. F., Petersen, R. C., & Jack, C. R. (2011). Time-to-event voxel-based techniques to assess regional atrophy associated with MCI risk of progression to AD. *NeuroImage*, 54, 985–991.

Zhou, H. & Hanson, T. (2018). A unified framework for fitting Bayesian semiparametric models to arbitrarily censored survival data, including spatially referenced data. *Journal of the American Statistical Association*, 113, 571–581.

Received 1 January 2020
Accepted 29 June 2020