



Modeling and Prediction of Multiple Correlated Functional Outcomes

Jiguo CAO , Kunlaya SOIAPORN, Raymond J. CARROLL, and David RUPPERT

We propose a copula-based approach for analyzing functional data with correlated multiple functional outcomes exhibiting heterogeneous shape characteristics. To accommodate the possibly large number of parameters due to having several functional outcomes, parameter estimation is performed in two steps: first, the parameters for the marginal distributions are estimated using the skew t family, and then the dependence structure both within and across outcomes is estimated using a Gaussian copula. We develop an estimation algorithm for the dependence parameters based on the Karhunen–Loève expansion and an EM algorithm that significantly reduces the dimension of the problem and is computationally efficient. We also demonstrate prediction of an unknown outcome when the other outcomes are known. We apply our methodology to diffusion tensor imaging data for multiple sclerosis (MS) patients with three outcomes and identify differences in both the marginal distributions and the dependence structure between the MS and control groups. Our proposed methodology is quite general and can be applied to other functional data with multiple outcomes in biology and other fields.

Supplementary materials accompanying this paper appear online.

Key Words: Diffusion tensor imaging; Gaussian copulas; Multiple sclerosis; Skewed functional data; Tractography data.

1. INTRODUCTION

Functional data analysis (FDA) is a powerful tool for modeling data observed at various time points or locations to uncover the underlying features that are assumed to vary over a continuum of time points or locations. Ramsay and Silverman (2005) and Ferraty and Romain (2010) provide a comprehensive review of functional data analysis. See Tsing and

Jiguo Cao and Kunlaya Soiaporn have contributed equally to this article.

Jiguo Cao (✉), Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby, BC V5A1S6, Canada (E-mail: jiguo_cao@sfu.ca). Kunlaya Soiaporn, Capital One, Vienna, VA 22180, USA (E-mail: jeab.soiaporn@gmail.com). Raymond J. Carroll, Department of Statistics, Texas A&M University, College Station, TX 77843, USA and School of Mathematical and Physical Sciences, University of Technology Sydney, Broadway, NSW 2007, Australia (E-mail: carroll@stat.tamu.edu). David Ruppert, Department of Statistical Science and School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14850, USA (E-mail: dr24@cornell.edu).

© 2018 International Biometric Society

Journal of Agricultural, Biological, and Environmental Statistics, Volume 24, Number 1, Pages 112–129
<https://doi.org/10.1007/s13253-018-00344-0>

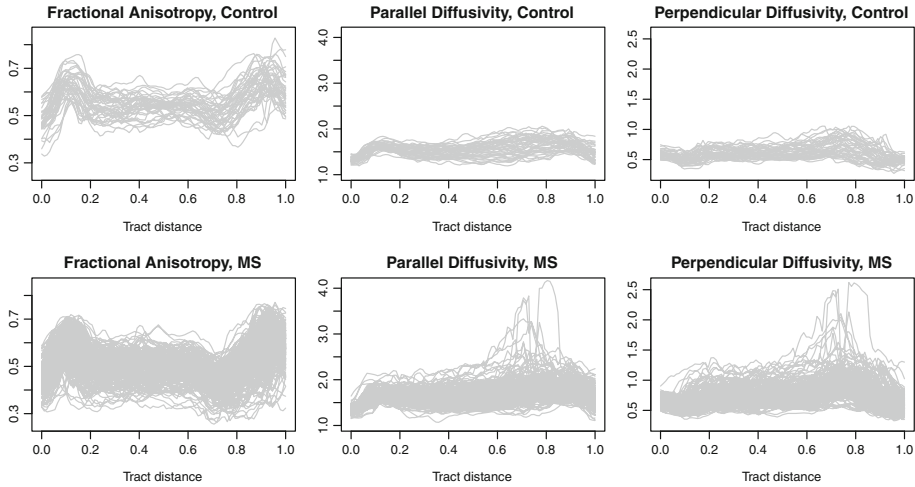


Figure 1. Plot of DTI data. The top row shows the three outcomes from 42 healthy controls. The bottom row shows the same three outcomes from 162 MS patients. The tract distance is normalized to $[0,1]$. See Sect. 5 for a full discussion of these data.

Eubank (2015) for the theoretical development of FDA. Often we obtain multiple functional observations and we wish to study them simultaneously. In this case, it is important to have tools to study the dependence structure between different outcomes. The methodology should be efficient in reducing the dimension of the data to accommodate the possible large size of functional data, especially in the case of multiple outcomes.

This study proposes a method for studying functional data with multiple outcomes that exhibits some shape characteristics that vary with spatial or temporal locations. Figure 1 illustrates an example of data that motivated this study. The data were obtained from a diffusion tensor imaging (DTI) study of neuronal tracts in 162 multiple sclerosis (MS) patients and 42 healthy controls. DTI is a magnetic resonance imaging technique for measuring the diffusion of water that can be used to detect abnormalities in brain tissues (see Basser et al. 1994, 2000). DTI tractography can be summarized by 3 functional outcomes, fractional anisotropy, parallel diffusivity and perpendicular diffusivity, each along the normalized tract distance; see Sect. 5 for definitions of these outcomes. As seen in Fig. 1, for each outcome, the pointwise mean and variance of both groups vary along the tract location and are different between the two groups. Another interesting characteristic of this dataset is that the amount of skewness seems to vary along the tract location and are greater among patients than among controls. An appropriate methodology to study the data should be able to capture this variation in the skewness, especially because skewness is related to disease status and could provide insights into the disease process.

Much of the FDA literature assumes Gaussian marginal distributions, but non-Gaussian functional data are not uncommon. Fortunately, models for non-Gaussian data are now being developed. Staicu et al. (2012) developed a copula-based approach for analyzing functional data with one outcome and applied their study to the parallel diffusivity measurement from this DTI data. The data were assumed to have pointwise marginal distributions in a parametric family with shape parameters such as a skew normal or skew t family (Azzalini 1985; Azzalini and Capitanio 2003). The mean, standard deviation and shape parameter functions were

modeled nonparametrically as functions of location. The dependence structure was estimated assuming an underlying Gaussian or t copula using methods based on principal component analysis (PCA) or the sample Kendall's tau matrix. The results in Staicu et al. (2012) showed that the marginal mean, variance, and skewness functions were different between the healthy and MS groups, while the correlation within the outcome is remarkably similar between the two groups. In contrast, we found the interesting result that the correlation between two different outcomes can differ between the MS patients and controls; see Sect. 5.

There are relatively few previous studies that analyze the relationship between multiple functional outcomes. Ramsay and Silverman (2005) developed a PCA methodology to study the simultaneous variation of more than one functional outcomes. Dubin and Müller (2005) defined dynamical correlation for multiple functional outcomes. Zhou et al. (2008) used a functional data approach to analyze paired longitudinal data based on a mixed-effect model framework. Li et al. (2014) considered the correlation between mixed continuous and discrete functional outcomes. Cao et al. (2017) developed a functional mapping method to detect quantitative trait loci (QTLs) that simultaneously control multiple dynamic traits.

One of the main challenges in analyzing functional data with multiple outcomes is the number of parameters that need to be estimated. As the number of outcomes grows, the traditional method of moment estimators might not be practical. We overcome this challenge with two important tools: (1) using the copula approach introduced by Staicu et al. (2012) so that the marginal distributions of the data and its dependence structure can be modeled separately, and, (2) under the Gaussian copula assumption, applying the Karhunen–Loève (KL) expansion for Gaussian processes to write the data processes as a sum of a small number of components. The dependence between outcomes can be modeled through the dependence between their components. The resulting covariance structure of these components is similar to those in the study by Zhou et al. (2008) under a mixed-effect model framework. We develop a similar EM algorithm as done by Zhou et al. (2008) to calculate maximum likelihood estimators of the copula parameters. Our algorithms are fast and do not require separated steps to model the dependence within each outcome and between different outcomes. This allows the use of a bootstrapping technique to obtain confidence intervals for the estimates. We also explain the steps needed to predict an outcome when all of the other outcomes are known, along with its prediction intervals.

The article is organized as follows. We briefly explain the process for estimating marginal distributions and present the model for the dependence structure in Sect. 2. In Sect. 3, the EM and prediction algorithms are presented. The finite-sample performance of the proposed method is evaluated by simulation studies in Sect. 4. The proposed method is demonstrated by analyzing the DTI data in Sect. 5. Section 6 concludes with a discussion including possible extensions.

2. MODELING MULTIPLE OUTCOMES

Let $\{Y_{ip}(t_{ij}); t_{ij} \in \mathcal{T}\}$ with $j = 1, 2, \dots, m_i$ be the data from outcome p , $p = 1, 2, \dots, P$ for subject i , $i = 1, 2, \dots, N$, observed at a grid points $\{t_{i1}, \dots, t_{im_i}\} \in \mathcal{T}$, where \mathcal{T} is a compact interval. Following the copula approach introduced by Staicu et al. (2012), we suppose that

$$Y_{ip}(t) = \mu_p(t) + \sigma_p(t)G^{-1}\{W_{ip}(t); \alpha_p(t)\} \quad (1)$$

where $\mu_p(t)$ is the mean function and $\sigma_p(t)$ is the standard deviation function of outcome p . Here $W_{ip}(t)$ is a latent process such that for each t , $W_{ip}(t)$ is uniformly $(0,1)$ distributed, $G(\cdot, \alpha)$ is a cumulative distribution function (CDF) of the skew-normal or skew-t distribution with zero mean, unit variance, and shape parameter α , and $G^{-1}(\cdot, \alpha)$ is the inverse of $G(\cdot, \alpha)$. In the case of the skew normal, the shape parameter is the skewness parameter, while in the case of the skew-t distribution the shape parameters are the skewness parameter and the degree of freedom parameter, which determines tail weight (Azzalini 1985; Azzalini and Capitanio 2003). For example, the CDF of the skew-normal distribution with zero mean, unit variance, and skewness parameter α is $G(\cdot, \alpha) = \Phi(\cdot) - 2T(\cdot, \alpha)$, where $\Phi(\cdot)$ is the CDF of the standard normal distribution, and $T(\cdot, \alpha)$ is the Owen's T function (Owen 1956); see Eq. (2.37) of Azzalini (2014). As commented by one reviewer, model (1) is similar to a linear model with errors distributed according to $G(\cdot, \alpha)$, although, unlike a linear model, $\mu_p(t)$ is model nonparametrically. If $G(\cdot, \alpha)$ is the CDF of the skew-normal or skew-t distribution with zero mean, unit variance, and shape parameter α , $Y_{ip}(t)$ also has a skew-normal or skew-t distribution with mean $\mu_p(t)$, variance $\sigma_p(t)$, and shape parameter α .

We assume that $\mu_p(t)$, $\sigma_p(t)$ and $\alpha_p(t)$ vary smoothly with t . The main objective is to estimate these marginal parameter functions and the dependence structure within each outcome and across different outcomes. We take a two-step approach. In the first step, we estimate the marginal parameter functions $\mu_p(t)$, $\sigma_p(t)$, and $\alpha_p(t)$ for each outcome p . In the second step, we estimate the correlation function of $(W_{i1}(t), \dots, W_{iP}(t))$.

2.1. MODELING MARGINAL DISTRIBUTION

This section summarizes the method used to estimate the marginal parameter functions. For details, see Staicu et al. (2012). The estimation is done in two steps.

Step 1 For simplicity, suppose the data are observed on a common dense grid of points so that $t_{ij} = t_j$ for all i and $j = 1, 2, \dots, m$. In this first step, we obtain an undersmoothed estimates $\tilde{\mu}_p(t)$, $\tilde{\sigma}_p(t)$ and $\tilde{\alpha}_p(t)$ for $\mu_p(t)$, $\sigma_p(t)$, and $\alpha_p(t)$ by maximizing, for each j , the pointwise likelihood function

$$\ell\{\mu_p(t_j), \sigma_p(t_j), \alpha_p(t_j)\} = \sum_{i=1}^N \log \left[g \left\{ \frac{Y_{ip}(t_j) - \mu_p(t_j)}{\sigma_p(t_j)} \right\}; \alpha_p(t_j) \right] - \log\{\sigma_p(t_j)\} \quad (2)$$

where $g(y; \alpha) = \partial G(y; \alpha) / \partial y$ is the density function corresponding to the distribution function G . When G is assumed to be a skew-normal or skew-t distribution, the estimates $\tilde{\mu}_p(t)$, $\tilde{\sigma}_p(t)$ and $\tilde{\alpha}_p(t)$ can be computed using the functions `sn.mle` or `st.mle` from R package `sn` (Azzalini 2011). See Staicu et al. (2012) for a discussion for the case of sparse data and alternative methods.

Step 2 The estimates from step 1 are smoothed further using penalized splines. Write the mean function as $\mu_p(t) = B(t)^T \beta_{\mu,p}$, where $B(t)$ is a vector of spline functions evaluated at t and $\beta_{\mu,p}$ is a vector of spline coefficients. The estimates $\hat{\beta}_{\mu,p}$ for $\beta_{\mu,p}$ are obtained by minimizing the penalty criterion

$$\text{PL}_\mu(\beta_{\mu,p}) = \sum_{j=1}^m \left\{ \mu_p(t_j) - B_j^T \beta_{\mu,p} \right\}^2 + \lambda_{\mu,p} \Omega_{\mu,p}(\beta_{\mu,p}), \quad (3)$$

where $\Omega_{\mu,p} = \beta_{\mu,p}^T D_{\mu,p} \beta_{\mu,p}$, $B_j = B(t_j)$, and $D_{\mu,p}$ is a penalty matrix. See Ruppert et al. (2003) for a discussion of penalized splines and penalty matrices. The final estimates $\hat{\mu}_p$ is $B(t)^T \hat{\beta}_{\mu,p}$. The estimates for the standard deviation function can be computed similarly. For the shape parameter function, Staicu et al. (2012) proposed a penalized marginal pseudo-likelihood criterion where the mean and variance parameter functions are fixed at estimates. That is, the criterion to be minimized is

$$\text{PL}_\alpha(\beta_{\alpha,p}) = -2 \sum_{i=1}^N \sum_{j=1}^m [\ell_{ij}(\beta_{\alpha,p}; \hat{Y}_{ip}(t_j))] + \lambda_{\alpha,p} \Omega_{\alpha,p}(\beta_{\alpha,p}) \quad (4)$$

where $\hat{Y}_{ip}(t_j) = \{Y_{ip}(t_j) - \hat{\mu}_p(t_j) / \hat{\sigma}_p(t_j)\}$ are the standardized observations, $\ell_{ij}(\beta_{\alpha,p}; \hat{Y}_{ip}(t_j)) = \log[g\{\hat{Y}_{ip}(t_j); B_j^T \beta_{\alpha,p}\}]$, and $\alpha_p(t) = B(t)^T \beta_{\alpha,p}$.

2.2. MODELING THE DEPENDENCE ACROSS DIFFERENT OUTCOMES

We assume Gaussian copulas for simplicity because they have proved adequate for our purposes. After the estimates of the parameters of marginal distributions for each outcome $\hat{\mu}_p(t)$, $\hat{\sigma}_p(t)$ and $\hat{\alpha}_p(t)$ are obtained, we transform the observed outcome by

$$\hat{R}_{ip}(t) = \Phi^{-1} \left[G \left\{ \frac{Y_{ip}(t) - \hat{\mu}_p(t)}{\hat{\sigma}_p(t)}; \hat{\alpha}_p(t) \right\} \right], \quad (5)$$

where Φ is the cumulative distribution function of the standard normal distribution, and $G(\cdot; \alpha)$ is the distribution function of a distribution with mean 0, variance 1, and shape parameter α . As pointed out by a reviewer, except for the estimation errors in $\hat{\mu}_p(t)$, $\hat{\sigma}_p(t)$, and $\hat{\alpha}_p(t)$, $\hat{R}_{ip}(t)$ has a standard normal distribution, because $G \left\{ \frac{Y_{ip}(t) - \mu_p(t)}{\sigma_p(t)}; \alpha_p(t) \right\}$ is distributed in $U(0, 1)$.

We model \hat{R}_{ip} as the sum of two independent components: (1) a finite Karhunen–Loève (KL) expansion with a small number of components, and (2) a white noise process with variance $\sigma_{\epsilon_p}^2$. That is,

$$\hat{R}_{ip}(t) = \sum_{k=1}^{K_p} Z_{ipk} f_{kp}(t) + \epsilon_{ip}(t), \quad (6)$$

where $\{Z_{ipk}, k = 1, 2, \dots, K_p\}$ are normally distributed with mean 0 and are independent across k , and f_{kp} are the eigenfunctions of the covariance function of the process $\hat{R}_{ip}(t)$ (Ramsay and Silverman 2005). To have the marginal variance of $\hat{R}_{ip}(t)$ equal to 1, we require that for all t ,

$$\sum_{k=1}^{K_p} f_{kp}^2(t) \text{var}(Z_{ipk}) + \sigma_{\epsilon_p}^2(t) = 1. \quad (7)$$

For identifiability, we assume that the variances of $\{Z_{ikp}, k = 1, 2, \dots, K_p\}$ are in decreasing order.

The eigenfunctions f_{kp} can be estimated using splines as follows. Let $b(t) = \{b_1(t), \dots, b_q(t)\}^T$ be an orthonormal spline basis, so that $\int_{\mathcal{T}} b_k(t)b_l(t)dt = \delta_{kl}$ where δ_{kl} is the delta function. We approximate $f_p = (f_{1p}, \dots, f_{K_p,p})^T$ using $b(t)$ as

$$f_p(t)^T = b(t)^T \Theta_p \quad (8)$$

where Θ_p is a matrix of spline coefficients with dimension $q \times K_p$ with orthogonal columns. The orthogonality of $b(t)$ and Θ_p implies the orthogonality of the principal component curves f_{kp} .

Define $Z_{ip} = (Z_{ip1}, \dots, Z_{ipK_p})^T$, $Z_i = (Z_{i1}^T, \dots, Z_{iP}^T)^T$, and $D_p = \text{cov}(Z_{ip})$. D_p is diagonal since the Z_{ipk} are independent across k . For identifiability, we also have that the diagonal elements of D_p are in a decreasing order. Define $C_{pp'} = \text{cov}(Z_{ip}, Z_{ip'})$. Then we can write the covariance matrix Σ of Z_i as

$$\Sigma = \text{cov}(Z_i) = \begin{pmatrix} D_1 & C_{12} & \cdots & C_{1P} \\ C_{21} & D_2 & \cdots & C_{2P} \\ \vdots & \vdots & \ddots & \vdots \\ C_{P1} & C_{P2} & \cdots & D_P \end{pmatrix}. \quad (9)$$

Suppose that we have the same observation times $\{t_j, j = 1, 2, \dots, m\}$ for every subject i . Let $\widehat{R}_{ip} = \{\widehat{R}_{ip}(t_1), \dots, \widehat{R}_{ip}(t_m)\}^T$, $B = \{b(t_1), \dots, b(t_m)\}^T$, and $\epsilon_{ip} = \{\epsilon_{ip}(t_1), \dots, \epsilon_{ip}(t_m)\}^T$. Putting everything together, we have the following reduced-rank model for \widehat{R}_{ip} ,

$$\widehat{R}_{ip} = B\Theta_p Z_{ip} + \epsilon_{ip} \quad (10)$$

$$\epsilon_{ip} \sim N(0, \sigma_{\epsilon_p}^2 I_m), \quad Z_{ip} \sim N(0, D_p), \quad \text{cov}(Z_{ip}, Z_{ip'}) = C_{pp'}, \quad \text{for } p \neq p'.$$

Our reduced-rank model for the latent processes (10) has a similar dependence structure as the model for two outcomes introduced by Zhou et al. (2008). Here we adopt a similar approach using an iterative EM algorithm for estimating the parameters.

To ensure identifiability, it is sufficient to require that the first nonzero elements of each column of Θ_p are positive. With finite samples, it is best to determine the sign using the elements with the largest magnitude in each column of Θ_p , since this choice is least influenced by finite-sample random fluctuation (for details, see Zhou et al. 2008). In our estimation algorithm, we require that, in every iteration, the element with the largest magnitude in each column of Θ_p is positive. In the next section, we develop an algorithm for estimating parameters in model (10).

3. ESTIMATION ALGORITHMS

This section explains the criterion and algorithm for estimating parameters, along with a discussion of tuning parameter specification. We also explain how to predict an outcome after we have observed each of the other outcomes.

3.1. PARAMETER ESTIMATION ALGORITHM

The estimates of the dependence parameters can be obtained by maximizing the pseudo-likelihood. Let $L_i(\Lambda)$ denote the contribution to the likelihood from individual i , where Λ is the covariance matrix of $\widehat{R}_i = (\widehat{R}_{i1}^T, \dots, \widehat{R}_{ip}^T)^T$. The loglikelihood for individual i is given as

$$\log L_i(\Lambda) = -\frac{Pm}{2} \log(2\pi) - \frac{1}{2} \log |\Lambda| - \frac{1}{2} \widehat{R}_i^T \Lambda^{-1} \widehat{R}_i, \quad (11)$$

where Λ can be computed in terms of $(\{\Theta_p\}, \Sigma, \{\sigma_{\epsilon p}^2\})$ as

$$\text{cov}(\widehat{R}_i) = \Lambda = \begin{pmatrix} \Lambda_{11} & \Lambda_{12} & \cdots & \Lambda_{1P} \\ \Lambda_{21} & \Lambda_{22} & \cdots & \Lambda_{2P} \\ \vdots & \cdots & \ddots & \vdots \\ \Lambda_{P1} & \Lambda_{P2} & \cdots & \Lambda_{PP} \end{pmatrix}$$

where

$$\Lambda_{pp} = B\Theta_p D_p \Theta_p^T B^T + \sigma_{\epsilon p}^2 I_m \quad \text{and} \quad \Lambda_{pp'} = B\Theta_p C_{pp'} \Theta_{p'}^T B^T, \quad (12)$$

for $p = 1, \dots, P$. Even though the main interest here is to estimate the covariance matrix Λ of \widehat{R}_i , not the principal component functions which can be obtained from Θ_p , we use the reduced-rank model for \widehat{R}_{ip} so that our estimation algorithm involves estimating D_p , $C_{pp'}$ and Θ_p which are of much smaller size than Λ . The KL expansion also provides a nice diagonal structure for D_p . Our iterative algorithm estimates D_p and $C_{pp'}$ sequentially. Once the estimates for all of the parameters are obtained, the estimated covariance of $\widehat{R}_{ip}(t)$ can be computed by plugging in the estimates to Eq. (12).

Extending the approach as in Zhou et al. (2008), we use the penalized pseudo-likelihood and minimize

$$\begin{aligned} & -2 \sum_{i=1}^N \log L_i(\Lambda) + \sum_{p=1}^P \lambda_p \sum_{k=1}^{K_p} \Theta_{pk}^T \int b''(t) b''(t)^T dt \Theta_{pk} \\ & = \sum_{i=1}^N \left\{ Pm \log(2\pi) + \log |\Lambda| + \widehat{R}_i^T \Lambda^{-1} \widehat{R}_i \right\} + \sum_{p=1}^P \sum_{k=1}^{K_p} \lambda_{pk} \Theta_{pk}^T \int b''(t) b''(t)^T dt \Theta_{pk}, \end{aligned} \quad (13)$$

where Θ_{pk} denotes the k^{th} column of Θ_p , and λ_{pk} are the smoothing parameters. In general, we can have different λ_{pk} for each f_{pk} . Since smoothing parameters are mostly determined

by the scale of the data, to simplify the computation we assume that with suitable rescaling we can use $\lambda_{pk} = \lambda_p$ for all $k = 1, \dots, K_p$. Minimizing this expression can be complicated. Instead, we treat the Z_i as missing values and use the EM algorithm (Dempster et al. 1977). The joint loglikelihood is given as

$$\begin{aligned} \log L(\hat{R}_i, Z_i) &= \log \{f(\hat{R}_{i1}|Z_{i1}) \dots f(\hat{R}_{ip}|Z_{ip})f(Z_i)\} \\ &= \sum_{p=1}^P \left\{ \frac{-m}{2} \log(\sigma_{\epsilon p}^2) - \frac{1}{2\sigma_{\epsilon p}^2} (\hat{R}_{ip} - B\Theta_p Z_{ip})^T (\hat{R}_{ip} - B\Theta_p Z_{ip}) \right\} \\ &\quad - \frac{1}{2} \log(|\Sigma|) - \frac{1}{2} Z_i^T \Sigma^{-1} Z_i. \end{aligned} \quad (14)$$

The EM algorithm is performed as follows:

The E-step: Compute the conditional distribution of Z_i given \hat{R}_i

The M-Step: Update the parameter estimates by minimizing

$$-2E \left\{ \sum_{i=1}^N \log L(\hat{R}_i, Z_i) \middle| \hat{R}_i \right\} + \sum_{p=1}^P \lambda_p \sum_{k=1}^{K_p} \Theta_{pk}^T \int b''(t)b''(t)^T dt \Theta_{pk} \quad (15)$$

The details for each step are given in Section A of the supplementary file.

3.2. SPECIFICATION OF SPLINES AND TUNING PARAMETERS

Recall that we require the splines $b(t)$ defined by (8) and used to estimate the eigenfunctions $f_p(t)$ to be orthonormal. In the definition of orthogonality, we will replace the integral used to define the inner product between the splines by a sum over t_1, \dots, t_m . This implies that the $m \times q$ matrix $B = \{b(t_1), \dots, b(t_m)\}^T$ must have orthogonal columns, so that B is semi-orthogonal. As in Zhou et al. (2008), to obtain an orthogonal matrix B , we start by choosing any spline basis $\tilde{b}(t)$, such as a truncated power basis, and then we evaluate these functions at the observation points to obtain the matrix $\tilde{B} = \{\tilde{b}(t_1), \dots, \tilde{b}(t_m)\}^T$. Let $\tilde{B} = QR$ be the QR decomposition of \tilde{B} , where Q is a semi-orthogonal matrix and R is an upper triangular matrix. Then the semi-orthogonal matrix B is obtained as $B = \tilde{B}R^{-1} = Q$. R package `orthogonalsplinebasis` (Redd 2011) provides functions for orthogonalizing a spline basis.

The two sets of tuning parameters that need to be specified are the spline smoothing parameters, λ_p , and the number of components in the KL expansion, K_p . For penalized splines, the knots are typically placed at fixed quantiles of the observation times or locations, and the number of knots are not critical provided that it is sufficient large; therefore, the number of knots is not considered to be a tuning parameter. The smoothness of the curves is mainly determined by the smoothing parameters λ_p (Ruppert 2002; Ruppert et al. 2003). We start by performing the analysis for each individual outcome. For each outcome p , we specify a range of λ_p for the grid search. For each λ_p fixed at a value in the grid search range, we vary the number of components starting with only one component. The EM algorithm (as explained in Sect. 3.1) is used to fit each 1-outcome curve. We add components and

continue the EM algorithm until the variance of the newly added principal component score is less than a prespecified small fraction of that of the previously added component. The number of components before adding this last component is K_p chosen for this λ_p . The corrected AIC (Ruppert et al. 2003) is computed for this (λ_p, K_p) . The pair (λ_p, K_p) with the minimum corrected AIC is selected for this outcome. Other criteria such as CV and AIC can also be used. Once we analyze multiple outcomes, the number of components can be fine tuned further. We do not expect the values of λ to change from the individual to the multiple outcome case, as λ should depend on the scale of individual outcome data more than the dependence structure between outcomes.

3.3. CONFIDENCE INTERVALS FOR ESTIMATES

Bootstrapping can be easily applied to obtain pointwise confidence intervals for the estimates for the marginal distribution parameters and the correlations. Bootstrap samples are obtained by resampling the subjects from the original data $\{Y_{ip}\}$. For each bootstrap sample, we calculate estimates for the marginal distribution parameters (as explained in Sect. 2.1), use them to transform the data as in Eq. (5), and then use the EM algorithm to obtain estimates for the correlations by plugging in the parameter estimates to Eq. (12). The pointwise confidence intervals for the estimates of both the marginal distribution parameters and the correlations are computed based on sample quantiles of the estimates for the bootstrap samples.

3.4. PREDICTION OF AN OUTCOME USING ALL THE OTHER OUTCOMES

In addition to simplifying the computation for parameter estimation, our reduced-rank model for multiple outcomes also provides a straightforward way of predicting an outcome for an individual when all the other outcomes of the same subject are known. This can be done using the conditional distribution property of the multivariate normal random variables. Suppose we have the observation Y_{ip} of a subject i for all outcomes $p \neq 1$ and we want to predict outcome 1 for this individual. Let \widehat{R}_{ip} denote the transformed process as in Eq. (5). Since we have that $(\widehat{R}_{i1}, \dots, \widehat{R}_{iP}) \sim N(0, \Lambda)$, we can compute the conditional distribution of \widehat{R}_{i1} as

$$\widehat{R}_{i1} | \widehat{R}_{i2}, \dots, \widehat{R}_{iP} \sim N(\bar{\mu}_1, \bar{\Sigma}_1) \quad (16)$$

where

$$\bar{\mu}_1 = (\Lambda_{12} \ \dots \ \Lambda_{1P}) \begin{pmatrix} \Lambda_{22} & \dots & \Lambda_{2P} \\ \vdots & \ddots & \vdots \\ \Lambda_{2P} & \dots & \Lambda_{PP} \end{pmatrix}^{-1} \begin{pmatrix} \widehat{R}_{i2} \\ \vdots \\ \widehat{R}_{iP} \end{pmatrix},$$

and

$$\bar{\Sigma}_1 = \Lambda_{11} - (\Lambda_{12} \ \dots \ \Lambda_{1P}) \begin{pmatrix} \Lambda_{22} & \dots & \Lambda_{2P} \\ \vdots & \ddots & \vdots \\ \Lambda_{2P} & \dots & \Lambda_{PP} \end{pmatrix}^{-1} \begin{pmatrix} \Lambda_{21} \\ \vdots \\ \Lambda_{P1} \end{pmatrix}.$$

Then $\bar{\mu}_1$ can be used as a predictor for \widehat{R}_{i1} . A prediction \widehat{Y}_{i1} for Y_{i1} is obtained by transforming $\bar{\mu}_1$ back as

$$\widehat{Y}_{i1}(t) = G^{-1} \{ \Phi(\bar{\mu}_1); \widehat{\alpha}_1(t) \} \widehat{\sigma}_1(t) + \widehat{\mu}_1(t), \quad (17)$$

using the estimates $\widehat{\mu}_1(t), \widehat{\sigma}_1(t), \widehat{\alpha}_1(t)$ obtained earlier. The standard deviation of the prediction error can be computed by bootstrapping and used for a prediction interval. The prediction for other outcomes can be done similarly.

4. SIMULATION STUDIES

In this section, we demonstrate our methodology using simulated data. As far as we are aware, there exist no other estimators for multiple functional outcomes of the type we are considering, that is, with spatially varying non-Gaussian characteristics. Therefore, there are no competing estimators to compare with ours.

We simulated 100 datasets, each containing $P = 3$ outcomes on $N = 200$ subjects. The observations are taken at the common time points $t_j, j = 1, 2, \dots, 80$. The data are generated using the finite KL expansion for Gaussian processes as follows:

1. For $p = 1, 2, 3$, generate

$$Q_{ip}(t) = \sum_{k=1}^{K_p} Z_{ipk} h_{kp}(t) + \epsilon_{ip}(t), \quad (18)$$

where $K_1 = K_2 = 2, K_3 = 3, \epsilon_{ip}(t) \sim N(0, \tau_p^2), \tau_1^2 = 0.2, \tau_2^2 = 0.3, \tau_3^2 = 0.4$, $Z_i = (Z_{i11}, Z_{i12}, Z_{i21}, Z_{i22}, Z_{i31}, Z_{i32}, Z_{i33})$ is normal with mean 0 and covariance matrix

$$\begin{pmatrix} 6.47 & 0 & -4.46 & 1.53 & 2.93 & 0.31 & 1.76 \\ 0 & 4.08 & 0.02 & 0.89 & 0.99 & -2.39 & -1.14 \\ -4.46 & 0.02 & 20.03 & 0 & -3.52 & 2.61 & -4.05 \\ 1.53 & 0.89 & 0 & 7.46 & 3.51 & -2.01 & -1.43 \\ 2.93 & 0.99 & -3.52 & 3.51 & 16.34 & 0 & 0 \\ 0.31 & -2.39 & 2.61 & -2.01 & 0 & 7.63 & 0 \\ 1.76 & -1.14 & -4.05 & -1.43 & 0 & 0 & 2.98 \end{pmatrix}$$

The functions $h_{kp}(t)$ are the normalized version of $\sin(\pi t)$ and $\cos(\pi t)$ for $p = 1, t^2 - 13t + 43/6$ and $t + 0.5$ for $p = 2$, and $e^t, t - (e - 1)^{-1}$ and $t^2 - 1.0025t + 0.1654$ for $p = 3$. Then we transform $Q_{ip}(t)$ to

$$R_{ip}(t) = \frac{Q_{ip}(t)}{\sigma_p(t)}, \text{ where } \sigma_p^2(t) = \sum_{k=1}^{K_p} \text{var}(Z_{ipk}) h_{kp}^2(t) + \tau_p^2, \quad (19)$$

so that $R_{ip}(t)$ has a marginal $N(0, 1)$ distribution.

2. Generate the observations

$$Y_{ip}(t) = \mu_p(t) + \sigma_p(t)G^{-1}[\Phi\{R_{ip}(t)\}; \alpha_p(t)], \quad (20)$$

where $\Phi(\cdot)$ and $G(\cdot; \alpha)$ denote the cumulative distribution functions of the standard normal and skew-normal distribution with mean 0, variance 1 and shape parameter α , respectively. The mean and shape parameter functions are

$$\begin{aligned} \mu_1(t) &= 5, \mu_2(t) = -3t^5 + 4.5t^3 - 1.5t + 15, \mu_3(t) = 10e^{-t^2/2} + 20t^2 + 10, \\ \alpha_1(t) &= 0, \alpha_2(t) = 10 \sin(2\pi t), \alpha_3(t) = 48t^2 - 48t + 6 \end{aligned} \quad (21)$$

The plots of simulation results are shown in Figures S1–S4 in Section B of the supplementary file. The estimates of the mean, standard deviation and shape parameter functions of each outcome from all 100 datasets are shown in Figure S1. We assumed the skew-normal distribution. These estimates were obtained by maximizing the pointwise likelihood at each point t_j using the R package `sn` (Azzalini 2011). The estimates were then smoothed further using penalized truncated cubic polynomial splines. The smoothing parameters were selected using the REML criterion implemented in R package `mgcv` (Wood 2006). The black lines indicate the true functions. Even though there seems to be high variation in the estimates when $\alpha \approx 10$, the density of skew normal for α equal to 10 is not much different from when α is much higher. In fact, as $\alpha \rightarrow \infty$ the skew-normal distribution converges to the half-normal distribution, and when α is 10 it is already similar to the half normal. This means that even though the estimate of α in this range is not very accurate, we can still obtain a good estimate for the distribution.

The eigenfunctions in the Gaussian copula [see Eq. (6)] were estimated using the orthonormal transform of cubic splines with 9 knots. The number of components and smoothing parameters were chosen based on AIC. The transformation and parameter selection are explained in Sect. 3.2. Figure S2 shows the true and estimates of the covariance from 2 datasets. The true covariances within an outcome and cross-covariances between outcomes are shown in the first and fourth rows, respectively. The estimated covariances within an outcome from 2 datasets are shown in rows 2 and 3. The estimated cross-covariance between outcomes from the same 2 datasets are shown in rows 5 and 6.

Estimated variances of each latent process are shown in Figure S3. The values close to 1 indicate the closeness to our assumption that the latent processes have marginal variance of 1. The square root of the integrated mean square error (IMSE), integrated square bias (IBIAS) and integrated variance (IVAR) for the marginal parameter functions and covariance parameters are shown in Table S1 in Section B of the supplementary file. These quantities are calculated, for example, for the mean function by $\text{IMSE} = \int_0^1 \widehat{E} \{ \widehat{\mu}(t) - \mu(t) \}^2 dt$, $\text{IVAR} = \int_0^1 \widehat{\text{Var}} \{ \widehat{\mu}(t) \} dt$ and $\text{IBIAS} = \int_0^1 [\widehat{E} \{ \widehat{\mu}(t) \} - \mu(t)]^2 dt$. Here \widehat{E} and $\widehat{\text{Var}}$ denote the sample mean and sample variance, respectively. Overall, the bias and variance are low for the mean and standard deviation estimates. The bias for the shape parameter estimates is somewhat high, especially, as seen from Figure S1, at locations with high α (in absolute value). This should not affect our estimated marginal distributions too much as explained in Section B

of the Supplementary Materials. The covariance estimates have small IMSE. The contour plots of pointwise square root of the mean square error for the covariance estimates are shown in Figure S4.

5. APPLICATION TO DIFFUSION TENSOR IMAGING (DTI)

In this section, we demonstrate our methodology using the DTI data from the healthy and MS groups. Our data consists of measurements from 42 healthy controls and 162 MS patients taken at 93 locations along the corpus callosum tract. Figure 1 displays the 3 outcomes for both groups. A subset of our data set is freely available as the DTI data set in R's `refund` package Crainiceanu et al. (2012).

DTI is a magnetic resonance imaging technique that measures the diffusion of water in tissue. The anisotropy of water diffusion allows images of the white matter in the brain to be generated. White matter tracts are made up of axons that transmit signals between different regions of the brain. These axons are surrounded and insulated by myelin, a fatty substance facilitating rapid signal transmission. Multiple sclerosis is an autoimmune disease associated with damage to myelin and can lead to significant disabilities in patients (Goldsmith et al. 2012).

There have been many other studies using the DTI tractography for MS patients. For instance, Reich et al. (2005) analyzed various indices obtained from DTI tractography to identify abnormalities in brain tissues of MS patients. Goldsmith et al. (2011a,b, 2012) studied DTI data using penalized functional regression. Longitudinal functional PCA (LFPCA) was introduced by Greven et al. (2010) to account for the longitudinal variability of tractography data in MS. LFPCA extended the MFPCA approach of Di et al. (2009) who focused on replicated functional data. See Staicu et al. (2010) for the application of PCA for complex multilevel spatially correlated functional data. McLean et al. (2014) introduced a functional generalized additive model and applied it to the relationship between a cognitive test score and DTI tractography.

DTI provides many measurements of water diffusion. In this study, we consider 3 measurements: (1) fractional anisotropy, (2) parallel diffusivity and (3) perpendicular diffusivity. Parallel and perpendicular diffusivities are diffusion parallel and perpendicular to the long axis of a fiber bundle, respectively, while fractional anisotropy is a contrast between diffusion in the two directions. Mathematically, at each location, DTI tractography can be described as a 3×3 symmetric, positive definite matrix. Denoting the eigenvalues of the matrix by $\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$, parallel and perpendicular diffusivities are λ_1 and $\frac{1}{2}(\lambda_2 + \lambda_3)$, respectively. The fractional anisotropy is

$$\left[\frac{3 \{(\lambda_1 - \bar{\lambda})^2 + (\lambda_2 - \bar{\lambda})^2 + (\lambda_3 - \bar{\lambda})^2\}}{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)} \right]^{1/2}, \quad (22)$$

where $\bar{\lambda} = (\lambda_1 + \lambda_2 + \lambda_3)/3$ (McLean et al. 2014).

Our objective is to identify the neuronal tract locations where the abnormalities in brain tissues occur and also explain the characteristics of the abnormalities. The intention is to

use differences between controls and diseased subjects in any DTI characteristics to help understand the disease process. These differences should be useful for detecting MS and for monitoring the progress of a treatment. We found the same result as in Staicu et al. (2012) for parallel diffusivity. Each of the mean, variance, and skewness functions differs between the two groups. For fractional anisotropy and perpendicular diffusivity, only the mean and variance functions are significantly different. The correlations within each outcome are similar for the two groups.

Our most important finding is that the cross-correlation between fractional anisotropy and parallel diffusivity is different between the MS and control groups in the middle part of the tract (roughly from tract locations 0.2–0.8). The cross-correlation in this tract section is slightly positive in the control group, while it is negative in the MS group. Bootstrapping confirms that the difference is significant. The difference between the two groups in the cross-correlation between the fractional anisotropy and the perpendicular diffusivity is less apparent, while the cross-correlation between the perpendicular and parallel diffusivities is similar in most parts of the tract.

5.1. MARGINAL PARAMETER ESTIMATION FOR DTI DATA

We assumed that each measurement from each group follows a skew-t marginal distribution. We assumed that the mean, standard deviation, and skewness parameter vary along the tract, while the degrees of freedom are constant along the tract locations. (There were insufficient observations to estimate spatially varying degrees of freedom parameters.) Using R package `sn` (Azzalini 2011), we obtained the maximum likelihood estimates for the degree of freedom for the fractional anisotropy, parallel diffusivity and perpendicular diffusivity equal to 24.93, 11,112 and 16.30, respectively, for the control group, and 11,732, 5.98 and 8.47 for the MS group. We used the skew-normal distribution to model any outcome with the estimated degrees of freedom > 100 . The undersmoothed estimates for the mean, standard deviation, and skewness parameters were obtained via pointwise maximum likelihood estimation using function `st.mle()` in the R package `sn` (Azzalini 2011). The mean and standard deviation estimates were smoothed further by penalized spline fitting using generalized additive model implemented in package `mgcv` in R (Wood 2006). The skewness parameters were smoothed further by fitting a penalized spline to maximize the penalized pseudo-likelihood computed using the estimates of the mean and standard deviation as in Eq. (4).

The estimated mean, variance and skewness functions for the three outcomes for both groups are shown in Figure S5 in the supplementary file. The bootstrap 90% pointwise confidence intervals for the differences of the marginal parameters between the two groups are shown in Fig. 2. The mean, variance, and skewness functions are all statistically significantly different between the two groups for parallel diffusivity, while only the mean and variance are significantly different for fractional anisotropy and perpendicular diffusivity. The mean of fractional anisotropy for the MS group is higher than the control group, but the means of parallel and perpendicular diffusivities are lower in the MS group. The variances for the MS group are higher in all of the three outcomes. The skewness of the fractional anisotropy is close to zero for both groups. The parallel diffusivity of the MS group is positively skewed

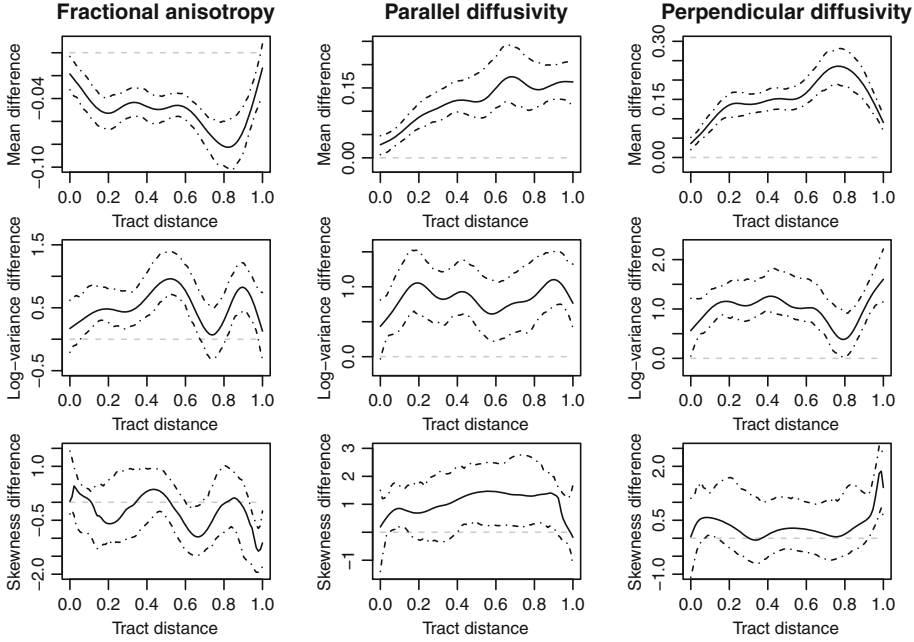


Figure 2. Estimated differences (black solid lines) in the mean, log-variance and skewness between the healthy and MS groups. The black dashed lines show 90% confidence intervals computed using bootstrapping with 1000 samples. The gray dashed horizontal lines are through 0.

while the skewness is close to zero in most part of the tract in the control group. The difference in the skewness between the two groups is significant. The perpendicular diffusivity is positively skewed in most part of the tract in both groups. There is no significant difference in the skewness between the two groups.

5.2. DEPENDENCE STRUCTURE OF DTI DATA

After the estimates $\hat{\mu}_p$, $\hat{\sigma}_p$, $\hat{\alpha}_p$ for the mean, standard deviation and skewness parameters for the outcome $p = 1, 2, 3$ for each group were obtained, the data were transformed using Eq. (5). Here, G is the skew-t distribution with mean 0, variance 1, skewness parameter $\hat{\alpha}_p$ and degree of freedom for each outcome and each group as stated earlier. The method explained in Sect. 3 was applied to the transformed data. The transformation implies that the marginal variance of the transformed process is 1, and, hence, the covariance function for the transformed process is also its correlation function. In our study, estimates of the marginal variance vary from 0.6 to 1.3 for the MS group, and 0.6–1.5 for the control group, due to the small sample sizes. We verified this by a simple simulation as follows. We generated data with 42 observations for control group, and 162 observations for MS group according to the skew-t distribution using the estimated parameters obtained earlier. Then we transformed the simulated data as in Eq. (5) and computed the marginal sample variances. The variances vary similarly to those of the DTI data. The sample variances are noticeably closer to 1 as

the number of observations for each group increases to 200. Most of the sample variances lie well within $[0.9, 1.1]$ as the number of observations increases to 500.

The estimated correlations for the DTI data were obtained by converting the estimated covariances for the transformed process (i.e., scale the covariances by the marginal variances). Figure S6 in the supplementary file displays the estimated correlations within each of the three outcomes for the healthy and MS groups. Figure S7 in the supplementary document displays the estimated cross-correlations between different outcomes for the healthy and MS groups. The correlations within the outcomes look somewhat similar for the two groups, while there is a noticeable difference between the two groups for the cross-correlations, especially for the cross-correlation between parallel diffusivity and fractional anisotropy. The pointwise confidence intervals for the estimated correlations can be computed using bootstrapping. Figure S8 in the supplementary document shows the significance levels of the difference of the correlations between the two groups calculated from the 90% and 95% bootstrap pointwise confidence intervals based on 1000 samples.

The bootstrapping results confirm that there is no significant difference in the correlations within fractional anisotropy and parallel diffusivity. The correlation within perpendicular diffusivity is different in various small regions of the tract. There is a significant difference between the two groups in the cross-correlation between the fractional anisotropy and the parallel diffusivity in a large middle part of the tract (at around locations 0.2–0.8; see the bottom, left panel of Figure S8). The cross-correlation is slightly positive in the control group, while it is negative in the MS group. The difference in the cross-correlation between the fractional anisotropy and the perpendicular diffusivity is significant in a smaller part of the tract (at around locations 0.4–0.8; see the bottom, middle panel of Figure S8). The cross-correlation is more negative in the MS group in this region. The cross-correlation between the parallel and perpendicular diffusivities is somewhat similar between the two groups with the parallel diffusivity at around location 0.2 being more positively-correlated with the overall perpendicular diffusivity in the MS group.

5.3. CASE STATUS PREDICTION FOR DTI DATA

We study the prediction accuracy of our model using a leave-one-out analysis. To predict the case status for subject i , we obtain estimates for marginal and dependence parameters using all of the subjects but subject i . We then obtain the transformed process $\hat{R}_{ip}^{(MS)}, \hat{R}_{ip}^{(C)}$, $p = 1, \dots, P$, defined in Eq. (5) for subject i using the obtained parameter estimates for MS and control groups, respectively. Because we have that $(\hat{R}_{i1}^{(MS)}, \dots, \hat{R}_{iP}^{(MS)}) \sim N(0, \Lambda^{(MS)})$ and $(\hat{R}_{i1}^{(C)}, \dots, \hat{R}_{iP}^{(C)}) \sim N(0, \Lambda^{(C)})$, where $\Lambda^{(MS)}$ and $\Lambda^{(C)}$ are the covariance matrix estimates for MS and control groups, respectively, we can calculate the likelihood functions $L(\hat{R}_{i1}^{(MS)}, \dots, \hat{R}_{iP}^{(MS)})$ and $L(\hat{R}_{i1}^{(C)}, \dots, \hat{R}_{iP}^{(C)})$. To predict the case status, we compare the two likelihood functions and predict this subject to be an MS patient if $L(\hat{R}_{i1}^{(MS)}, \dots, \hat{R}_{iP}^{(MS)}) > L(\hat{R}_{i1}^{(C)}, \dots, \hat{R}_{iP}^{(C)})$. As suggested by reviewers, we also compare our joint model with the single functional outcome's marginal model. The marginal model only considers one functional outcome, which predicts the i -th subject to be an MS patient if $L(\hat{R}_{ip}^{(MS)}) > L(\hat{R}_{ip}^{(C)})$, $p = 1, \dots, P$. In this DTI data set, our joint model using all

three functional outcomes predicts 100% MS patients to belong to the MS group, while the marginal model only predicts 76.5%, 86.4%, and 86.4% MS patients to belong to the MS group when using the single functional outcome: fractional anisotropy, parallel diffusivity and perpendicular diffusivity, respectively.

A classification method to identify patients would use as much information as available, and the cross-correlations between the three DTI functions would be only one piece of information. As we have seen, the amount of skewness and the variance of DTI signals differ between patients and control and would also be informative. Cognitive tests such as the PASAT provide additional information (Ozturk et al. 2010; McLean et al. 2014), since MS impairs cognition. Structural MRI can detect lesions and a lumbar puncture can detect abnormalities in antibodies. Structural MRI and the observations of symptoms are the most likely way that the disease is diagnosed. There are a number of symptoms such as numbness and weakness of the limbs, loss of vision, tremors, unsteady gait, slurred speech, and fatigue. The differences in cross-correlation that we found could be used to provide additional power to classify cases and controls, although DTI data are not normally collected on patients.

These data were collected in an effort to understand the disease and, ultimately, to improve treatment and perhaps develop a cure. MS can be treated, but treatment only slows the disease's progression and there is no cure at this time. Therefore, the goal here was not to develop a classification methodology per se, but rather to provide information to researchers that could help them understand the underlying biology of the disease.

6. DISCUSSION

We have developed a copula-based approach for modeling functional data with multiple outcomes. The copula approach, together with the KL expansion, significantly reduces the dimension of the model and allows fast estimation of the marginal distributions and the dependence structure. This is essential in the analysis of functional data, especially in the case of multiple outcomes where the data set can be very large.

Our methodology performed very well in our simulation study. We extended the study done by Staicu et al. (2012) to include three outcomes from DTI study and were able to identify the differences in the pointwise marginal distributions between the control and MS groups. Our dependence structure study shows that the correlations within the same outcome are not different between the two groups, while the cross-correlation between outcomes, especially between the parallel diffusivity and fractional anisotropy, are different.

One possible extension is to study functional data in a longitudinal setting. For example, the data used in this study, which contains only the data from the first visit of each subject, was taken from a larger dataset that contains measurements from subjects' multiple visits. This data set was studied by, for example, Goldsmith et al. (2012) and Greven et al. (2010). Goldsmith et al. (2012) studied the relation between cognitive disability in MS patients by incorporating DTI tractography measurement from multiple visits as functional covariates in a longitudinal penalized functional regression setting. Greven et al. (2010) developed a longitudinal functional principal component analysis approach to decompose the longitudinal DTI data into a time-dependent population average, baseline subject-specific variability,

longitudinal subject-specific variability, subject/visit-specific variability and measurement error. A longitudinal extension to our methodology would allow one to assess how the point-wise marginal distributions of each outcome and/or the dependence within and across the outcomes across tract locations of the subjects from each group change over time.

ACKNOWLEDGEMENTS

The authors are very grateful for the constructive comments of the Editor, the Associate Editor and two reviewers, which are extremely helpful for us to improve our work. The authors also thank Daniel Reich and Peter Calabresi and their research teams, who were instrumental in collecting the data for this study. Scans were funded by grants from the National Multiple Sclerosis Society and EMD Serono. We are grateful to Ciprian Crainiceanu for providing access to the data and for meaningful discussions and personal communications. Carroll was supported by a Grant U01-CA057030 from the National Cancer Institute (NCI). Ruppert was supported by the NCI Grant U01-CA057030 and a NSF Grant AST-1312903.

[Received November 2017. Accepted November 2018. Published Online December 2018.]

REFERENCES

- Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics* **12**, 171–178.
- (2011). *R Package ‘sn’: The skew-normal and skew-t distributions (version 0.4-17)*. URL <http://azzalini.stat.unipd.it/SN>, accessed August 1, 2012.
- (2014). *The Skew-Normal and Related Families*, Institute of Mathematical Statistics Monographs, Cambridge University Press.
- Azzalini, A. and Capitanio, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t distribution. *Journal of the Royal Statistical Society, Series B* **65**, 367–389.
- Basser, P. J., Mattiello, J., and LeBihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* **66**, 259–267.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., and Aldroubi, A. (2000). In vivo fiber tractography using dt-mri data. *Magnetic Resonance in Medicine* **44**, 625–632.
- Cao, J., Wang L., Huang, Z., Gai, J., Wu, R. (2017). Functional Mapping of Multiple Dynamic Traits. *Journal of Agricultural, Biological and Environmental Statistics* **22**, 60–75.
- Crainiceanu, C. M., Reiss, P., Goldsmith, J., Huang, L., Huo, L., Scheipl, F. et al. (2012). *R Package ‘refund’: Regression with functional data (version 0.1-6)*. URL <http://cran.r-project.org/web/packages/refund/index.html>, accessed 2010.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* **39**, 1–38.
- Di C., Crainiceanu C.M., Caffo B.S., and Punjabi N.M. (2009). Multilevel Functional Principal Component Analysis. *The Annals of Applied Statistics* **3**, 458–488.
- Dubin, J. A. and Müller, H. G. (2005). Dynamical correlation for multivariate longitudinal data. *Journal of the American Statistical Association* **100**(471), 872–881.
- Ferraty, F. and Romain, Y. (2010). *The Oxford Handbook of Functional Data Analysis*. New York: Oxford University Press.
- Goldsmith, J., Feder, J., Crainiceanu, C. M., Caffo, B., and Reich, D. S. (2011a). Penalized functional regression. *Journal of Computational and Graphical Statistics* **20**, 830–851.

- Goldsmith, J., Crainiceanu, C. M., Caffo, B., and Reich, D. S. (2011b). Penalized functional regression analysis of white-matter tract profiles in multiple sclerosis. *NeuroImage* **57**, 431–439.
- (2012). Longitudinal penalized functional regression for cognitive outcomes on neuronal tract measurements. *Journal of the Royal Statistical Society, Series C* **61**, 453–469.
- Greven, S., Crainiceanu, C. M., Caffo, B., and Reich, D. S. (2010). Longitudinal functional principal component analysis. *Electronic Journal of Statistics* **4**, 1022–1054.
- Li, H., Staudenmayer, J. and Carroll, R. J. (2014). Hierarchical functional data with mixed continuous and binary measurements. *Biometrics* **70**(4), 802–811.
- McLean, M. W., Hooker, G., Staicu, A., Scheipl, F., and Ruppert, D. (2014). Functional generalized additive models. *Journal of Computational and Graphical Statistics* **23**(1), 249–269.
- Owen, D. B. (1956). Tables for computing bivariate normal probabilities. *Annals of Mathematical Statistics* **27**, 1075–1090.
- Ozturk, A., Smith, S.A., Gordon-Lipkin E.M., Harrison, D.M., Shiee, N., Pham, D.L. et al. (2010). MRI of the corpus callosum in multiple sclerosis: association with disability. *Multiple Sclerosis* **16**, 166–177.
- Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. New York: Springer.
- Redd, A. (2011). *R Package ‘orthogonalsplinebasis’: Orthogonal bspline basis functions (version 0.1.5)*. URL <http://osplinebasis.r-forge.r-project.org>, accessed February 1, 2013.
- Reich, D. S., Smith, S. A., Zackowski, K. M., Gordon-Lipkin, E. M., Jones, C. K., Farrel, J. A.D. et al. (2005). Multiparametric magnetic resonance imaging analysis of the corticospinal tract in multiple sclerosis. *NeuroImage* **38**, 271–279.
- Ruppert, D. (2002). Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics* **11**, 735–757.
- Ruppert, D., Wand, M. P., and Carroll, R. J. (2003). *Semiparametric Regression*, Cambridge: Cambridge University Press.
- Staicu, A.-M., Crainiceanu, C. M., and Carroll, R. J. (2010). Fast methods for spatially correlated multilevel functional data. *Biostatistics* **11**, 177–194.
- Staicu, A., Crainiceanu, C. M., Reich, D. S., and Ruppert, D. (2012). Modeling functional data with spatially heterogeneous shape characteristics. *Biometrics* **68**, 331–343.
- Tsing, T. and Eubank, R. (2015) *Theoretical Foundations of Functional Data Analysis with an Introduction to Linear Operators*, Wiley.
- Wood, S. N. (2006). *Generalized Additive Models: An Introduction with R*. Boca Raton, FL: Chapman & Hall/CRC.
- Zhou, L., Huang, J. Z., and Carroll, R. J. (2008). Joint modelling of paired sparse functional data using principal components. *Biometrika* **95**, 601–619.