Methods for functional regression and nonlinear mixed-effects models with applications to PET data

Yakuan Chen

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2017

ABSTRACT

Methods for functional regression and nonlinear mixed-effects models with applications to PET data

Yakuan Chen

The overall theme of this thesis focuses on methods for functional regression and nonlinear mixed-effects models with applications to PET data.

The first part considers the problem of variable selection in regression models with functional responses and scalar predictors. We pose the function-on-scalar model as a multivariate regression problem and use group-MCP for variable selection. We account for residual covariance by "pre-whitening" using an estimate of the covariance matrix, and establish theoretical properties for the resulting estimator. We further develop an iterative algorithm that alternately updates the spline coefficients and covariance. Our method is illustrated by the application to two-dimensional planar reaching motions in a study of the effects of stroke severity on motor control.

The second part introduces a functional data analytic approach for the estimation of the IRF, which is necessary for describing the binding behavior of the radiotracer. Virtually all existing methods have three common aspects: summarizing the entire IRF with a single scalar measure; modeling each subject separately; and the imposition of parametric restrictions on the IRF. In contrast, we propose a functional data analytic approach that regards each subject's IRF as the basic analysis unit, models multiple subjects simultaneously, and estimates the IRF nonparametrically. We pose our model as a linear mixed effect model in which shrinkage and roughness penalties are incorporated to enforce identifiability and smoothness of the estimated curves, respectively, while monotonicity and non-negativity constraints impose biological in-

formation on estimates. We illustrate this approach by applying it to clinical PET data.

The third part discusses a nonlinear mixed-effects modeling approach for PET data analysis under the assumption of a compartment model. The traditional NLS estimators of the population parameters are applied in a two-stage analysis, which brings instability issue and neglects the variation in rate parameters. In contrast, we propose to estimate the rate parameters by fitting nonlinear mixed-effects (NLME) models, in which all the subjects are modeled simultaneously by allowing rate parameters to have random effects and population parameters can be estimated directly from the joint model. Simulations are conducted to compare the power of detecting group effect in both rate parameters and summarized measures of tests based on both NLS and NLME models. We apply our NLME approach to clinical PET data to illustrate the model building procedure.

Table of Contents

List of Figures

List of Tables

Acknowledgments

It is my pleasure to express the deepest appreciation to people who provided kind help and support during my doctoral study.

First and foremost, I am most grateful to my advisors, Dr. R. Todd Ogden and Dr. Jeff Goldsmith, for offering enormous effort and encouragement during my doctoral study. I acquired solid fundamentals from their classes in Theory of Statistical Inference and Linear Regression Models. In addition, the weekly event of the Functional Data Analysis Working Group they host broadens my knowledge in my research area tremendously. They provided inspiring advice and insightful criticisms throughout my thesis work. Without their incredible guidance and help, the accomplishment of my dissertation would not be possible.

Then, I would like to express my sincere gratitude to my dissertation committee members, Dr. Min Qian, Dr. Francesca Zanderigo and Dr. Gen Li for offering valuable thoughts and suggestions to help me improve my thesis. I would also like to thank the Biostatistics Department for providing me fellowship and Dr. John L. P. Thompson for supervising my research assisantship in the past two years.

Last but not least, I would like to thank my parents and friends for their incredible and unconditional support.

Chapter 1

Introduction

In recent years, data collected in various fields of study tend to be high dimensional and complicated in structure. In many situations, observations can be regarded as functions. As a result, functional data analysis (FDA), an important tool for analyzing functional data (Ramsay and Silverman, 2005), has received a great deal of attention.

Functional data are data that have functional form and vary over a continuum, where the continuum is often time, but may also be location, probability, etc. The scope of functional data includes one-dimensional curves and two- or three- dimensional images in which pixel or voxel intensities can be viewed as functions on spatial positions. A few examples of functional data include the classical Canadian weather data (Ramsay and Silverman, 2005), human growth data (Ramsay and Silverman, 2005), diffusion tensor imaging (DTI; McLean et al. 2014), electroencephalography (EEG; Di et al. 2009) and positron emission tomography (PET; Reiss and Ogden 2010).

1.1 Basic Tools for Analyzing Functional Data

Functional data differ from multivariate data in that functional data have natural ordering among the observations. Although in practice each functional observation

is recorded discretely at time or location points, it is assumed that functional data arise from some underlying smooth functions or processes, i.e.

$$y_{ij} = x_i(t_{ij}) + \epsilon_{ij}$$

where $x_i(\cdot)$ is a smooth underlying function or process that is observed on a grid $\{t_{ij}\}_{j=1}^{J_i}$.

As functional data are observed at a finite discrete grid and often with measurement error, a basic idea behind functional data analysis is to represent discrete data in terms of smooth functions. In terms of data representation, basis expansion and roughness penalization are often considered (Müller and Yao, 2008). Basis expansion involves projecting data onto a functional basis. A function $x(t)$ can be expanded as

$$x(t) = \sum_{j=1}^{K} c_j \phi_j(t) = \mathbf{\Phi}(t)\mathbf{c}$$

where $\mathbf{\Phi}(t) = (\phi_1(t), ..., \phi_K(t))$ is a pre-specified or data-driven basis and $\mathbf{c} = (c_1, ..., c_K)^T$ is the vector of coefficients. Common choices of a pre-specified basis include splines, wavelets, and Fourier bases which are independent of observed data. On the contrary, functional principal components analysis provides a data-driven basis which is estimated from data and the expansion of function is an approximation as long as finite number of basis functions are used. On the other hand, in order to penalize the roughness of functions, people can add smoothing penalties in the framework of least square estimation. In other words, estimated function $x(t)$ can be obtained by minimizing

$$\sum_{i=1}^{n} (y_i - x(t_i))^2 + \lambda \int [Lx(t)]^2 dt$$

where $Lx(t)$ measures the roughness of $x(t)$ and λ is the smoothing parameter. The roughness penalty based on the second derivative is the most commonly used in modern statistics literature, although the method can easily be adapted to penalties based on other derivatives.

1.1.1 Splines

Splines are piecewise-defined polynomial functions that are continuous and smooth at the knots. A spline of order M is constructed of piecewise order M polynomials and has continuous derivatives up to order $M-2$ at the knots (de Boor, 1978). For example, cubic splines, splines with order 4, have continuous first and second derivatives. In order to define a spline function, one needs to determine the order of polynomial, the number of knots, and their locations. One simple way of selecting knots is to place them at some pre-determined percentiles of the observations.

Functions can be represented using a spline basis. Among the many equivalent bases, B-spline bases are widely used since any function of a given order can be uniquely expressed as a linear combination of B-splines of that order. A B-spline basis includes piecewise polynomial functions that are defined over adjacent intervals spanned by the knots, each one having a local support. This structure leads to to a highly sparse design matrix, which is computationally favorable, especially when the number of knots is large.

Regularization, on the other hand, is necessary to control the complexity of fit. A smoothing spline estimate is defined as the minimizer $\hat{x}$ (over the class of twice differentiable functions) of the penalized residual sum of squares

$$\sum_{i=1}^{n}(y_i - x(t_i))^2 + \lambda \int [x''(t)]^2 dt$$

where λ is the tuning parameter that establishes a trade-off between closeness to the data and roughness of the function estimate. It has been shown that the minimizer is a natural cubic spline with knots at $t_i, i = 1, ..., n$ (de Boor, 1978). Natural cubic splines are cubic splines with additional constraint that the function is linear beyond the boundary knots. The tuning parameter λ can be chosen by cross-validation, generalized cross-validation (GCV) or restricted maximum likelihood (REML) when

connected with mixed model.

1.1.2 Wavelets

In contrast to splines, which often work well for smooth data, wavelets are well suited for describing functions with localized small scale components including jumps, spikes and peaks. One appealing property of wavelets is that they are capable of representing functions well in a sparse way, namely with few coefficients. This property makes wavelets useful for signal processing, especially in denoising and compression of signals and images.

Wavelet bases consist of functions with varying scales and locations. There are several families of wavelet functions. The Haar wavelet basis is the simplest wavelet basis since it produces a piecewise constant representation. However, it is not widely used in practice as it's neither continuous nor differentiable. As a contrast, other members of the Daubechies wavelet families are more popular since they give a smoother representation.

Due to its simplicity, the Haar basis is always a good example to illustrate how a wavelet basis is constructed. Suppose V_j is the space consisting of functions that are piecewise-constant over intervals of form $[2^{-j}k, 2^{-j}(k+1)]$, $j, k \in \mathbb{Z}$. If $\phi(t) = \mathrm{I}_{[0,1)}(t)$, then $\phi_{0,k}(t) = \phi(t-k), k \in \mathbb{Z}$ form a orthonormal basis for V_0. By dilations and translations of $\phi(t)$, also known as the father wavelet, we can obtain functions

$$\phi_{j,k}(t) = 2^{\frac{j}{2}}\phi(2^j t - k)$$

Then $\phi_{j,k}(t), k \in \mathbb{Z}$ form a orthonormal basis for V_j. In fact, the V_j spaces are nested, i.e. $V_{-1} \subset V_0 \subset V_1$ and they are also called approximation spaces since any function in $L^2(\mathbb{R})$ can be approximated by $\phi_{j,k}(t), k \in \mathbb{Z}$.

The wavelet function $\psi_{j,k}(t)$ is defined as

$$\psi_{j,k}(t) = \frac{1}{\sqrt{2}}\phi_{j+1,2k}(t) - \frac{1}{\sqrt{2}}\phi_{j+1,2k+1}(t)$$

Likewise $\psi_{j,k}(t), k \in \mathbb{Z}$ form an orthonormal basis for a space W_j, which is difference between successive approximation spaces V_{j+1} and V_j. On the other hand, the wavelet function $\psi_{j,k}(t)$ can also be generated by the mother wavelet $\psi(t) = \phi(2t) - \phi(2t-1)$ through dilations and translations

$$\psi_{j,k}(t) = 2^{\frac{j}{2}}\psi(2^j t - k)$$

One can show that the set of functions $\psi_{j,k}(t), k \in \mathbb{Z}$ form an orthonormal basis for $L^2(\mathbb{R})$ as well. Hence, any square-integrable function $x(t)$ can be expressed as

$$x(t) = \sum_j \sum_k d_{j,k}\psi_{j,k}(t)$$

where $\psi_{j,k}(t)$ is a wavelet function with dilation index j and translation index k ($j, k \in \mathbb{Z}$) and $d_{j,k}$ is the corresponding wavelet coefficient that can be computed as

$$d_{j,k} = \langle x, \psi_{j,k} \rangle = \int x(t)\psi_{j,k}(t)dt$$

Alternatively, given the basis $\phi_{j_0,k}(t)$, $\{\psi_{j,k}(t)\}_{j=j_0,\, j_0+1,\, \ldots}$, $k \in \mathbb{Z}$, $x(t)$ can also be represented as

$$x(t) = \sum_{k\in\mathbb{Z}} c_{j_0,k}\phi_{j_0,k}(t) + \sum_{j=j_0}^{\infty}\sum_{k\in\mathbb{Z}} d_{j,k}\psi_{j,k}(t)$$

where $c_{j_0,k} = \langle x, \phi_{j_0,k} \rangle$ and $d_{j,k} = \langle x, \psi_{j,k} \rangle$. If $x(t)$ is defined on the interval [0, 1], then

$$x(t) = \sum_{k=0}^{2^{j_0}-1} c_{j_0,k}\phi_{j_0,k}(t) + \sum_{j=j_0}^{\infty}\sum_{k=0}^{2^j-1} d_{j,k}\psi_{j,k}(t)$$

In practice, the multiresolution analysis above provides a good way of decomposing functions observed on a discrete grid in the wavelet basis space, which is known as

discrete wavelet transformation (DWT). Suppose the function is observed at $N = 2^J$ equally spaced time points on the interval [0, 1], denoted as $\boldsymbol{x} = \{x(t_j)\}_{j=1,2,3,\ldots,N}$. The DWT of $\boldsymbol{x}$ is

$$\boldsymbol{w} = \boldsymbol{W}\boldsymbol{x}$$

where $\boldsymbol{W}$ is an orthogonal $N \times N$ matrix associated with the pre-specified orthonormal wavelet basis. $\boldsymbol{w}$ is a $N \times 1$ vector with the wavelet coefficients of $\boldsymbol{x}$. It consists of the scaling coefficients $\mu_{j_0,k}$, $k = 0, \ldots, 2^{j_0} - 1$ associated with the scaling wavelet, and the "oscillation" coefficients $\nu_{j,k}$, $j = j_0, \ldots, J-1$, $k = 0, \ldots, 2^j - 1$ associated with the mother wavelet (Abramovich et al., 2000). For simplicity, we often time set $j_0 = 0$, in which case there is one scaling coefficients $\mu_{0,0}$ and $2^J - 1$ "oscillation" coefficients $\nu_{j,k}$, $j = 0, \ldots, J-1$, $k = 0, \ldots, 2^j - 1$. On the other hand, due to the orthonormality conditions, $\mu_{j_0,k}$ and $\nu_{j,k}$ are related to $c_{j_0,k}$ and $d_{j,k}$ via

$$\mu_{j_0,k} \approx \sqrt{n} c_{j_0,k} \qquad \nu_{j,k} \approx \sqrt{n} d_{j,k}$$

In contrast with decomposing a function, people can reconstruct a function using a wavelet basis as well. The reconstruction of $\boldsymbol{x}$ from $\boldsymbol{w}$, known as the inverse discrete wavelet transformation (IDWT) is simply given by

$$\boldsymbol{x} = \boldsymbol{W}^T \boldsymbol{w}$$

In practice, both DWT and IDWT can be performed through a fast $O(N)$ algorithm (Mallat, 1989), which is based on the two-scale relationship such that the wavelet coefficients at one level can always be computed using only the coefficients from another level without integration. In particular, each step of the algorithm involves the recursive application of some low- and high-pass filters.

A key property of wavelet analysis is that one can represent a function in terms of a relatively small number of coefficients (others are zero), namely in a sparse way. However, when the function is contaminated with noise, it yields a problem of

estimating N parameters using N data points. Clearly, some regularization is needed here. Consider the model

$$y_i = x(t_i) + \epsilon_i \qquad i = 1, \ldots, n$$

where $\epsilon_i \sim_{i.i.d.} N(0, \sigma^2)$. The goal is to estimate the function x using the noisy data y. Let $\hat{\boldsymbol{w}} = \{\hat{w}_{j,k}\}_{j=0,\ldots,J-1,\, k=0,\ldots,2^j-1}$ denote the vector of coefficients obtained by DWT of the observed $\{y_i\}_{i=1,\ldots,n}$. The regularization is accomplished by thresholding and/or shrinkage. Donoho and Johnstone (1994) suggested the hard and soft thresholding

$$\text{Hard thresholding:} \qquad \tilde{w}_{j,k} = \hat{w}_{j,k}\, I(|\hat{w}_{j,k}| > \lambda)$$

$$\text{Soft thresholding:} \qquad \tilde{w}_{j,k} = sgn(\hat{w}_{j,k})\, (|\hat{w}_{j,k} - \lambda|)_+$$

Then we can reconstruct $\hat{x}_i(t)$ by IDWT of the thresholded wavelet coefficient $\{\tilde{w}_{j,k}\}_{j=0,\ldots,J-1,\, k=0,\ldots,2^j-1}$. One common choice for the threshold λ is the universal threshold

$$\lambda = \sigma\sqrt{2\log n}$$

where σ may be estimated by the median absolute deviation method (Donoho and Johnstone, 1994):

$$\hat{\sigma} = \frac{\sqrt{n}\,\text{median}(|w_{J-1,k} - \text{median}(w_{J-1,k})|)}{0.6745}$$

1.1.3 Functional Principal Components

Another useful method to represent and analyze functional data is functional principal components analysis (FPCA). Unlike splines and wavelets, functional principal components analysis provides a data-driven basis and is ideally able to capture the major directions of variability in the data.

Suppose $x(t)$ is a square-integrable random function with a mean function $\mu(t)$ and a covariance function $\sigma(s,t) = E[(x(s) - \mu(s))(x(t) - \mu(t))]$. Then by the Karhunen-Loeve decomposition, covariance function can be expressed as

$$\sigma(s,t) = \sum_{j=1}^{\infty} \rho_j \nu_j(s) \nu_j(t)$$

where $\nu_j(t), j = 1, 2, ...$ are the principal component eigenfunctions that form an orthonormal basis and $\rho_j, j = 1, 2, ...$ are the corresponding eigenvalues. Both eigenfunctions and eigenvalues satisfy the equation

$$\int \sigma(s,t)\nu_j(t)dt = \rho_j \nu_j(s)$$

Additionally, the total variance of $x(t)$ is given as

$$\int var[x(t)]\, dt = \sum_{k=1}^{\infty} \rho_k$$

and the proportion of total variance "explained" by eigenfunction $\nu_j(t)$ is $\frac{\rho_j}{\sum_j \rho_j}$ (Ramsay and Silverman, 2005). In practice, we only need finitely many principal components to approximate the function $f(t)$. The number of principal components may be pre-specified or chosen by cross-validation.

A function $x(t)$ can be expanded as a linear combination of eigenfunctions

$$x(t) = \mu(t) + \sum_{j=1}^{\infty} \xi_j \nu_j(t)$$

where ξ_j is the functional principal component score which can be obtained by

$$\xi_j = \int \nu_j(t)[x(t) - \mu(t)]\, dt$$

In general, there are two approaches to estimate the eigenfunctions and eigenvalues. Suppose that $\boldsymbol{X}$ is a $n \times p$ matrix whose rows are functions $x_i(t)$, $i = 1, ..., n$ observed on a discrete grid of p points and all the functions have been centered.

The sample variance-covariance matrix $\hat{\Sigma} = n^{-1}\boldsymbol{X}^T\boldsymbol{X}$ is a $p \times p$ matrix with entries $\hat{\sigma}(t_k, t_l)$ where

$$\hat{\sigma}(t_k, t_l) = \frac{1}{n}\sum_{i=1}^{n} x_i(t_k)x_i(t_l)$$

Rao (1958) and Tucker (1958) applied multivariate principal components analysis to the observed functions, which yield the eigenvectors $\tilde{\nu}_j$ and eigenvalues $\tilde{\rho}_j$. Then the eigenfunctions $\hat{\nu}_j(t)$ are obtained by interpolating $\tilde{\nu}_j$ by applying some smoothing techniques.

Another approach involves basis function expansion of $x_i(t)$ and $\nu_j(t)$. Suppose that $\boldsymbol{\Phi}(t) = (\phi_1(t), ..., \phi_K(t))^T$ is a pre-defined basis and each function has basis expansion

$$x_i(t) = \sum_{k=1}^{K} c_{ik}\phi_k(t) \qquad \nu_j(t) = \sum_{k=1}^{K} b_{jk}\phi_k(t)$$

Then $\boldsymbol{X} = \boldsymbol{C}\boldsymbol{\Phi}$ and $\nu_j(t) = \boldsymbol{\Phi}^T(t)\boldsymbol{b}_j$, where $\boldsymbol{C}$ is the $n \times K$ coefficient matrix with entries c_{ik} and $\boldsymbol{b}_j = (b_{j1}, ..., b_{jK})^T$. The covariance function is given in matrix terms as

$$\sigma(s,t) = E[\boldsymbol{\Phi}^T(s)\boldsymbol{C}^T\boldsymbol{C}\boldsymbol{\Phi}(t)]$$

which yields

$$\begin{aligned}\int \sigma(s,t)\nu_j(t)dt &= E\left\{\int \boldsymbol{\Phi}^T(s)\boldsymbol{C}^T\boldsymbol{C}\boldsymbol{\Phi}(t)\boldsymbol{\Phi}^T(t)\boldsymbol{b}_j dt\right\} \\ &= E\left\{\boldsymbol{\Phi}^T(s)\boldsymbol{C}^T\boldsymbol{C}W\boldsymbol{b}_j\right\} \\ &= \rho_j\nu_j(s) = \rho_j\boldsymbol{\Phi}^T(s)\boldsymbol{b}_j\end{aligned}$$

where $W = \int \boldsymbol{\Phi}(t)\boldsymbol{\Phi}^T(t)\,dt$. If $\boldsymbol{\Phi}(t)$ is an orthonormal basis, $W = I$, where I is the identity matrix. Hence, $\hat{\boldsymbol{b}}_j$ can be obtained by solving

$$\frac{1}{n}\boldsymbol{C}^T\boldsymbol{C}W\boldsymbol{b}_j = \rho_j\boldsymbol{b}_j$$

and thereby $\hat{\nu}_j(t) = \boldsymbol{\Phi}^T(t)\hat{\boldsymbol{b}}_j$.

In addition, it is necessary to control the roughness of the estimated principal component functions by incorporating regularizations. One way is to maximize the sample variance of the principal component score with a penalty term included (Rice and Silverman, 1991), i.e.

$$\hat{\nu}_j(t) = \underset{\nu_j}{\operatorname{argmax}} \left\{ Var\left[\int x(t)\nu_j(t)\,dt\right] - \lambda_j \int \nu_j''(t)^2\,dt \right\}, \quad \text{subject to} \int \nu_j(t)^2\,dt = 1$$

where λ_j are the tuning parameters. The principal component functions can be estimated successively. Alternatively, Silverman (1996) proposed the penalized sample variance, which is defined as

$$\text{PCAPSV}(\nu_j) = \frac{Var[\int x(t)\nu_j(t)\,dt]}{\int \nu_j(t)^2\,dt + \lambda \int \nu_j''(t)^2\,dt}$$

Then ν_j can be estimated by maximizing PCAPSV(ν_j), subject to two constraints: i) $\int \nu_j(t)^2\,dt = 1$; and ii) $\int \nu_j(t)\nu_k(t)\,dt + \int \nu_j''(t)^2\nu_k''(t)^2\,dt = 0$.

1.2 Functional Regression Models

After a review of the techniques developed for expressing functional data, we now focus on building regression models that contain functional variables. There are generally three scenarios of regression modeling for functional data: scalar-on-function regression, function-on-scalar regression and function-on-function regression. For simplicity, we assume that the mean curve has been subtracted from each function in the following discussion.

1.2.1 Scalar-on-function Regression

Given observed data $(y_i, x_i(t)), i = 1, ..., n, t \in [0, 1]$ where y_i is a scalar response for subject i and $x_i(t)$ is the functional predictor, a scalar-on-function linear regression

model (Ramsay and Dalzell, 1991) is constructed as

$$y_i = \beta_0 + \int x_i(t)\beta(t)dt + \epsilon_i \tag{1.1}$$

where $\beta(t)$ is the coefficient function that determines the effect of $x_i(t)$ on y_i and ϵ_i is the error term. James (2002) and Müller and Stadtmüller (2005) extended it to a functional generalized linear model by incorporating a link function.

In practice, the functions are observed at finitely many points. Both functional predictors and coefficient function are treated as vectors of the same length. Then it becomes an ordinary multiple regression problem with model $y_i = \beta_0 + \boldsymbol{\beta^T X_i} + \epsilon_i$. However, it is oftentimes a $p > n$ problem where p is the number of predictors and n is the number of observations. Therefore, in order to get reasonable fits, dimension reduction and/or some regularity are required.

One basic approach to the estimation of the coefficient function is projecting coefficient function on some basis and converting it to a standard regression problem. Let

$$\beta(t) = \sum_{k=1}^{K} \eta_k B_k(t)$$

where $\{B_k(t), k = 1, ..., K\}$ is a K-dimensional basis, in this case

$$y_i = \beta_0 + \int x_i(t) \sum_{k=1}^{K} \eta_k B_k(t)dt + \epsilon_i = \beta_0 + \sum_{k=1}^{K} \eta_k \left(\int x_i(t) B_k(t)dt \right) + \epsilon_i$$

and then η_k can be estimated using least squares with smoothing constraint or other explicit penalties.

A common approach that has been widely used is functional principal component regression (FPCR), which is based on functional principal component analysis (FPCA) discussed above. The coefficient function $\beta(t)$ is expanded using K eigenfunctions $\hat{\nu}_k(t), k = 1, ..., K$ derived from $x_i(t)$, i.e., $x_i(t) \approx \sum_{k=1}^{K} c_{ik}\hat{\nu}_k(t)$ and $\beta(t) = \sum_{k=1}^{K} \eta_k \hat{\nu}_k(t)$. The number of principal components K may be determined by

cross validation or a pre-specified proportion of explained variance. The orthonormality of FPC basis guarantees that $\int x_i(t)\beta(t)dt = \sum_{k=1}^{K} c_{ik}\eta_k$. Then Model (1) turns out to be a multiple linear regression with FPC scores c_{ik} as predictors.

Alternatively, the penalized spline-based approach enforces smoothness of the coefficient function by imposing a roughness penalty. $\beta(t)$ is expanded using spline basis and estimate is obtained by minimizing $\sum(y_i - \beta_0 - \int x_i(t)\beta(t)dt)^2 + \lambda \int (\beta''(t))dt$. The tuning parameter, λ may be chosen by cross-validation, generalized cross validation or restricted maximum likelihood.

On the other hand, Zhao et al. (2012) proposed a wavelet-based approach. Similarly to FPCR, both functional predictor and coefficient function are expanded using an orthonormal wavelet basis. Then Model (1) becomes simply a multiple regression model. Lasso penalty is applied to perform variable selection and a sparse solution is obtained.

1.2.2 Function-on-scalar Regression

Suppose observed data are given as $(y_i(t), \boldsymbol{x}_i), i = 1, ..., n$ where $y_i(t)$ is a functional response for subject i and $\boldsymbol{x}_i$ is a p-vector representing the scalar predictors. One can construct a function-on-scalar regression model as

$$y_i(t) = \boldsymbol{x}_i^T \boldsymbol{\beta}(t) + \epsilon_i(t)$$

where $\boldsymbol{\beta}(t) = (\beta_1(t), ..., \beta_p(t))^T$ is a functional vector and $\epsilon_i(t)$ is the error function which is often assumed to be drawn from a stochastic process with expectation zero.

Ramsay and Silverman (2005) outlined an approach to estimate $\boldsymbol{\beta}(t)$ based on penalized ordinary least squares. The functional response $y_i(t)$ and coefficient functions $\beta_j(t)$ are both expanded using some functional basis, i.e.,

$$y_i(t) = \sum_{k=1}^{K} c_{ik}\phi_k(t) \qquad \beta_j(t) = \sum_{k=1}^{K} b_{jk}\phi_k(t)$$

The model thereby reduces to a multiple regression problem

$$\boldsymbol{C} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{E}$$

where $\boldsymbol{X}$ is the $n \times p$ design matrix; $\boldsymbol{C}$ is a $n \times K$ matrix with elements $\{c_{ik}\}_{i=1,\ldots,n;\, k=1,\ldots,K}$; and $\boldsymbol{B}$ is a $p \times K$ matrix with elements $\{b_{jk}\}_{j=1,\ldots,p;\, k=1,\ldots,K}$. In order to penalize roughness, $\boldsymbol{B}$ is chosen to minimize

$$||\boldsymbol{C} - \boldsymbol{X}\boldsymbol{B}||^{\boldsymbol{2}} + \lambda \boldsymbol{B}^{\boldsymbol{T}}\boldsymbol{P}\boldsymbol{B}$$

where $\boldsymbol{B}^{\boldsymbol{T}}\boldsymbol{P}\boldsymbol{B}$ approximates $\int (\boldsymbol{\beta}''(t))^2 dt$, where $\boldsymbol{\beta}(t) = (\beta_1(t), \ldots, \beta_p(t))$. The solution turns out to have a similar form as the generalized ridge regression estimator (Reiss et al., 2010).

Alternatively, rather than response in basis coefficient form, one can estimate $\boldsymbol{\beta}(t)$ based on raw response as well. Let $\boldsymbol{\Theta}$ denote the $D \times K$ matrix whose columns correspond to the K basis functions. We then express $\boldsymbol{\beta}(t)$ as $\boldsymbol{B}\boldsymbol{\Theta}^T$ where $\boldsymbol{B}$ is the $p \times K$ matrix of basis coefficients and the jth row corresponds to $\beta_j(t)$. Additionally let $\boldsymbol{Y}$ be the $n \times D$ matrix whose rows are functional outcomes observed on a grid of D points, $\boldsymbol{X}$ be the $n \times p$ design matrix and $\boldsymbol{\epsilon}$ be the $n \times D$ error matrix, our model becomes

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B}\boldsymbol{\Theta}^{\boldsymbol{T}} + \boldsymbol{\epsilon}$$

Estimation of $\boldsymbol{B}$ requires vectorizing both sides of equation. $\text{vec}(\boldsymbol{Y}^T)$ is the vector formed by concatenating the rows of $\boldsymbol{Y}$ and $\text{vec}((\boldsymbol{X}\boldsymbol{B}\boldsymbol{\Theta}^T)^T) = (\boldsymbol{X} \otimes \boldsymbol{\Theta})\text{vec}(\boldsymbol{B}^T)$, where $\otimes$ represents the Kronecker product of two matrices. Hence, $\text{vec}(\boldsymbol{B}^T)$ can be estimated by minimizing

$$||\text{vec}(\boldsymbol{Y}^T) - (\boldsymbol{X} \otimes \boldsymbol{\Theta})\text{vec}(\boldsymbol{B}^T)||^2 + \text{vec}(\boldsymbol{B}^T)^T \boldsymbol{P}_\Lambda \text{vec}(\boldsymbol{B}^T)$$

where $\boldsymbol{P}_\Lambda$ is the penalty matrix parameterized by $\Lambda = (\lambda_1, \ldots, \lambda_p)$. It is easy to obtain $\hat{\boldsymbol{B}}$ by rearranging $\text{vec}(\hat{\boldsymbol{B}}^T)$.

Additionally, Reiss et al. (2010) extends the model above to a penalized generalized least squares model and performs a fast automatic selection of multiple smoothing parameters.

1.2.3 Function-on-function Regression

Function-on-function regression may be used to study the association between one functional response and one or more functional predictors. Let $y_i(t)$ be the functional response for subject i and $x_i(s)$ be a functional predictor. The simplest model with one functional predictor is given as

$$y_i(t) = \beta_0(t) + \int \beta(t,s)x_i(s)ds + \epsilon_i(t)$$

where $\beta(t,s)$ is the coefficient function which in this case is a two-dimensional surface and $\epsilon_i(t)$ is a mean zero random stochastic process. Notice that $y_i(t)$ and $x_i(s)$ may be defined on different domains.

Several methods have been developed to fit the function-on-function regression model. Yao et al. (2005a) used principal component expansions for both functional predictor and the coefficient function. Ivanescu et al.(2013) proposed a penalized function-on-function regression method using mixed model representation of penalized regression. They expanded the coefficient surface in a bivariate basis and approximate $\int \beta(t,s)x_i(s)ds$ with Riemann sums on a fine grid. A penalty term was also added to enforce some amount of smoothness.

Chapter 2

Variable Selection in Function-on-Scalar Regression

2.1 Introduction

Regression models with functional responses and scalar predictors are routinely encountered in practice. These models face a challenge that also arises for traditional models: how to identify the important predictors among a potentially large collection. Functional-response models face the additional challenges of high dimensionality and residual correlation. The purpose of this article is to address the current lack of methods for variable selection in this class of models.

Our work is motivated by two-dimensional planar reaching data. As an assessment of upper extremity motor control, stroke patients and healthy controls made repeated reaching movements from a central point to eight targets arranged on a circle. The dataset consists of 57 subjects, including 33 patients suffering a unilateral stroke (meaning only one arm is affected) and 24 healthy controls, and contains motions made with both the dominant and non-dominant hands to each of the eight targets.

Our analytic goal is to explore the effects of the potential predictors of motor control on these motions and to identify the most essential ones using variable selection. Among the potential predictors, the Fugl-Meyer score is a quantity that measures the severity of arm motor impairment (Fugl-meyer et al., 1975). It ranges from 0 to 66 with smaller values indicating more severe impairment and 66 indicating healthy function. Other potentially important predictors include target direction, whether the hand used was the dominant or non-dominant, and whether the hand used was contralesional (directly affected by the stroke) or ipsilesional (indirectly affected or unaffected).

Figure 2.1 shows the observed reaching motions for three subjects: a stroke patient with contralesional dominant hand in the left column; a stroke patient with contralesional non-dominant hand in the center column and a heathy control in the right column. Reaching motions made by contralesional hand display deviation from straight paths from the starting point to each target; these deviations may be consistent for contralesional dominant or non-dominant hands. While deviation from straightness is not obvious in the ipsilesional arm, other effects, like over-reach, are observed. The potential for differential effects of stroke severity on reaching motions indicates the importance of allowing interactions between predictors of interest.

The observed data are horizontal and vertical coordinates of the hand position for each reaching motion as functions of time. We construct function-on-scalar regression models for the two outcome functions separately. Given scalar predictors x_{ij}, $i = 1, ..., n$, $j = 1, ..., p$ and functional responses $y_i(t)$, $i = 1, ..., n$, $t \in \mathcal{T}$, where $\mathcal{T}$ is some compact finite interval in $\mathbb{R}$, the linear function-on-scalar regression model is

$$y_i(t) = \beta_0(t) + \sum_{j=1}^{p} x_{ij}\beta_j(t) + \epsilon_i(t),\ i = 1, ..., n,\ t \in \mathcal{T} \tag{2.1}$$

where $\beta_j(\cdot), j = 0, ..., p$ are the $p + 1$ coefficient functions and $\epsilon_i(\cdot) \sim (0, \Sigma)$ is the error function drawn from a continuous stochastic process with expectation zero and

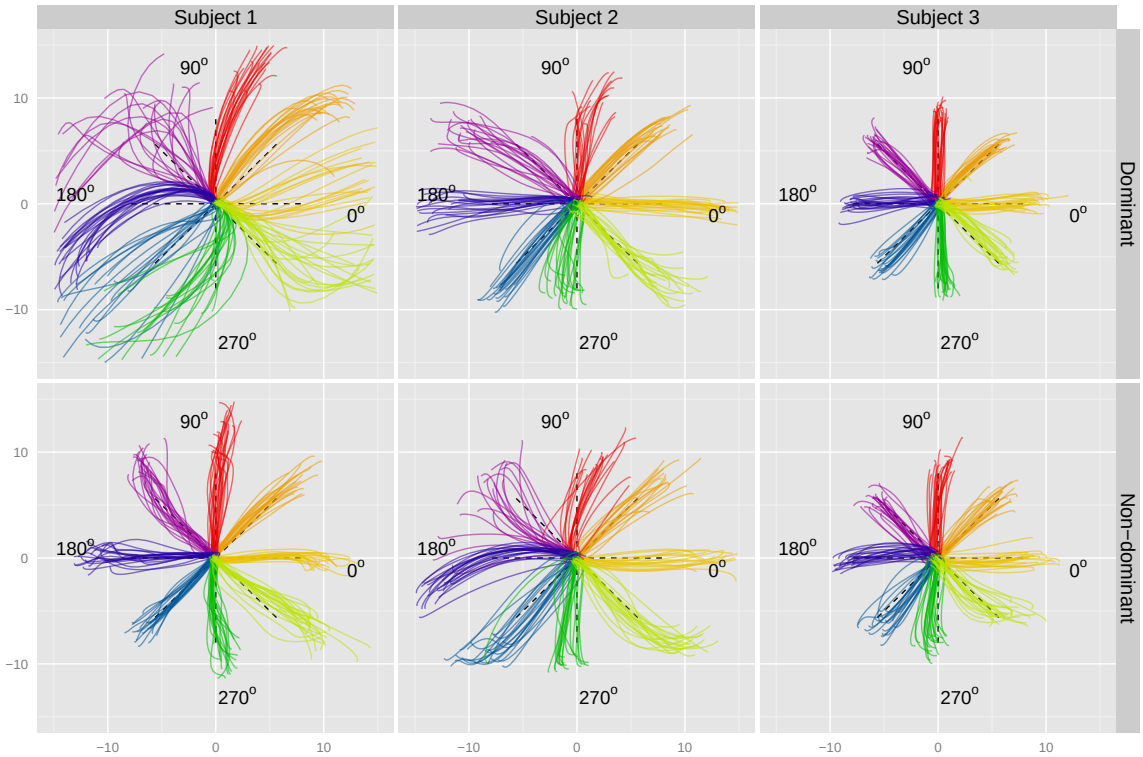

Figure 2.1: Observed reaching motions for three subjects. The top row shows the dominant hand and the bottom row shows the non-dominant hand of three subjects. The left column is a subject with a contralesional dominant hand. The center column is a subject with a contralesional non-dominant hand. The right column is a healthy control subject. Dashed lines are the straight paths to the eight targets terminating at the target location.

covariance function $\Sigma(s,t) = \text{cov}(\epsilon_i(s), \epsilon_i(t)),\ s,t \in \mathcal{T}$.

A common model fitting framework for function-on-scalar regression is outlined by Chapter 13 of Ramsay and Silverman (2005), in which the coefficient functions $\beta_j(\cdot)$ are expanded using some set of basis functions and basis coefficients are estimated using ordinary least squares. The imposition of quadratic roughness penalties to enforce smoothness of the estimated coefficient functions is also common. Reiss et al.

(2010) developed a fast automatic method for choosing tuning parameters in this model and accounted for correlated errors using generalized least squares. Goldsmith and Kitago (2015) develop a Bayesian approach that jointly models coefficient functions and the covariance structure, and applied their methods to the stroke kinematics dataset considered here.

When p is large, many scalar predictors may have no effect on the functional response and the corresponding coefficient functions would equal zero over all time points. In order to accurately identify the important predictors, we apply variable selection techniques when estimating the coefficient functions in Model (2.1). Since the coefficient functions are expanded using basis functions, the shape of each coefficient function is determined by a distinct group of basis coefficients. We therefore apply variable selection at the group level to include or exclude the vector of basis coefficients. The group lasso, proposed by Yuan and Lin (2006), is an extension of the classic lasso (Tibshirani, 1994) to the problem of selecting grouped variables. The lasso is known to induce biases in the included variables, so two alternative penalties, the smoothy clipped absolute deviation (SCAD) penalty (Fan and Li, 2001) and the minimax concave penalty (MCP) (Zhang, 2010), were proposed. These achieve consistency and asymptotic unbiasedness, and have been extended to grouped variable selection problem (Wang et al. (2007); Breheny and Huang (2013)).

Few approaches that consider variable selection in the context of functional regression models have been proposed in current literature. Wang et al. (2007) developed a penalized estimation procedure using group SCAD for variable selection in function-on scalar regression assuming errors $\epsilon_i(\cdot)$ are uncorrelated over their domain; this assumption is clearly violated in practice. Barber et al. (2015) presented Function-on-Scalar LASSO (FS-LASSO), a framework which extends the group LASSO to function-on-scalar regression; theory is developed for cases in which predictors are observed over dense or sparse grids. However, the bias for non-zero coefficients intro-

duced by LASSO was not addressed, and the method does not account for correlation among residual curves. Gertheiss et al. (2013) proposed a variable selection procedure for generalized scalar-on-function linear regression models, in which the predictors are in the form of functions and responses are scalar; though they also consider regression models for functional data, the structure of their models is very different from the one considered here.

We propose a method for variable selection in function-on-scalar regression that accounts for residual correction using tools from generalized least squares. We develop theory for this method and demonstrate its effectiveness in simulations that mimic our real-data application; direct comparisons with the method of Wang et al. (2007) and Barber et al. (2015) indicate superior performance of our proposed method for variable selection and prediction.

The rest of the article is organized as follows. In Section 2.2, we describe an estimation procedure for function-on-scalar regression models with errors that are uncorrelated over t using grouped variable selection methods. We then introduce our methods for the estimation of function-on-scalar regression models with correlated errors, including the development of an iterative method that refines the estimation of the error covariance and the variable selection. Simulations that resemble our motivating data examine and compare the numerical performance of competing methods in Section 2.3. An application of our method to the reaching motion data is given in Section 2.4. Finally, we present concluding remarks in Section 2.5. Our method is implemented in the user-friendly fosr.vs() function in the *refund* package (Ciprian Crainiceanu et al., 2014), available on CRAN, and code for our simulations is included in the supplementary material.

2.2 Methodology

2.2.1 Estimation for models with i.i.d. errors

Suppose $\{\phi_1(\cdot), ..., \phi_K(\cdot)\}$ is a set of pre-specified basis functions. The coefficient functions $\beta_j(\cdot)$, $j = 0, ..., p$ can be expanded as

$$\beta_j(\cdot) = \sum_{k=1}^{K} b_{jk}\phi_k(\cdot). \tag{2.2}$$

Hence, Model (2.1) is expressed as

$$y_i(t) = \sum_{k=1}^{K} b_{0k}\phi_k(t) + \sum_{j=1}^{p} x_{ij}\left(\sum_{k=1}^{K} b_{jk}x_{ij}\phi_k(t)\right) + \epsilon_i(t) \tag{2.3}$$

The problem is thereby reduced to estimating the basis coefficients $\{b_{jk}\}_{j=0,...,p;\, k=1,...,K}$. Functional basis should be chosen based on the properties of estimated curves. For instance, smooth basis, such as orthogonal polynomials or Fourier basis, is preferred when the estimated functions are assumed to be smooth. For implementations of our method in Section 2.3 and Section 2.4, we use the popular B-spline basis with pre-specified number of basis functions that implicitly determines the smoothness of the curves.

In practice, functions are observed on a discrete grid. For simplicity, we assume that the grid, denoted $\{t_1, ..., t_D\}$, is shared across subjects. Let $\boldsymbol{Y}$ be the $n \times D$ matrix whose rows are vector-valued functional responses; $\boldsymbol{\Phi}$ be the $D \times K$ matrix whose columns correspond to the K basis functions evaluated at $\{t_1, ..., t_D\}$; and $\boldsymbol{B}$ be the $(p+1) \times K$ matrix with jth row being the vector of basis coefficients for $\beta_j(\cdot)$. Then Model (2.3) can be expressed as

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B}\boldsymbol{\Phi}^T + \boldsymbol{E} \tag{2.4}$$

where $\boldsymbol{X}$ is the $n \times (p+1)$ design matrix and $\boldsymbol{E}$ is the $n \times D$ matrix containing vector-valued error functions.

Model (2.4) can be posed as a standard linear model in the following way. Let $\text{vec}(\boldsymbol{Y}^T)$ be the vector formed by concatenating the rows of $\boldsymbol{Y}$, and note that $\text{vec}\left((\boldsymbol{X}\boldsymbol{B}\boldsymbol{\Phi}^T)^T\right) = (\boldsymbol{X} \otimes \boldsymbol{\Phi})\text{vec}(\boldsymbol{B}^T)$ where $\otimes$ represents the Kronecker product of two matrices. Then

$$\text{vec}(\boldsymbol{Y}^T) = (\boldsymbol{X} \otimes \boldsymbol{\Phi})\text{vec}(\boldsymbol{B}^T) + \text{vec}(\boldsymbol{E}^T) \tag{2.5}$$

and $\text{vec}(\boldsymbol{B}^T)$ can be estimated using least squares. An estimate of $\hat{\boldsymbol{B}}$ is obtained by rearranging $\text{vec}\left(\hat{\boldsymbol{B}}^T\right)$.

To accurately identify the zero coefficient functions, we apply variable selection techniques when estimating $\text{vec}(\boldsymbol{B}^T)$ in Model (2.5). Let $\boldsymbol{B}_j$ be the vector of coefficients associated with the jth coefficient function $\beta_j(\cdot)$, specifically the jth row of $\boldsymbol{B}$. Note that the "zeroth" row of $\boldsymbol{B}$ corresponds to the intercept function $\beta_0(t)$, which we do not penalize. Setting the entire $\beta_j(\cdot)$ function to 0 is equivalent to setting all the entries of $\boldsymbol{B}_j$ to zero. Therefore, we apply variable selection techniques at the group level.

Variable selection can be achieved by penalizing the estimates of the coefficients. The general form of a group penalty is $\sum_{j=1}^{p} p_{\lambda,\gamma}(||\boldsymbol{B}_j||)$, where $p_{\lambda,\gamma}(\cdot)$ is the penalty function for the specific method and λ and γ are the tuning parameters. Therefore, the penalized estimator is obtained by minimizing

$$\frac{1}{2}\left\|\text{vec}(\boldsymbol{Y}^T) - (\boldsymbol{X} \otimes \boldsymbol{\Phi})\text{vec}(\boldsymbol{B}^T)\right\|^2 + nD\sum_{j=1}^{p} p_{\lambda,\gamma}(||\boldsymbol{B}_j||). \tag{2.6}$$

We use group MCP to perform variable selection; the penalty has the form

$$p_{\text{mcp}}(||\boldsymbol{B}_j||) = \begin{cases} \lambda||\boldsymbol{B}_j|| - \frac{||\boldsymbol{B}_j||^2}{2\gamma} & \text{if } ||\boldsymbol{B}_j|| \leq \gamma\lambda, \\ \frac{1}{2}\gamma\lambda^2 & \text{if } ||\boldsymbol{B}_j|| > \gamma\lambda \end{cases}$$

where λ and γ are tuning parameters. When $||\boldsymbol{B}_j||$ is small, the MCP penalty behaves exactly as lasso, but as $||\boldsymbol{B}_j||$ increases the amount of penalization is reduced

until there is no penalization at all, thereby avoiding bias in the estimate of large coefficients.

In terms of tuning parameter selection, γ is set to be 3 as recommended in Zhang (2010) and λ is chosen by cross-validation. Another parameter to be determined is K, the number of basis functions used in the expansion of the coefficient functions. In the following implementations of our method, a cubic B-spline basis with 10 basis functions was used. However, since we do not explicitly penalize the roughness of the estimated coefficient functions, the exact choice of K will vary from application to application and should be chosen with care.

2.2.2 Estimation for models with correlated errors

The estimation framework discussed in Section 2.2.1 assumes that errors are independent and identically distributed over the entire domain, and is similar to the framework of Wang et al. (2007). In most cases, however, within-function errors are correlated. Let $\boldsymbol{\Sigma}$ denote the $D \times D$ covariance matrix for discretely observed data. For estimation of the Model (2.4) with correlated errors, we use techniques from generalized least squares. If $\boldsymbol{\Sigma}$ is known, one can "pre-whiten" both sides of (2.4) with the lower triangular matrix $\boldsymbol{L}$ obtained by Cholesky decomposition of $\boldsymbol{\Sigma}$, i.e., $\boldsymbol{\Sigma} = \boldsymbol{L}\boldsymbol{L}^T$, to construct a new model

$$\boldsymbol{Y}^* = \boldsymbol{X}\boldsymbol{B}\boldsymbol{\Phi}^{*T} + \boldsymbol{E}^* \tag{2.7}$$

where $\boldsymbol{Y}^* = \boldsymbol{Y}(\boldsymbol{L}^{-1})^T$, $\boldsymbol{\Phi}^* = \boldsymbol{L}^{-1}\boldsymbol{\Phi}$ and the error $\boldsymbol{E}^* = \boldsymbol{E}(\boldsymbol{L}^{-1})^T$ is independent. Similarly, parameters in model (2.7) can be estimated by minimizing

$$\frac{1}{2}\left\|\mathrm{vec}(\boldsymbol{Y}^{*T}) - (\boldsymbol{X} \otimes \boldsymbol{\Phi}^*)\mathrm{vec}(\boldsymbol{B}^T)\right\|^2 + nD\sum_{j=1}^{p} p_{\lambda,\gamma}(||\boldsymbol{B}_j||). \tag{2.8}$$

For a given $\boldsymbol{\Sigma}$, the minimizer of (2.8) can be obtained using existing software by pre-whitening as described; our implementation is publicly available and uses the *grpreg* function in the *grpreg* package (Breheny and Huang, 2013).

The covariance matrix $\boldsymbol{\Sigma}$ is unknown in practice and it is necessary to obtain an estimate $\hat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$ and to use this estimate to pre-whiten data. To obtain this estimate, we first fit Model (2.5) using ordinary least squares under the assumption of independence; this provides an unbiased estimate $\hat{\boldsymbol{B}}$ of the coefficient matrix $\boldsymbol{B}$. From this model fit, we obtain the estimated residual matrix $\hat{\boldsymbol{E}} = \boldsymbol{Y} - \boldsymbol{X}\hat{\boldsymbol{B}}\boldsymbol{\Phi}^T$. Using $\hat{\boldsymbol{E}}$, we consider two approaches for estimating $\boldsymbol{\Sigma}$. The first, which we refer as the raw estimate, is constructed using a method-of-moments approach based on the residual matrix. The second approach uses functional principal component analysis (Yao et al., 2005b). Here, the off-diagonal elements of the raw covariance are smoothed and an eigen-decomposition of the resulting matrix is obtained. Our estimate is

$$\hat{\boldsymbol{\Sigma}} = \sum_{l=1}^{L} \hat{\lambda}_l \hat{\boldsymbol{\psi}}_l \hat{\boldsymbol{\psi}}_l^T + \hat{\sigma}^2 \boldsymbol{I} \tag{2.9}$$

where $\hat{\boldsymbol{\psi}}_1, ..., \hat{\boldsymbol{\psi}}_L$ are the estimated eigenfunctions over the grid $\{t_1, ..., t_D\}$, $\hat{\lambda}_l$, $l = 1, 2, ..., L$ are the corresponding eigenvalues, $\hat{\sigma}^2$ is the estimated measurement error variance and I is the identity matrix. The truncation level L is determined by the cumulative proportion of variability explained by eigenfunctions. This approach separates $\boldsymbol{\Sigma}$ into a smooth covariance over the observed grid and an additional uncorrelated measurement error process. The FPCA-based approach can also be applied to sparse data with irregular and unequal spaced grid (Yao et al., 2005b). But more parameters will be introduced in this scenario when estimating the subject-specific covariance matrices, which is more computationally intensive and time consuming. Although we focus on these methods for estimating $\boldsymbol{\Sigma}$, others that provide consistent estimators can be substituted.

2.2.3 Oracle properties of generalized group MCP estimator

We now discuss the theoretical properties of the method described in Section 2.2.2. Without loss of generality, we assume $\beta_0(t) = 0 \, \forall \, t \in \mathcal{T}$. We also assume the true coefficient functions $\beta_j(t)$ are in the space spanned by the set of basis functions $\boldsymbol{\Phi}$. Additionally, we assume the first s groups of coefficients, $\boldsymbol{B}_+ = (\boldsymbol{B}_1^T, ..., \boldsymbol{B}_s^T)^T$, are nonzero and the remaining $p - s$ groups of coefficients, $\boldsymbol{B}_0 = (\boldsymbol{B}_{s+1}^T, ..., \boldsymbol{B}_p^T)^T$, are zero. Let $(\boldsymbol{X} \otimes \boldsymbol{\Phi})_+$ denote the design matrix associated with $\boldsymbol{B}_+$ and $(\boldsymbol{X} \otimes \boldsymbol{\Phi})_0$ denote the one associated with $\boldsymbol{B}_0$. Therefore, we have $\boldsymbol{X} \otimes \boldsymbol{\Phi} = [(\boldsymbol{X} \otimes \boldsymbol{\Phi})_+ | (\boldsymbol{X} \otimes \boldsymbol{\Phi})_0]$ and $\boldsymbol{B} = [\boldsymbol{B}_+^T, \boldsymbol{B}_0^T]^T$. The additional assumptions required for the theorems are

1. $\lim_{n \to \infty} \frac{1}{nD} (\boldsymbol{X} \otimes \boldsymbol{\Phi})^T (\boldsymbol{X} \otimes \boldsymbol{\Phi})$ is a positive definite matrix;
2. $\lambda_n \to 0$ and $\sqrt{n}\lambda_n \to \infty$ as $n \to \infty$;
3. there exists a $\sqrt{n}$-consistent estimate $\hat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$;
4. the tuning parameter γ of the penalty is fixed.

Then we have the following results:

Theorem 1 (Estimation consistency)**.** *Under assumptions 1-4, there exists a local minimizer $\hat{\boldsymbol{B}}$ of*

$$Q(\boldsymbol{B}) = \frac{1}{2}\Big[\text{vec}(\boldsymbol{Y}^T) - (\boldsymbol{X} \otimes \boldsymbol{\Phi})\text{vec}(\boldsymbol{B}^T)\Big]^T (\boldsymbol{I}_n \otimes \hat{\boldsymbol{\Sigma}})^{-1} \Big[\text{vec}(\boldsymbol{Y}^T) - (\boldsymbol{X} \otimes \boldsymbol{\Phi})\text{vec}(\boldsymbol{B}^T)\Big] + nD \sum_{j=1}^{p} p_{\lambda_n, \gamma}(||\boldsymbol{B}_j||)$$

such that $||\text{vec}(\hat{\boldsymbol{B}}^T) - \text{vec}(\boldsymbol{B}^T)|| = O_p(n^{-1/2})$.

Theorem 2 (Oracle property)**.** *Under assumptions 1-4, the $\sqrt{n}$-consistent local minimizer $\hat{\boldsymbol{B}} = [\hat{\boldsymbol{B}}_+^T, \hat{\boldsymbol{B}}_0^T]^T$ must satisfy*

(1) Sparsity: $\hat{\boldsymbol{B}}_0 = \mathbf{0}$*, with probability tending to 1;*

(2) Asymptotic Normality:

$$\sqrt{n}(\mathrm{vec}(\hat{\boldsymbol{B}}_+^T) - \mathrm{vec}(\boldsymbol{B}_+^T)) \xrightarrow{D} \mathcal{N}\left(\mathbf{0}, \left(\lim_{n\to\infty} \frac{1}{n}(\boldsymbol{X} \otimes \boldsymbol{\Phi})_+^T (\boldsymbol{I}_n \otimes \boldsymbol{\Sigma})^{-1} (\boldsymbol{X} \otimes \boldsymbol{\Phi})_+\right)^{-1}\right).$$

The proof of these theorems is provided in Appendix A. Note that the length of the grid D and the number of basis functions K are considered fixed in the theorems above. Theorems for the scenario in which D and K are unfixed need to be considered thoroughly and derived separately.

2.2.4 Iterative algorithm for models with correlated errors

The method described in Section 2.2.2 uses ordinary least squares to estimate basis coefficients and obtains an estimate $\hat{\boldsymbol{\Sigma}}$ of the covariance $\boldsymbol{\Sigma}$; this estimate is then used to pre-whiten the data prior to the application of variable selection techniques. However, re-estimating the covariance after variable selection may give a refined estimate which can, in turn, be used to pre-whiten the data. This intuition suggests an iterative algorithm:

1. Fit a model using ordinary least squares to obtain an initial estimate $\hat{\boldsymbol{B}}^{(0)}$;
2. Compute residuals and obtain an estimate $\hat{\boldsymbol{\Sigma}}^{(0)}$ of $\boldsymbol{\Sigma}$;
3. For $k > 0$, iterate the following steps until convergence:
   (a) Pre-whiten using the covariance $\hat{\boldsymbol{\Sigma}}^{(k-1)}$;
   (b) Minimize (2.8) to obtain $\hat{\boldsymbol{B}}^{(k)}$;
   (c) Use $\hat{\boldsymbol{B}}^{(k)}$ to construct fitted values and residual curves, and use these to construct $\hat{\boldsymbol{\Sigma}}^{(k)}$.

Various criteria of convergence can be used to monitor convergence of this iterative algorithm; one possible criterion is $\left|\left|\hat{\boldsymbol{B}}^{(k+1)} - \hat{\boldsymbol{B}}^{(k)}\right|\right|^2 < \delta$, which we use in our implementations. This iterative method will be compared to the one-step approach of Section 2.2.2 in simulations.

2.3 Simulation

We conducted simulation studies to examine the properties of the proposed approach. Specifically, we constructed 500 training samples, each consisting of 100 random curves, and 1 test sample containing 1000 random curves. All curves are generated from the model

$$y_i(t) = \sum_{j=1}^{20} x_{ij}\beta_j(t) + \epsilon_i(t)$$

where $x_{ij} \overset{\text{i.i.d.}}{\sim} \mathrm{N}(0, 10)$, $\beta_1(t)$, $\beta_2(t)$, $\beta_3(t)$ are non-zero functions, and the remaining coefficient functions are zero. All functions are observed on a equally spaced grid of length 25. Errors $\epsilon_i(t_d)$ are generated from a multivariate Gaussian distribution with mean zero and covariance $\boldsymbol{\Sigma} = \boldsymbol{G} + \boldsymbol{I}$ where $\boldsymbol{G}$ is the error covariance and $\boldsymbol{I}$ is the identity matrix. the non-zero coefficient functions $\beta_1(t)$, $\beta_2(t)$ and $\beta_3(t)$ are derived from the motivating data in the following way. Focusing on y position curves for reaching motions made to the target at 0 degrees, we estimated motions made by healthy controls, moderately affected stroke patients, and severely affected patients (stroke severity was defined by thresholding the Fugl-Meyer score). These estimated motions were the non-zero coefficients, and are shown in the middle panel of Figure 2.2. The error covariance G was constructed using an FPCA decomposition of residual curves after subtracting the group-specific means.

Four implementations of our proposed method are considered: one-step approaches

as described in Section 2.2.2 using raw and FPCA-based covariance matrix estimates, and iterative approaches as described in Section 2.2.4 using raw and FPCA-based covariance matrix estimates. For the FPCA-based covariance matrix estimate, we used two different values, 0.5 and 0.99, as the cumulative proportion of variance explained (PVE) threshold to determine L. For comparison, we include an approach that pre-whitens using true covariance matrix, as well as ordinary least squares, a variational Bayes method that includes a smoothness penalty (Goldsmith and Kitago, 2015), the FS-LASSO method that uses group LASSO but does not account for residual correlation or biases due to the LASSO penalty, and a group MCP method that assumes uncorrelated error curves, analogously to Wang et al. (2007).

Table 2.1 reports the true positive (TP) and true negative (TN) rates of the estimates of both zero and non-zero coefficient functions. We define functions estimated to be non-zero as "positive" while functions estimated to be zero as "negative". Our iterative approach using a FPCA-based covariance matrix estimate with PVE=0.99 outperforms most competing approaches in terms of correctly identifying the zero functions; its performance is comparable to the approach that uses the true covariance matrix. The approaches using PVE=0.5 perform less well because the estimate of the covariance matrix omits important structure. Our proposed methods substantially outperform the method that assumes uncorrelated errors in accurately identifying zero functions. FS-LASSO has the highest true negative rate but the lowest true positive rate for $\beta_1(t)$, potentially indicating a tendency to over-shrink coefficients to zero. All methods are able to identify $\beta_2(t)$ and $\beta_3(t)$ as non-zero.

Estimates of zero and non-zero coefficient functions obtained using the iterative algorithm with FPCA-based covariance matrix estimate using PVE=0.99 are shown in the left and middle panels of Figure 2.2, respectively. Because their coefficients are relatively large, the estimate of $\beta_2(\cdot)$ and $\beta_3(\cdot)$ are approximately unbiased owing to the structure of the penalty. For $\beta_1(\cdot)$, coefficients are shrunk toward and sometimes

set equal to zero. We show the mean squared error MSE $= E\left(\beta_j(t) - \hat{\beta}_j(t)\right)^2$ and squared bias $E\left(\beta_j(t) - \bar{\hat{\beta}}_j(t)\right)^2$ as functions of t in the right panel of Figure 2.2, where $\bar{\hat{\beta}}_j(t)$ is the average curve across all the simulation datasets. For $\beta_1(\cdot)$, both the MSE and squared bias curves present a sinusoidal shape, which is driven by the sinusoidal shape of the coefficient function itself and by the shrinkage to zero. There is an increasing trend in general as t increases for the MSE of $\beta_2(\cdot)$ and $\beta_3(\cdot)$, which is mostly caused by the increased variability of curves at the end of the distribution as the biases are relatively small. This plot further emphasizes the lack of bias for large coefficients stemming from the use of the group MCP penalty, especially in the case of $\beta_3(\cdot)$.

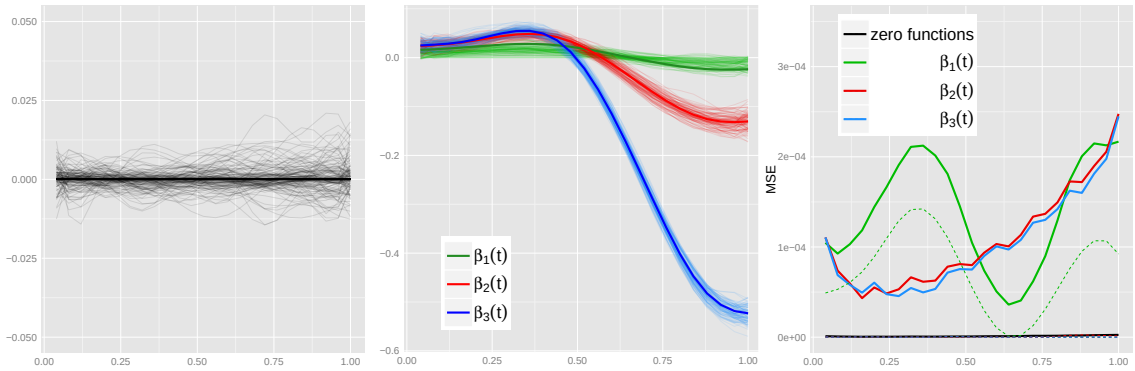

Figure 2.2: Estimates of zero functions (*left*) and non-zero functions (*middle*) obtained using the iterative approach with FPCA-based covariance matrix estimate using PVE=0.99 across all simulated datasets. The true functions are overlaid (bold curves). The right panel shows the both MSE (solid) and squared bias (dashed) as functions of time for all the coefficient functions.

The left and middle columns of Figure 2.3 display the root mean integrated squared error (RMISE), $\sqrt{\int_0^1 \left(\beta_j(t) - \hat{\beta}_j(t)\right)^2 dt}$ for zero and non-zero functions, respectively; in the top row, the FPCA-based covariance estimate is based on PVE=0.99

and in the bottom row based on PVE=0.5. The iterative approach with FPCA-based covariance matrix estimate compares favorably to other approaches, reinforcing the results from Table 2.1. Indeed, the RMISE of our iterative method is comparable to pre-whitening using the true covariance for both zero and non-zero functions. Although FS-LASSO is comparable for zero functions, it has substantially higher RMISE for non-zero functions. Prediction errors on the test sample are shown in the right panel of Figure 2.3. These errors reflect a combination of RMISEs for zero and non-zero functions, and display similar patterns: our proposed methods, in particular when using the FPCA-based estimate of the covariance, have excellent numerical performance. Although there is a slight decline in performance when PVE=0.5, the proposed method still outperforms OLS, FS-LASSO and the method that assumes uncorrelated errors.

Additional simulations that generate uncorrelated errors are presented in detail in Appendix B. In this case, there is no noticeable disadvantage to using our proposed approach, which outperforms competing methods in prediction error.

2.4 Application

We now apply our iterative algorithm using the FPCA-based covariance matrix estimate described in Section 2.2.4 to our motivating dataset. The X and Y position functions are the outcomes of interest, and potential predictors include the Fugl-Meyer score, whether the hand was dominant or non-dominant, whether the hand was contralesional or ipsilesional, target direction (as a categorical predictor) and the interactions between these variables. We analyze the X and Y position functions separately, using the same models and steps.

First, we perform a cross validation analysis to evaluate the algorithm in terms

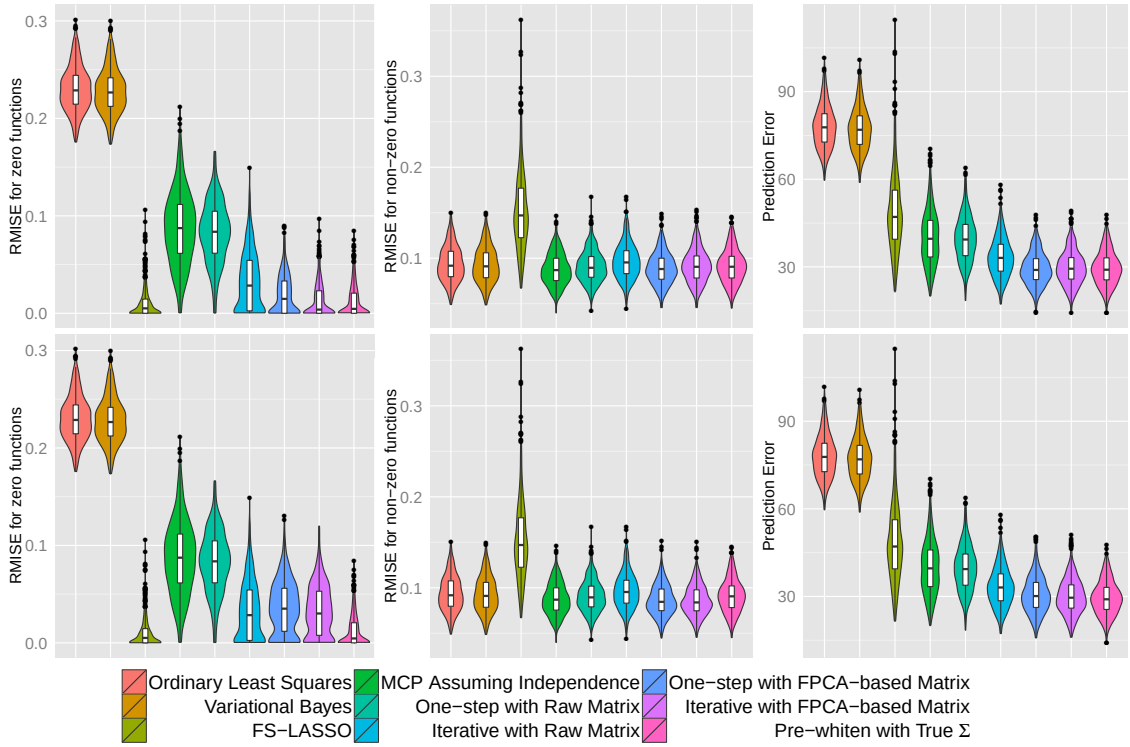

Figure 2.3: The top row shows the comparison among the algorithms when PVE = 0.99 while the second row shows the comparison when PVE = 0.5. The three columns show RMISE for zero functions (*left*) and non-zero functions (*middle*); and prediction error (*right*).

of prediction error on the motivating data. Training and test sets are generated in the following way. For each subject and each hand, we randomly select one motion to each of the eight target directions. These motions are partitioned so that four are in the training set and four are in the test set. Previous work on this dataset (Goldsmith and Kitago, 2015) indicates little or no correlation between motions to different targets made by the same subject, and so our training and test sets are approximately independent even though they contain data from the same subjects. This procedure results in 452 curves in the training set and 452 curves in the test set;

an example is shown in Figure 2.4.

A function-on-scalar regression model is then constructed on the training sample, and prediction errors are obtained for the test sample. Four predictors of interest, the target direction (a categorical variable with eight levels), Fugl-Meyel score (a continuous variable), hand used (dominant/non-dominant) and arm affectedness (contralesional/ipsilesional), are considered in these models. In addition to main effects, all the possible interactions are included to maximize flexibility and scientific interpretation. Thus, the model has 64 coefficient functions to estimate. Rather than the typical design that assigns a reference level for each categorical predictor, a constraint is imposed to the construction of design matrix so that target-specific interpretations are available. This design matrix is equivalent to building the following model for each target:

$$\begin{aligned} y(t) = &\beta_0(t) + \beta_1(t) * \text{Ips.Non.} + \beta_2(t) * \text{Con.Dom.} + \beta_3(t) * \text{Con.Non.} + \beta_4(t) * \text{Fugl-Meyer} \\ &+ \beta_5(t) * \text{Fugl-Meyer} * \text{Ips.Non.} + \beta_6(t) * \text{Fugl-Meyer} * \text{Con.Dom.} \\ &+ \beta_7(t) * \text{Fugl-Meyer} * \text{Con.Non.} + \epsilon(t) \end{aligned} \tag{2.10}$$

where we use the ipsilesional (unaffected) dominant hand of a healthy control as the reference $\beta_0(t)$. Coefficients $\beta_1(t)$, $\beta_2(t)$ and $\beta_3(t)$ compare ipsilesional nondominant, contralesional dominant, and contralesional nondominant to the reference, respectively. The effect of increasing motor impairment in the ipsilesional dominant arm is estimated by $\beta_5(t)$, while differences in the effect of increasing motor impairment comparing other groups to baseline are given by $\beta_6(t)$, $\beta_7(t)$ and $\beta_8(t)$.

The complete procedure described above, consisting of generating training and test sets, fitting the full model to the training set, and producing predictions for the test set, is repeated 100 times. We fit the model using 5, 10, 15 and 20 basis functions, and found that $K = 15$ gave the smallest cross-validated prediction errors. The right panel

of Figure 2.4 presents the prediction errors obtained using our iterative algorithm with FPCA-based covariance matrix estimate; we compare to the variational Bayes approach (without variable selection but with a standard second-derivative penalty). Our iterative algorithm decreases mean prediction error by around 10% (X direction: 163.8 vs. 144.8; Y direction: 143.6 vs. 132.9) compared to the variational Bayes approach. In addition, the iterative algorithm seems to be more stable than the variational Bayes approach as it has fewer outliers and lower median prediction error.

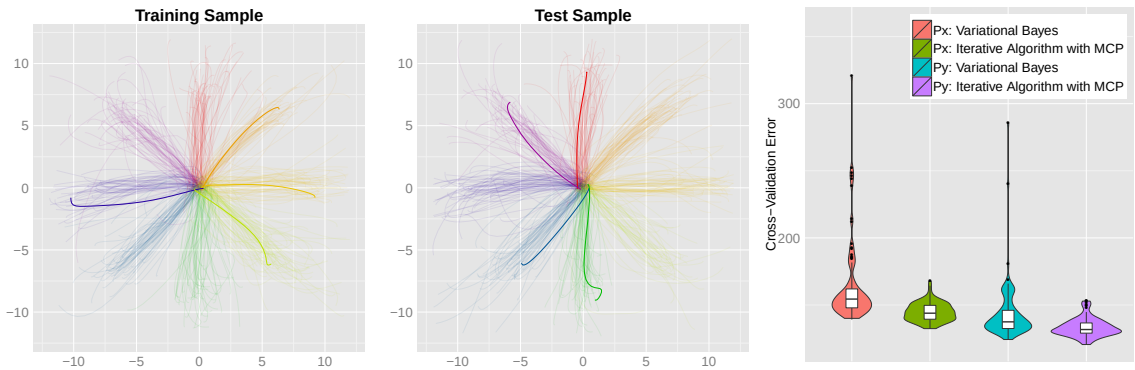

Figure 2.4: One training sample ($left$) and one test sample ($middle$) generated from the planar reaching data. Highlighted curves are from one subject and show how each subject contributes to the training and test sets. Violin plots ($right$) of cross validation errors using the variational Bayes approach and iterative algorithm.

We next conduct our analysis without splitting data into training and test sets. The function-on-scalar regression model is estimated using one motion for each subject and hand to each target with motions drawn randomly for each target and hand. We repeat this analysis 100 times, and Table 2.2 presents the proportion of times selected by the algorithm for each of the 64 coefficient functions. Each row of Table 2.2 corresponds to coefficients $\beta_0(t), \beta_1(t)$, ..., $\beta_7(t)$ in Model (2.10) for a specific target. For instance, the value 0.24 in the third entry of the first row indicates that, in 24 of

100 datasets, there was an estimated difference between contralesional and ipsilesional dominant hands when reaching to the target at 0°.

Large numbers in the table suggest consistent non-zero effects or differences in effect across datasets. Targets at 90° and 270° may have zero effects in the X trajectories, since for those targets the X position is roughly constant over time. The same is true for Targets at 0° and 180° for the Y trajectories. The results in Table 2.2 indicate relatively few differences between ipsilesional and contralesional dominant arms for very mild strokes (Fugl-Meyer = 66), and some differences between the non-dominant arms and the ipsilesional dominant arm. An effect of increasing stroke severity is relatively rarely found for the ipsilensional arms but, as expected, is much more frequently found for the contralesional arms. The conclusions are further reinforced by Figure 2.5, where the predicted motions of subjects with different combinations of Fugl-Meyel Score(66/26), hand used (dominant/non-dominant) and arm affectedness (contralesional/ipsilesional) are presented.

2.5 Discussion

We proposed a model fitting framework that performs variable selection in the context of function-on-scalar regression allowing within-function correlation. This work was motivated by two-dimensional planar reaching data gathered to understand the mechanisms of motor deficit following stroke. We developed an iterative algorithm that alternatively estimates the coefficient functions and covariance structure. Our method relies on a reasonable estimate of the covariance structure, and in our simulations and application we found that an estimation procedure based on FPCA works well. Results from the simulation studies demonstrate the effectiveness of our proposed method in identifying the true zero functions. Indeed, our proposed method has

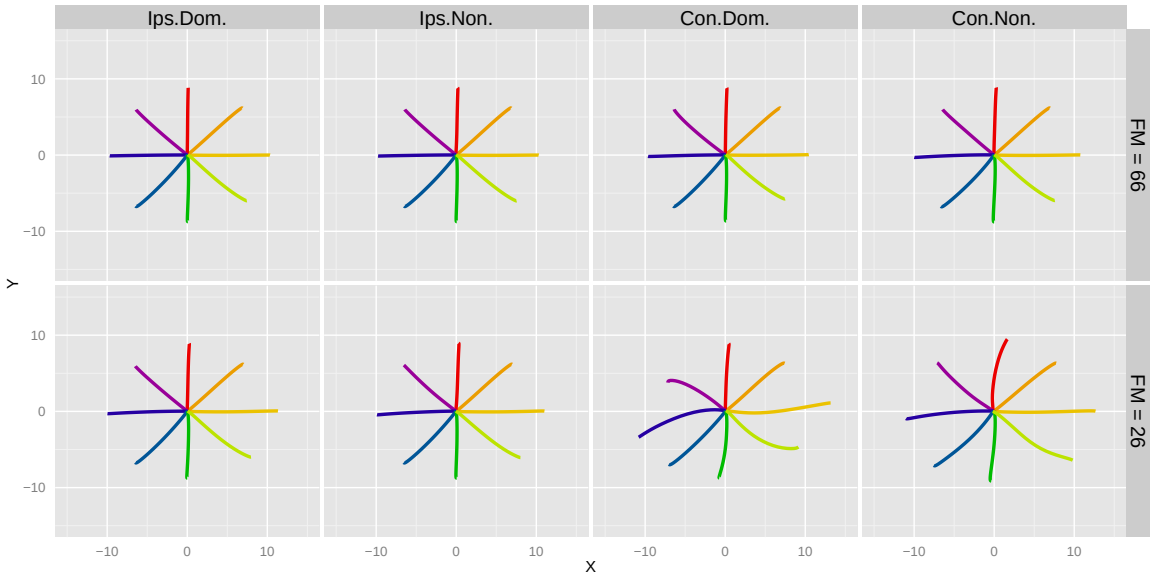

Figure 2.5: Predicted reaching motions for eight subjects with different combinations of Fugl-Meyel Score(66/26), hand used (dominant/non-dominant) and arm affectedness (contralesional/ipsilesional). Motions to different targets are distinguished by colors.

performance comparable to performing variable selection using the true covariance. The application to the motivating data indicates our proposed iterative algorithm makes a significant improvement in terms of decreasing prediction errors and identifying true zero functions.

Future extension of our methodology may take several directions. Quadratic roughness penalties are often applied to enforce smoothness of the coefficient functions in spline-based estimation frameworks. It would be worthwhile to incorporate an explicit roughness penalty in addition to the variable selection penalty to reduce sensitivity to the size of the basis expansion. Motivated by our application (in which repeated motions are made to each target by each subject), the development of methods that account for subject- and target-specific random effects is necessary.

Table 2.1: true positive (TP) and true negative (TN) rates of estimated coefficient functions, where FN is false negative and FP is false positive. They are estimated across all the training samples.

	$\frac{TN}{FP+TN}$	$\frac{TP}{TP+FN}(\beta_1)$	$\frac{TP}{TP+FN}(\beta_2)$	$\frac{TP}{TP+FN}(\beta_3)$
FS-LASSO	0.953	0.850	1.000	1.000
MCP Assuming Independent Errors	0.567	1.000	1.000	1.000
One-step with Raw Matrix	0.370	1.000	1.000	1.000
Iterative with Raw Matrix	0.813	0.962	1.000	1.000
One-step with FPCA-based Matrix (PVE = 0.5)	0.755	0.996	1.000	1.000
Iterative with FPCA-based Matrix (PVE = 0.5)	0.779	0.996	1.000	1.000
One-step with FPCA-based Matrix (PVE = 0.99)	0.863	0.986	1.000	1.000
Iterative with FPCA-based Matrix (PVE = 0.99)	0.915	0.956	1.000	1.000
Pre-whiten with True $\mathbf{\Sigma}$	0.913	0.964	1.000	1.000

Table 2.2: Proportions of 64 coefficient functions being selected, obtained from models with X trajectories (*top*) and Y trajectories (*bottom*).

Target	Fugl-Meyer = 66				Δ Fugl-Meyer = -1			
Direction	Ips.Dom.	Ips.Non.	Con.Dom.	Con.Non.	Ips.Dom.	Ips.Non.	Con.Dom.	Con.Non.
0°	1.00	0.21	0.24	0.35	0.41	0.38	0.58	0.34
45°	1.00	0.20	0.03	0.15	0.16	0.08	0.46	0.37
90°	0.31	0.65	0.31	0.33	0.30	0.17	0.23	0.60
135°	1.00	0.22	0.16	0.24	0.11	0.48	0.64	0.57
180°	1.00	0.35	0.13	0.40	0.18	0.38	0.48	0.35
225°	1.00	0.18	0.16	0.33	0.08	0.22	0.45	0.36
270°	0.83	0.37	0.19	0.34	0.14	0.35	0.52	0.33
315°	1.00	0.41	0.20	0.33	0.28	0.35	0.57	0.67

Target	Fugl-Meyer = 66				Δ Fugl-Meyer = -1			
Direction	Ips.Dom.	Ips.Non.	Con.Dom.	Con.Non.	Ips.Dom.	Ips.Non.	Con.Dom.	Con.Non.
0°	0.34	0.10	0.14	0.11	0.13	0.18	0.66	0.20
45°	1.00	0.02	0.02	0.05	0.01	0.20	0.36	0.12
90°	1.00	0.07	0.16	0.15	0.08	0.21	0.17	0.43
135°	1.00	0.05	0.22	0.29	0.05	0.40	0.79	0.43
180°	0.40	0.13	0.11	0.31	0.15	0.25	0.79	0.27
225°	1.00	0.05	0.03	0.11	0.01	0.12	0.33	0.21
270°	1.00	0.03	0.15	0.14	0.06	0.24	0.13	0.22
315°	1.00	0.03	0.22	0.18	0.06	0.36	0.52	0.54

Chapter 3

Functional Data Analysis of Dynamic PET Data

3.1 Introduction

3.1.1 Background on PET

Positron emission tomography (PET) is a nuclear imaging technique that allows the study of basic mechanisms of the human body. The application of PET imaging in neuroscience has proven to be a valuable tool to better our current understanding of changes during brain stimulation, cognitive activation, and metabolic processes associated with mental illnesses and neurological disorders. One particular application of PET imaging aims to estimate the density of various proteins throughout the brain. For instance, investigators use PET imaging to study the density of β-amyloid plaque that plays a key role in the pathogenesis of Alzheimer's disease (Zeng and Goodman, 2013); another example is the examination of the serotonin (5-HT) neurotransmitter system in the pathophysiology of depression (Miller et al., 2013) and bipolar disorder

(Sullivan et al., 2009), among many others.

The application of PET in such a neuroimaging study begins with the injection of a radiolabeled compound that has affinity for a particular protein in the human brain. This radiolabeled compound, or radiotracer, is designed to bind preferentially to that target protein. Once it is introduced into the bloodstream, the radiotracer is continuously delivered to the brain by the vascular system. Within the brain, each tracer molecule exist in one of three biomedical states: it may be "free" in the synapse, i.e., not bound to any biomolecules; it may be bound "specifically" to the target protein; or it may be "nonspecifically" associated with other macromolecular components. While in the brain, tracer molecules can change from one state to another, potentially making many such transitions during the PET scan. Additionally, because the tracer molecules can cross the blood-brain barriers in both directions, they may also exit the brain and be delivered by the bloodstream to other organs, or back to the brain again.

All radiotracer molecules, no matter their biomedical state, undergo radioactive decay (i.e., emitting positrons) throughout the scan. By detecting the radiation emitted over a given time interval, a three-dimensional image may be obtained via a reconstruction algorithm. Thus, dynamic PET data consist of a sequence of these 3-dimensional images, each voxel of which is a measurement of the concentration of the radiotracer at the corresponding time and location. This concentration depends on the amount of tracer that has been available for delivery in the bloodstream and on the binding behavior of tracer molecules in the brain. Neglecting the noise for the moment, the concentration of the radiotracer may generally be expressed as the convolution between two functions:

$$c(t) = \int_0^t f(s)g(t-s)\,ds \tag{3.1}$$

The function g is the concentration of the radiotracer in the arterial plasma over

time, corrected to account for radioactive metabolites of the tracer; this is termed the "input function" since it represents the amount of tracer available to enter the brain at each time. The function f is the location-specific "impulse response function" (IRF) that represents what the hypothetical concentration of the tracer would be over time if the input function were an instantaneous bolus spike (Dirac delta function). Biologically, if the tracer were to be delivered as an instantaneous bolus spike at time 0, the density of the tracer in the brain would be highest at time 0 and gradually decrease as tracer molecules exit the brain. Therefore, it is expected that the IRF will be non-negative and non-increasing over time.

Because the tracer is designed to bind to the target protein, the IRF is related to the density of that protein in the corresponding location. For instance, in a target-protein-rich region, the IRF decreases slowly because the tracer molecules tend to spend much of the time bound to the target protein. In contrast, in a region with no target proteins, the IRF will decrease at a higher rate. In such a region, because it is completely devoid of the target protein, tracer molecules can only be free or associated with macromolecular components other than the target. If such a region exists, it is termed a "reference region". The binding capacity in the reference region thus represents only "non-specific binding", typically assumed to be uniform throughout the brain; and as a result, specific binding may be estimated based on the difference between IRFs of the region of interest (ROI) and the reference region (Innis et al. (2007); Slifstein and Laruelle (2001)).

For a given voxel, the sequence of concentrations across time is termed the time activity curve (TAC). In practice, the observed TAC is contaminated with noise and may be expressed

$$y(t) = \int_0^t f(s)g(t-s)\,ds + \epsilon(t) \tag{3.2}$$

where f and g are, respectively, the IRF and the input function, and ϵ represents

the errors observed due to radioactive decay, the detection process, processing errors, and other sources of error. For each subject, the input function is common across all voxels, but each voxel or region of interest has its specific IRF. In many PET studies, samples of arterial blood are drawn during the scanning; with each sample, the concentration of the tracer is measured and a metabolite analysis is performed. In this way, the input function g can be measured. Although there is some uncertainty in the measured input function, this is generally small relative to the PET noise so it is typically considered "known" in expressions like (3.2). Hence, any PET modeling technique that involves an input function measured from blood samples must involve deconvolution of the TAC data using this input function to recover the IRF f, which contains information about the density of the target protein.

3.1.2 Overview of our proposed nonparametric modeling approach

Many approaches for dynamic PET modeling have been proposed. The preponderance of these methods have three characteristics in common. First, once the estimated IRF is obtained, it is summarized using a single scalar measure, and subsequently standard univariate analyses, such as t tests or linear mixed models, are performed on the scalar measure. While the scalar summaries have straightforward biological interpretation, by summarizing the entire IRF using a single scalar, it is possible to lose some important features of this function. Second, all these methods focus on estimating the IRF for one subject and one region at a time. This tends to limit the complexity of models that can be fit to PET data. Third, most of the approaches impose strong parametric assumptions on the model of estimating IRF, and some of these assumptions may not hold in real data applications.

In this paper, we propose an alternative analytic approach to explore the kinetics

of the tracer that will improve upon existing techniques in all three of these areas by estimating the IRF nonparametrically using functional data analytic (FDA) techniques. First, in our approach, the entire functions, rather than just the summarized scalars, can be compared across subjects/regions. Comparing the entire estimated IRFs prevents the loss of important information, such as local features of the function. Second, in contrast to the current state of the art, our approach models TACs from multiple subjects simultaneously, which can help capture patterns for subjects with common characteristics. Third, we construct a nonparametric model by applying FDA techniques, which are an important tool set to analyze data that have functional form, such as the TAC in dynamic PET imaging studies. Specifically, the effects of multiple scalar covariates, including continuous covariates, can be incorporated into our model fitting framework. After incorporating the scalar covariates, which we will discuss in detail in Section **??**, Model (3.2) can be treated as a regression problem with functional responses and scalar predictors, i.e., a function-on-scalar regression (Ramsay and Silverman (2005); Reiss et al. (2010); Morris (2015)). Therefore, we can convert this model to a multivariate regression model by extending current FDA techniques.

The advances we propose – to emphasize the IRF as the fundamental unit of interest, rather than a scalar summary; to jointly model IRFs from all subjects; and to model IRFs using a functional data approach in place of the standard parametric model – are a direct response to the shortcomings of existing tools for analyzing dynamic PET data. While our literature review in Section 3.1.3 will identify some related research, the comprehensive analytic framework developed in this manuscript is a major departure from available tools.

3.1.3 Brief overview of current estimation methods for dynamic PET data

Traditional approaches for estimating the IRF impose a parametric form on the IRF, which is motivated by a physiological model for tracer distribution. Among these, compartment modeling is the most widely used method to describe the uptake and clearance of a tracer in the tissue (Slifstein and Laruelle, 2001). The compartments of a system can be defined in our application as biomedical states in which each radiotracer molecule can exist: "free", specifically bound to the target protein, nonspecifically associated with other macromolecular components. For many radiotracers in neuroreceptor mapping, the kinetics of the tracer in the brain can be approximated using a three-tissue compartment model (in which compartments are "free" tracer, tracer specifically bound to the target protein, and tracer nonspecifically associated with other macromolecular components) or a more common two-tissue compartment model (in which "free" tracer and tracer nonspecifically associated with other macromolecular components are considered as comprising a single compartment). Basic assumptions of compartment modeling include that all injected tracer molecules will be in exactly one compartment at any given time and that the rates of transfer between compartments are constant over time. These assumptions ensure that the IRF can be expressed as a sum of exponential functions whose time constants and coefficients are functions of the rate parameters. Rate constants involved in the model can be routinely estimated by solving ordinary differential equations and applying nonlinear regression modeling techniques (Cunningham and Jones (1993); Gunn et al. (2001)).

Although kinetic models are well-established and almost universally applied, it is generally understood that they are inadequate for modeling many radiotracers, and therefore many alternative modeling strategies have been proposed. For example,

"spectral analysis" (Cunningham and Jones, 1993) characterizes the IRF in terms of a set of basis functions and fits the model using non-negative least squares. Gunn et al. (2002) extended this basis function framework by imposing an L_1 penalty. Similarly, Jiang and Ogden (2008) and Lin et al. (2014) proposed a mixture modeling procedure in which each IRF is represented in terms of a smaller number of basis functions. In addition, Logan et al. (1990) introduced "graphical analysis" which is based on the two tissue compartment model. This approach does not estimate the IRF directly but instead estimates a scalar summary. Still, all these approaches have their basis in the standard compartment model.

As with all parametric models, compartment models rely on assumptions about the data generating process and can perform poorly when these are violated in practice. One key assumption is the assumed compartmental structure itself, which is generally understood to be a simplification of a more realistic (but more complex) model. Additionally, the non-linear least squares methods commonly used to fit compartment models tend to have bias that depends on the parameter values, and these tools can also be somewhat numerically unstable (Peng et al., 2008).

The limitations of parametric methods have helped to motivate the development of nonparametric approaches that allow model-free estimation. O'Sullivan et al. (2009) proposed nonparametric "residue analysis" of dynamic PET data based on the indicator dilution theory originally put forth by Meier and Zierler (1954). They base their modeling on Equation (4.1) but place no parametric restrictions on the IRF f. Instead, they express it as $f(t) = C\left(1 - \int_0^t h(s)ds\right)$, where C is a proportionality constant that is interpretable as an overall flow and h is a probability density function. The term $1 - \int_0^t h(s)ds$ is called the *tissue residue function*, and it reflects the fraction of radiotracer that remains in the system at time t. With this formulation, the IRF is constrained to be non-negative and non-increasing over time. In O'Sullivan et al. (2009), the probability density function h is estimated nonparametrically. It is

expressed in terms of a natural cubic B-spline basis, and a weighted second derivative penalty is employed to control the roughness of the estimated curve.

Zanderigo et al. (2015) proposed a nonparametric method that approximates the problem in terms of a discrete deconvolution operation, which can be solved by using a singular value decomposition (SVD). This method is, however, rather sensitive to noise, potentially causing the estimated curve to oscillate considerably, although these effects can be minimized by eliminating diagonal elements below a certain threshold in the diagonal matrix constructed from SVD. Jiang et al. (2015) presented a nonparametric approach for estimating the IRF based on a functional principal component analysis (FPCA). They smoothed the observed PET curves for all voxels using a pre-specified kernel smoother and subsequently applied FPCA on the pre-smoothed curves. Deconvolution was only required on the mean function and the eigenfunctions, rendering it more computationally efficient. The IRF is then recovered using the functions obtained from the deconvolution operator. Regularization is achieved by selecting the number of components using an *ad hoc* measure of goodness-of-fit. Note that each of these methods estimates the IRFs one at a time.

Whether the IRF is estimated using parametric or nonparametric methods, current practice involves summarizing the estimated IRF using a single scalar measure and then comparing this measure across subjects/regions in subsequent analysis. In the parametric approaches discussed above, these summary measures are related to some aspect of the density of the target protein. For instance, the total volume of distribution of tracer (V_T) is defined as a functional of the IRF; specifically, $V_T = \int_0^\infty f(t)\,dt$. Even with the nonparametric methods for estimating the IRF, interest generally lies in computing some scalar measure that can then be compared across subjects/regions. One option is to calculate $\int_0^{t_{end}} f(t)\,dt$, the area under the IRF until the end of the scan. This is a nonparametric analogue to V_T, although it does not have the same clear biological interpretation. Another option is to calibrate

the nonparametric estimator with a specific compartment model so that the resulting summary measure will have the same interpretation in the case that the parametric model holds (O'Sullivan et al. (2009); Zanderigo et al. (2015)).

In contrast to the preceding, we develop a flexible non-parametric approach to dynamic PET data that 1) focuses on the IRFs, rather than a single summary measure of these functions, as the basis for comparisons; 2) models data from all subjects simultaneously; and 3) estimates IRF for each subject and includes covariate effects using FDA techniques. The first two of these are novel contributions to the PET literature, and final point required new developments in functional data analysis.

The rest of the article is organized as follows. In Section 3.2, we describe a functional approach to nonparametrically estimate the IRFs of all subjects simultaneously and to compare the entire estimated IRFs across subjects. We conduct a simulation study and present the results in Section 3.3. An application of our method on some clinical PET data is given in Section 3.4. Finally, we present some concluding remarks in Section 3.5.

3.2 Methodology

3.2.1 Conceptual model

In this article, we consider only the situation in which the input function is observed for each subject. Following the description given in Section 3.1, our model for subject i is

$$y_i(t) = \int_0^t f_i(s) g_i(t-s)\, ds + \epsilon_i(t), \qquad i = 1, \ldots, n \tag{3.3}$$

In principle, both y and f should be indexed by subject and region because the analysis can be performed on any voxel/region. For simplicity, we restrict our attention to a single region and suppress the related index.

For the purpose of estimating f_i in Model (3.3), we assume f_i can be separated into a population-level fixed effect that may depend on some measured covariates and a subject-level random effect. This formulation allows for the effects of multiple scalar covariates, including continuous covariates, on the IRF to be directly estimated from the model. Then

$$\begin{aligned} f_i(s) &= \beta_0(s) + \sum_{j=1}^{p} x_{ij}\beta_j(s) + \delta_i(s) \\ &= \boldsymbol{x}_i^T \boldsymbol{\beta}(s) + \delta_i(s),\ s \in \mathcal{T} \end{aligned} \tag{3.4}$$

where $\boldsymbol{x}_i = (1, x_{i1}, \cdots, x_{ip})^T$ is the vector of covariate values for subject i and $\boldsymbol{\beta}(s) = (\beta_0(s), \beta_1(s), \cdots, \beta_p(s))^T$. Next, by expanding the functions $\beta_0(\cdot)$, $\beta_1(\cdot)$, $\cdots$, $\beta_p(\cdot)$, and $\delta_i(\cdot)$ in terms of a pre-specified set of basis functions $\phi_1, \cdots, \phi_k$, Model (3.4) becomes

$$\begin{aligned} f_i(s) &= \boldsymbol{x}_i^T \begin{pmatrix} \sum_{k=1}^{K} \beta_{0k}\phi_k(s) \\ \vdots \\ \sum_{k=1}^{K} \beta_{pk}\phi_k(s) \end{pmatrix} + \sum_{k=1}^{K} \delta_{ik}\phi_k(s) \\ &= \boldsymbol{x}_i^T \boldsymbol{B}\boldsymbol{\phi}(s) + \boldsymbol{\delta}_i^T \boldsymbol{\phi}(s) \\ &= \left(\boldsymbol{x}_i^T \boldsymbol{B} + \boldsymbol{\delta}_i^T\right) \boldsymbol{\phi}(s) \end{aligned} \tag{3.5}$$

where $\boldsymbol{\phi}(s) = (\phi_1(s), \cdots, \phi_K(s))^T$; $\boldsymbol{\delta}_i = (\delta_{i1}, \cdots, \delta_{iK})^T$ are the basis function coefficients of $\delta_i(\cdot)$; and $\boldsymbol{B} = (\boldsymbol{\beta}_0, \boldsymbol{\beta}_1, \cdots, \boldsymbol{\beta}_p)^T$, where $\boldsymbol{\beta}_j = (\beta_{j1}, \cdots, \beta_{jK})^T$, $j = 0, \cdots, p$ are the basis function coefficients corresponding to $\beta_0(\cdot)$, $\beta_1(\cdot)$, $\cdots$ and $\beta_p(\cdot)$, respectively. We then replace $f_i(s)$ in Model (3.3) with the expression in (3.5), which gives

$$y_i(t) = \int_0^t \left(\boldsymbol{x}_i^T \boldsymbol{B} + \boldsymbol{\delta}_i^T\right) \boldsymbol{\phi}(s) g_i(t-s)\, ds + \epsilon_i(t). \tag{3.6}$$

3.2.2 Model for the observed PET data

In practice, the measured concentration values are derived from the decay counts observed over a given time interval across a grid of time points $\{t_{i\ell}\}_{\ell=1,\cdots,L_i}$. By design, the time frames are gradually longer over time during the scan, because of the radioactive decay process and because of decreasing concentration of the tracer. The change in time frames over the scan has two practical implications. First, the discrete grid of time points on which TACs are observed is taken to be the midpoints of the frames, which can be irregular. Second, since data are observed over consecutive time frames of different lengths, the frame duration, the radioactive decay and the overall concentration affect the variability of the response. As a result, weighting schemes that account for these factors are necessary. In the simulation and real data analyses below, weights are set to be the duration of the time-frame corresponding to $t_{i\ell}$ (Zanderigo et al., 2015), although our methodology allows other weighting schemes to be used. Errors are assumed to be uncorrelated over time as they arise originally from decay count data, which are naturally independent, and then are reconstructed and registered separately for each time interval.

With these considerations, the observed data within Model (3.6) can be expressed in matrix form

$$\begin{aligned}\boldsymbol{y}_i &= \boldsymbol{Z}_i\left(\boldsymbol{B}^T\boldsymbol{x}_i + \boldsymbol{\delta}_i\right) + \boldsymbol{\epsilon}_i \\ &= \boldsymbol{Z}_i\boldsymbol{B}^T\boldsymbol{x}_i + \boldsymbol{Z}_i\boldsymbol{\delta}_i + \boldsymbol{\epsilon}_i \\ &= \boldsymbol{Z}_i\left(\boldsymbol{x}_i^T \otimes \mathbf{I}_K\right)\boldsymbol{\beta} + \boldsymbol{Z}_i\boldsymbol{\delta}_i + \boldsymbol{\epsilon}_i \end{aligned} \tag{3.7}$$

where $\boldsymbol{y}_i = (y_i(t_{i1}), \cdots, y_i(t_{iL_i}))^T$; $\boldsymbol{Z}_i$ is a $L_i \times K$ matrix with the (ℓ, k)th element $\int_0^{t_{i\ell}} \phi_k(s) g_i(t_{i\ell} - s)\, ds$; $\mathbf{I}_n$ denotes the $n \times n$ identity matrix; $\boldsymbol{\beta} = \text{vec}(\boldsymbol{B}^T)$ is the vector obtained by stacking the rows of $\boldsymbol{B}$, i.e., $\boldsymbol{\beta} = (\boldsymbol{\beta}_0^T, \boldsymbol{\beta}_1^T, \cdots, \boldsymbol{\beta}_p^T)^T$; and $\boldsymbol{\epsilon}_i = (\epsilon_i(t_{i1}), \cdots, \epsilon_i(t_{iL_i}))^T \sim \mathcal{N}(\mathbf{0}, \sigma^2\boldsymbol{W}_i^{-1})$. $\boldsymbol{W}_i$ is a diagonal matrix with diagonal

elements $\{w_{i\ell}\}_{l=1,\cdots,L_i}$, where $\{w_{i\ell}\}$ are as fixed and known observation weights.

Based on our construction, Model (3.7) can be viewed as a linear mixed effects model where $\boldsymbol{\beta}$ are the population-level fixed effects and $\boldsymbol{\delta}_i \sim \mathcal{N}(\mathbf{0}, \sigma_\delta^2 \mathbf{I}_K)$ are the subject-level random effects. Let $\oplus_{i=1}^n \boldsymbol{D}_i$ denote a block diagonal matrix with diagonal matrix elements $\{\boldsymbol{D}_1, \cdots, \boldsymbol{D}_n\}$, i.e.,

$$\oplus_{i=1}^n \boldsymbol{D}_i = \mathrm{diag}(\boldsymbol{D}_1, \cdots, \boldsymbol{D}_n) = \begin{pmatrix} \boldsymbol{D}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{D}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \boldsymbol{D}_n \end{pmatrix}$$

and let $N = \sum_{i=1}^n L_i$. By combining the equations for all subjects in Model (3.7), we now have

$$\begin{aligned} \boldsymbol{y} &= \boldsymbol{Z}\left(\boldsymbol{X} \otimes \mathbf{I}_K\right)\boldsymbol{\beta} + \boldsymbol{Z}\boldsymbol{\delta} + \boldsymbol{\epsilon} \\ &= \boldsymbol{U}\boldsymbol{\beta} + \boldsymbol{Z}\boldsymbol{\delta} + \boldsymbol{\epsilon}, \end{aligned} \tag{3.8}$$

where $\boldsymbol{y} = (\boldsymbol{y}_1^T, \cdots, \boldsymbol{y}_n^T)^T$, $\boldsymbol{Z} = \oplus_{i=1}^n \boldsymbol{Z}_i$, $\boldsymbol{X} = (\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n)^T$, $\boldsymbol{\delta} = (\boldsymbol{\delta}_1^T, \cdots, \boldsymbol{\delta}_n^T)^T$, $\boldsymbol{\epsilon} = (\boldsymbol{\epsilon}_1^T, \cdots, \boldsymbol{\epsilon}_n^T)^T$ and $\boldsymbol{U} = \boldsymbol{Z}\left(\boldsymbol{X} \otimes \mathbf{I}_K\right)$.

3.2.3 Constrained estimation

As discussed in Section 3.1, the IRF is non-negative and non-increasing over time. Formally,

1. $f_i(s) \geq 0$, $i = 1, \ldots, n$ (non-negativity);

2. $f_i(s) - f_i(t) \geq 0$, $\forall\, s \leq t$, $i = 1, \ldots, n$ (monotonicity).

In addition, we incorporate a roughness penalty on the estimated curves to prevent overfitting.

Based on the likelihood function of Model (3.8) and by incorporating the preceding constraints and penalties, $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$ can be estimated by minimizing

$$\frac{1}{2}(\boldsymbol{y}-\boldsymbol{U}\boldsymbol{\beta}-\boldsymbol{Z}\boldsymbol{\delta})^T\boldsymbol{W}(\boldsymbol{y}-\boldsymbol{U}\boldsymbol{\beta}-\boldsymbol{Z}\boldsymbol{\delta})+\lambda_1\boldsymbol{\beta}^T\boldsymbol{P}_\beta^T\boldsymbol{P}_\beta\boldsymbol{\beta}+\lambda_1\boldsymbol{\delta}^T\boldsymbol{P}_\delta^T\boldsymbol{P}_\delta\boldsymbol{\delta}+\lambda_2\boldsymbol{\delta}^T\boldsymbol{\delta}, \quad (3.9)$$

subject to $\boldsymbol{C}(\boldsymbol{X}\otimes\mathbf{I}_K)\boldsymbol{\beta}+\boldsymbol{C}\boldsymbol{\delta}\geq\mathbf{0}$, where $\boldsymbol{W}=\oplus_{i=1}^n\boldsymbol{W}_i$, $\boldsymbol{P}_\beta=\mathbf{I}_{p+1}\otimes\boldsymbol{\Delta}_2\boldsymbol{\Phi}$ and $\boldsymbol{P}_\delta=\mathbf{I}_n\otimes\boldsymbol{\Delta}_2\boldsymbol{\Phi}$ and

$$\boldsymbol{C}=\left(\frac{\mathbf{I}_n\otimes\boldsymbol{\Phi}}{\mathbf{I}_n\otimes\boldsymbol{\Delta}_1\boldsymbol{\Phi}}\right).$$

Throughout, $\otimes$ represents the Kronecker product of two matrices; $\boldsymbol{\Phi}$ is the matrix consisting of basis functions evaluated at a pre-specified dense grid $\{\tau_d\}_{d=1,\cdots,D}$, which is equally spaced and lies in the range determined by the irregularly spaced and subject-specific time points at which the data are observed (i.e., the (d,k)th entry of $\boldsymbol{\Phi}$ is $\phi_k(\tau_d)$ for $d=1,\ldots,D$ and $k=1,\ldots,K$; and $\boldsymbol{\Delta}_1$ and $\boldsymbol{\Delta}_\mathbf{2}$ are the first and second order difference matrices, respectively.

In the loss function (3.9), the non-negativity and monotonicity constraints are implemented by the inequality $\boldsymbol{C}(\boldsymbol{X}\otimes\mathbf{I}_K)\boldsymbol{\beta}+\boldsymbol{C}\boldsymbol{\delta}\geq\mathbf{0}$, and the upper and lower blocks of $\boldsymbol{C}$ correspond to these two constraints, respectively. The terms $\lambda_1\boldsymbol{\beta}^T\boldsymbol{P}_\beta^T\boldsymbol{P}_\beta\boldsymbol{\beta}$ and $\lambda_1\boldsymbol{\delta}^T\boldsymbol{P}_\delta^T\boldsymbol{P}_\delta\boldsymbol{\delta}$ control the smoothness of fixed and random effects, respectively, where $\boldsymbol{P}_\beta=\mathbf{I}_{p+1}\otimes\boldsymbol{\Delta}_2\boldsymbol{\Phi}$ and $\boldsymbol{P}_\delta=\mathbf{I}_n\otimes\boldsymbol{\Delta}_2\boldsymbol{\Phi}$. The magnitudes of both terms are controlled by the same tuning parameter λ_1 since we expect the smoothness of both effects to be similar; separate tuning parameters could also be used, but doing so would be more computationally intensive when choosing the values of the parameters. The term $\lambda_2\boldsymbol{\delta}^T\boldsymbol{\delta}$ in the loss function (3.9) controls the magnitude of the variance of the individual-specific random effects and implicitly guarantees that the model is identifiable; this term also reflects the random effects specification in our mixed model representation.

The algorithm that minimizes the loss function (3.9) is implemented in the `pcls`

function in the `mgcv` package (Wood, 2011), which solves least squares problems with quadratic penalties subject to linear equality and/or inequality constraints using quadratic programming.

3.2.4 Tuning parameter selection

The values of the tuning parameters λ_1 and λ_2 may be chosen by cross validation through a process we describe below. Another parameter to be determined is K, the number of basis functions used in the expansion of the IRF, which could also be determined by cross validation. Provided that the basis set is rich enough to capture all the details of the functions to be estimated, the choice of K is not crucial (Ruppert, 2002). However, the exact choice of K may vary from application to application and some examination of this choice is necessary.

To choose λ_1 and λ_2 we use a bivariate grid search. In each iteration, we generate the test sample by randomly selecting two points from the observed TAC for each subject and treat the unselected data as the training sample. While it is common in some functional data applications to leave out the entire curves (i.e., performing a leave-one-out cross validation at the subject level), our model contains an unobserved subject-level random effect, and therefore randomly leaving out two points from each curve can help assess the performance of the subject-specific effects. For our data, this procedure amounts to leaving out roughly 10% of the observations in each training-test split. The main purpose of performing this "regression-style" cross validation is to strike a good bias/variance tradeoff at the subject and population level. For each split, a full model is fit to the training sample, and prediction error is obtained by applying the fitted model to the test data.

3.3 Simulation

In this section we conduct a simulation study to assess the quality of our proposed nonparametric method. Throughout, simulated datasets are generated based on the motivating DASB data. We also perform a systematic comparison to an existing parametric method; Frankle et al. (2006) and Ogden et al. (2007) concluded that the one-tissue compartment model has good performance on DASB binding by evaluated common competing methods using test-retest data, and therefore we use this as a comparison approach.

The two methods are compared in two scenarios: first when the data follow the one-tissue compartment parametric model, and second when they do not. We generate realistic datasets under both scenarios by first fitting both our proposed method and the one-tissue compartment model to the motivating data. Estimated IRFs are computed for both methods and the integrated squared difference is used to assess agreement between methods. To generate datasets under the parametric model, we identify the 20 subjects for whom the methods have the best overall agreement. From these, we randomly choose two parametric estimates and take a weighted average to obtain the new subject's "true" IRF; weights for the weighted average are α and $1 - \alpha$, where where α is sampled from a Uniform$[0, b]$ distribution with b chosen to ensure the resulting data has the same mean and variance as the original data. The simulated TAC is produced using the "true" IRF and an input function randomly sampled from the observed data according to Model 3.2, with errors generated from a mean-zero Gaussian distribution with variance equal to the error variance in our real data analysis. A similar process is used to generate data that does not follow the parametric model, except that the two subjects are selected from among those with the least agreement between methods and the "true" IRF is simulated by taking the

weighted average of their nonparametric estimates. The simulated dataset in each scenario consists of 100 subjects.

We fit both our proposed model and the one-tissue compartment model to each simulated dataset. The values of tuning parameters are determined by a bivariate cross-validation as described in Section 3.2.4, using a 10×10 grid of possible tuning parameter value and evaluating each combination on 100 training-test splits.

The root integrated mean square errors of the estimated IRFs obtained from both scenarios are presented in Figure 3.1. Both methods perform well under the scenario in which the IRFs used to generate data come from the parametric model, although there is a small but expected decrease in performance for the non-parametric approach stemming from the increase in model complexity. For the scenario in which the truth comes from a non-parametric model, however, the proposed approach substantially outperforms the parametric method. Indeed, the performance of the proposed approach is broadly similar across data generating mechanisms, while the performance of the parametric approach suffers when the assumed model is not true.

3.4 PET data analysis

Impaired serotonergic function has been implicated in the pathophysiology of major depressive disorder (MDD) and bipolar disorder (BPD). Both have been associated with suicidal behavior and completed suicide. In these studies, the tracer [^{11}C]DASB is frequently used to examine the binding capacity of the serotonin transporter in the serotonin (5-HT) neurotransmitter system in the human brain. The one-tissue model generally provides reasonable fit to DASB data and also results in reproducible estimates of binding parameters (Ogden et al., 2007). This model involves a single brain compartment which exchanges tracer molecules with the blood compartment, with tracer particles crossing the blood brain barrier (BBB) into the brain at constant

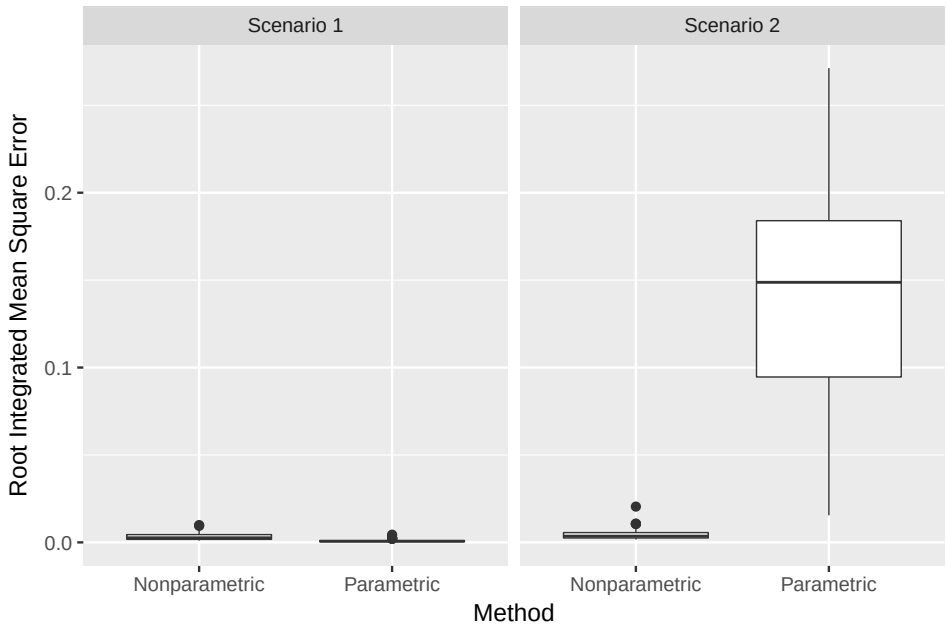

Figure 3.1: The root integrated mean square errors of the estimated IRFs obtained from both our proposed method and the one-tissue compartment model when data follow parametric model (left) and when they do not (right).

rate K_1 and flowing in the other direction with rate k_2. For this model, the total volume of distribution can be shown to be $V_T = K_1/k_2$ (Innis et al., 2007). Binding potential, a measure related to the density of the target protein in the brain, is typically calculated indirectly with this model, by comparing total distribution volume in a region of interest with that in a reference region.

Our data consist of PET scans using the [^{11}C]DASB tracer of 137 subjects belonging to three diagnostic groups: BPD (20), MDD (83), and normal control (34). Details of the data acquisition are given in Miller et al. (2013) and Miller et al. (2016). To summarize, injected dose averaged approximately 16mCi for each of the groups with a standard deviation of approximately 2 for all groups. Average injected mass ranged from 4μg to 5μg for the groups with standard deviation approximately 2. PET

scanning was done on an Siemens ECAT HR+, and reconstruction was done using the filtered back projection algorithm. Any subject head motion during scanning was corrected by applying the FMRIB linear image registration tool (FLIRT). Regions of interest were identified on a T1-weighted MRI for each subject and transferred to the PET imaging space by coregistering to the subjects' corresponding sequence of PET images. The regions that we consider in this analysis are relatively large and easy to identify. Arterial samples were drawn every 10 seconds for the first 2 minutes, every 20 seconds for the next two minutes, and then at time points 6, 8, 12, 16, 20, 30, 40, 50, 60, 80, 90, 100, and 120 minutes. Blood samples drawn at 2, 12, 20, 50, 80, and 100 minutes were also analyzed to determine unmetabolized parent compound levels and the arterial data were corrected accordingly (Parsey et al., 2006b). The model used to describe the arterial concentration data is linear to the peak and a sum of three exponentials after the peak, and the fitted model was used at the input function for each subject.

In this section, we apply our method for estimating IRFs while accounting for covariate effects and constraints, described in Section 3.2, to this dataset. We first conduct an analysis on the midbrain, an ROI whose importance in the development of depression has been previously demonstrated. Serotonin transporter availability in the midbrain has been shown to be different in depressed subjects by Parsey et al. (2006a) and Malison et al. (1998) and has been studied in other PET studies of depression (Miller et al. (2013); Sullivan et al. (2009)). Subsequently, we conduct another analysis that focuses more closely on the binding specific to the target receptor. Raw TACs of the ROI and reference region as well as the input functions for all subjects are shown in Figure 3.2.

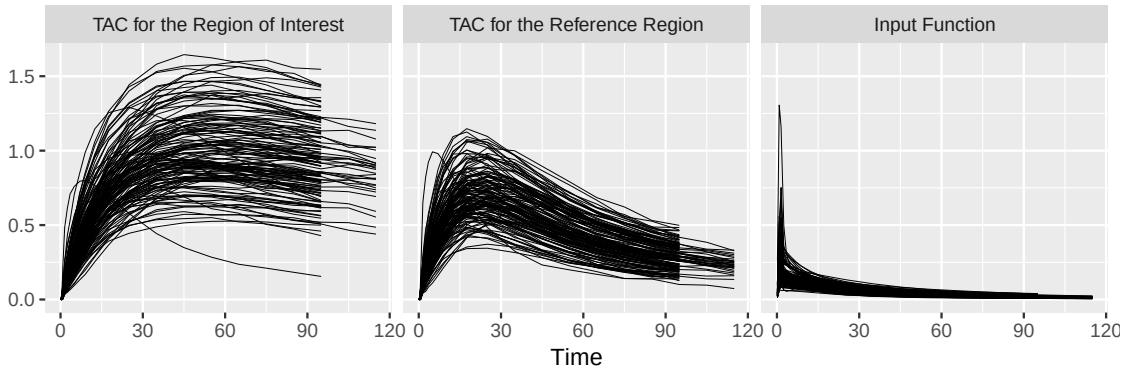

Figure 3.2: Raw TACs of the ROI (left) and reference region (middle) and the input functions (right) for all subjects.

3.4.1 Analysis of the midbrain data

Focusing first on the midbrain, a model with diagnosis group as the only predictor is fit to the entire dataset. A cubic B-spline basis with 10 basis functions is used to model IRFs; no appreciable difference is observed when we repeat the analysis with either $K = 5$ or 15. Values of the tuning parameters λ_1 and λ_2 are determined by cross validation, as discussed in Section 3.2.4, using a 20×20 grid of possible tuning parameter value and evaluating each combination on 100 training-test splits. Figure 3.3 provides the cross validated prediction error over all values of the tuning parameters.

After choosing λ_1 and λ_2, we estimate model parameters using the complete dataset. The fitted group mean and individual IRFs are shown in the left panel of Figure 3.4. These results indicate that the mean IRF of the patients with bipolar disorder tends to be lower than that of the patients with major depression throughout the entire study period. In addition, the mean IRF of the healthy controls is lower than the other two groups at the beginning of the scan, but decreases at a slower rate and is higher than the other groups for most of the scan duration. When the

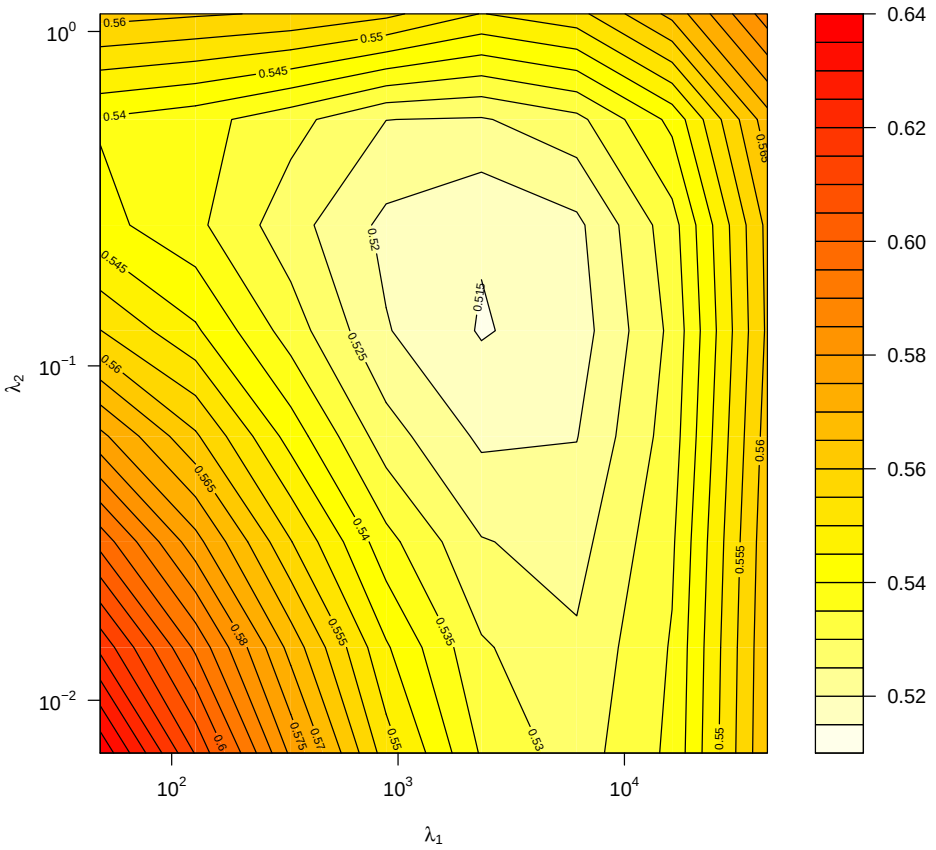

Figure 3.3: A heat map of cross validated errors, used to determine the tuning parameters λ_1 and λ_2, with contours overlaid.

IRF is estimated according to assumptions required by kinetic models, any difference between two IRFs can only be attributed to differences in the set of rate parameters (which combine to determine the density of target proteins), and that is the extent of the interpretation of such functions. By estimating the IRF nonparametrically, however, a much more flexible interpretation is possible. The IRF reflects the density of the target proteins, to be sure, but going well beyond that it also represents the

rate at which the tracer molecules leave the system, which may be time-varying.

We next construct pointwise confidence bands of the group mean differences using a bootstrap algorithm in which 1000 bootstrap samples are generated in the following way. Within each diagnosis group, subjects are chosen with replacement with the sample size of the bootstrap sample of each group set to be the same as the original sample size. Then apply our proposed approach on each bootstrap sample to estimate the group mean curves and take the differences between healthy controls and the other two groups.

The right panel of Figure 3.4 shows the 95% pointwise bootstrap confidence bands based on the 1000 bootstrap samples. Due to the relatively large sample size of the major depression group, the confidence band for the difference between major depression patients and healthy controls is narrower than that for the difference between bipolar disorder patients and controls. Figure 3.4 suggests that the mean IRF of major depression patients is lower than the controls between 75 and 90 minutes and the mean IRF of the bipolar disorder patients is lower than the controls between 75 and 105 minutes.

To provide a frame of reference, we also model these data using the one-tissue compartment model, which has been deemed a reasonable compartment model structure for this tracer (Frankle et al. (2006); Ogden et al. (2007)). IRFs for four selected subjects, estimated using our approach and the compartment model, are shown in the top row of Figure 3.5. For the first subject the estimates are similar, but differences for the remaining subjects can be clearly observed. This suggests that parametric models may be appropriate for some subjects but not others. The bottom row of Figure 3.5 shows the observed data and the estimated TACs for the same four subjects using both methods. Visual inspection of these panels suggest that our method can substantially outperform the parametric approach in terms of fitted values. Lastly, Figure 3.6 shows the residuals obtained using both our nonparametric approach and

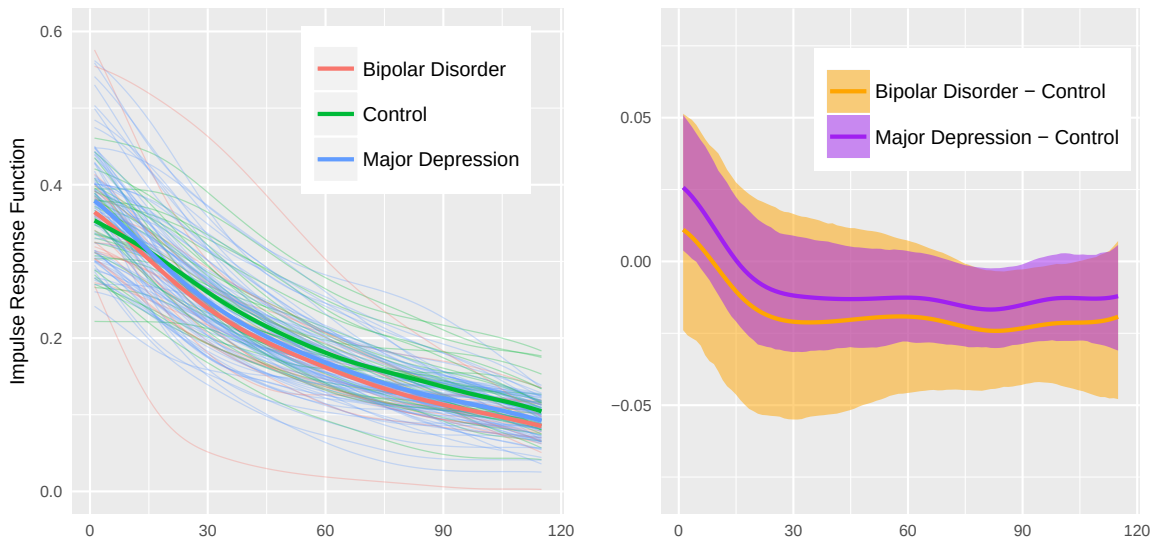

Figure 3.4: The left panel shows the estimated IRFs obtained using our method. Group mean (thick) and individual curves (thin) of different groups are presented in different colors. The right panel shows the 95% pointwise bootstrap confidence intervals (shaded areas) with the group mean differences (solid) estimated from the original sample overlaid.

the parametric method. The residuals obtained using our approach have mean zero and roughy constant variance, while the residuals obtained through parametric modeling appear to miss trends in the data and have larger variability.

3.4.2 Analysis of the difference between the midbrain and the reference region

Next, we focus on isolating the binding capacity that is specific to the target protein. As mentioned in Section 3.1, tracer molecules may bind to their target protein ("specifically bound") or they may be associated with other macromolecular components ("non-specifically associated"). However, the observed PET data can measure

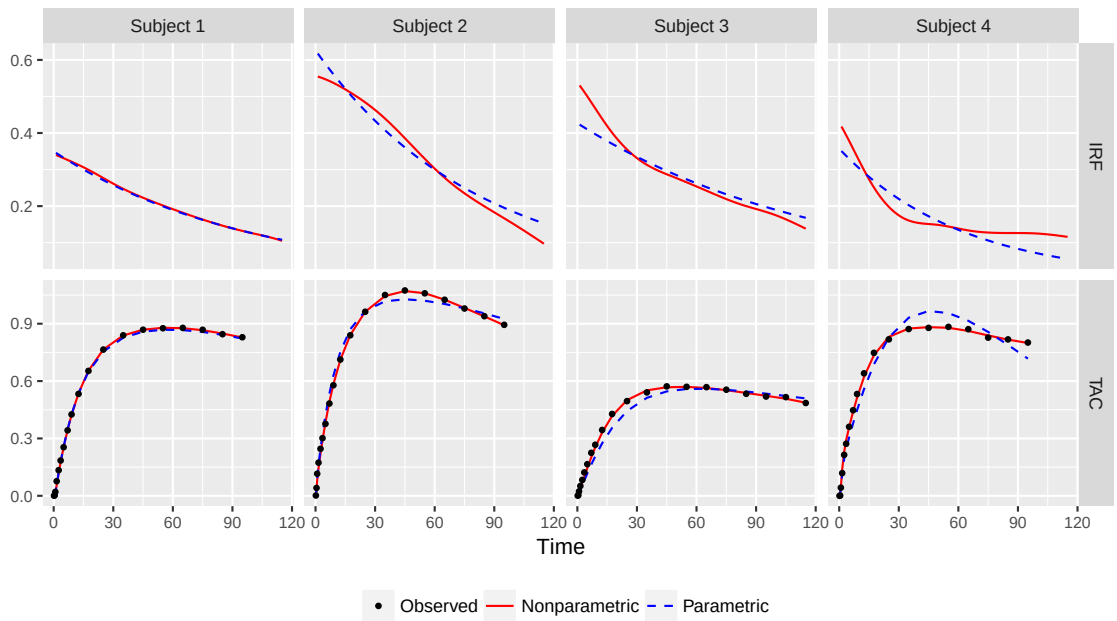

Figure 3.5: The top row shows comparisons between estimated IRFs using our proposed approach and the one-tissue compartment model for four selected subjects. The bottom row shows comparisons between estimated TACs using our proposed approach and the one-tissue compartment model for the same subjects with observed curves overlaid.

only total concentration, consisting of unbound molecules as well as those bound to either type of protein, and is unable to discriminate among those states. Thus, as mentioned in Section 3.1.1, it is common in practice to designate a region that is devoid of the target protein as a "reference region". If non-specific association is uniform across the brain (as is always assumed), a reference region, which will allow only non-specific association, will allow better focus on the binding of the tracer to the specific target protein.

As we discussed in Section 3.1.3, the total volume of distribution of the tracer (V_T) is a commonly used summary measure for the parametric approaches. It is

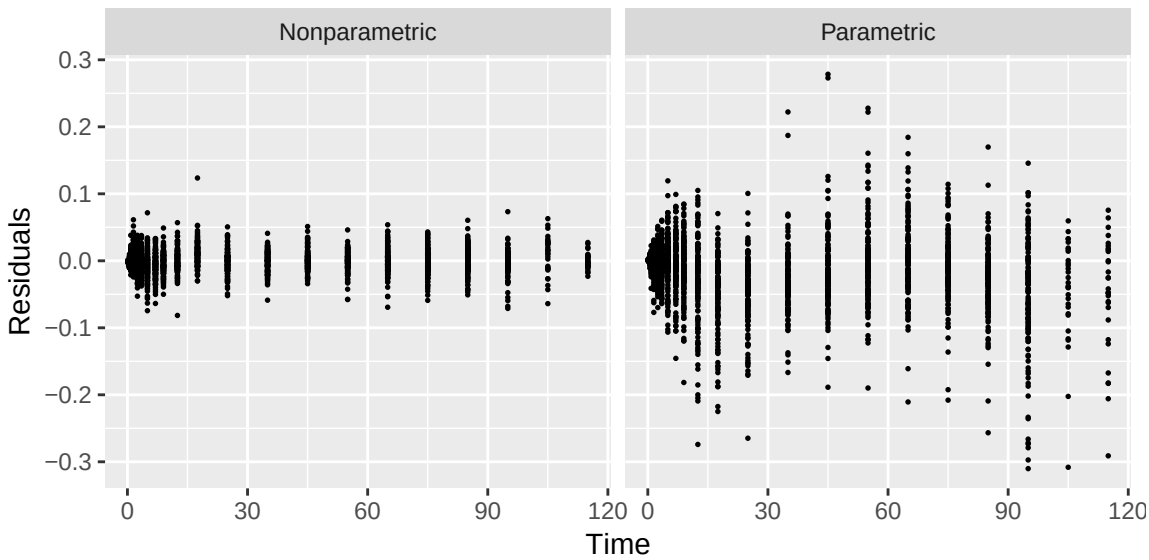

Figure 3.6: Residuals obtained using our proposed approach (left) and residuals obtained using the 1TC model (right).

made up of two components: one that is only involved with specific binding to the target; and the other includes everything else, including volume of unbound tracer and tracer that is associated with other macromolecular components. The V_T of the region of interest represents the total volume of the two components while the V_T of a reference region consists only the second component. Thus, a standard measure of binding is $V_{T\_\text{region}} - V_{T\_\text{ref}} = \int_0^\infty f_{\text{region}}(t)dt - \int_0^\infty f_{\text{ref}}(t)dt = \int_0^\infty \left(f_{\text{region}} - f_{\text{ref}}\right)(t)dt$, i.e., that the binding measure is based on a functional of the difference between two IRFs. Although V_T refers only to parametric analysis, it is still reasonable to take the difference between the IRFs obtained from nonparametric approaches because the difference pertains only to specific binding component of the IRF.

In this analysis, we designate the midbrain as the region of interest and the cerebellar gray matter as the reference region. We modeled IRFs in both the region of interest and the reference region using our approach described in Section 3.2, with

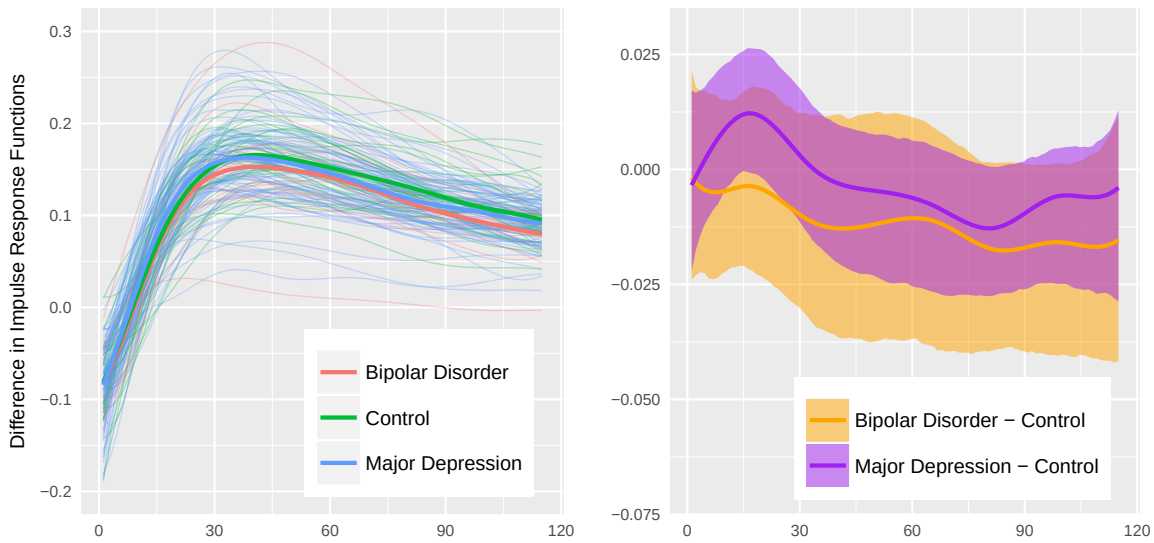

Figure 3.7: The left panel shows the estimated differences in IRFs, including both group mean and individual curves, obtained using our approach. The right panel shows the 95% pointwise bootstrap confidence bands (shaded areas) with the group mean differences (solid) estimated from the original sample overlaid.

values for the tuning parameters determined separately. The differences between regions within each subject are estimated by subtracting the IRFs of the reference region from those of the region of interest.

The left panel of Figure 3.7 displays the estimated difference between the IRFs of the region of interest and reference region using this approach. Results indicate that the difference curve starts negative and, as time goes on, reaches a peak and then decreases. The negative difference at the beginning may be due to faster initial uptake in the reference region than in the midbrain region; this rate is unrelated to receptor availability. As time goes on, precisely because of the specific binding of the target protein in the region of interest, tracer molecules in the reference region would tend to be washed out relatively early compared to those in the midbrain. Comparisons

across groups indicate that the mean difference IRF of the healthy controls is higher than the other two groups after 40 minutes and that the mean difference IRFs of the patients with bipolar disorder and the patients with major depression are not quite distinguishable.

As in Section 3.4.1, we perform a bootstrap analysis to construct pointwise confidence bands of the group differences of the mean difference IRFs. For each bootstrap sample, we fit the models on the region of interest and on the reference region. Therefore, the group differences of the estimated mean difference IRFs can be obtained for all the bootstrap samples. The 95% pointwise confidence bands shown in the right panel of Figure 3.7 are constructed based on the bootstrap estimates of the mean difference curves. In contrast to what we observed with the midbrain TACs in Section 3.4.1, both confidence bands cover 0 for the entire time range, indicating an insignificant difference between health controls and the other two groups in terms of the mean difference IRFs.

3.5 Discussion

We proposed a nonparametric model fitting framework that estimates the IRF using functional data analytic techniques. For the first time ever, our method models dynamic PET data from multiple subjects simultaneously. In our approach, IRFs are estimated using a linear mixed effects functional data model with population-level fixed effects and subject-level random effects. In accordance with our biological understanding of the IRF, we imposed appropriate non-negativity and monotonicity constraints on the estimates when fitting the model. Because of its flexibility, our model can be used generally for data with any tracer. The application of our approach to clinical PET data indicates that it successfully captures the structure in IRF, both when we model the region of interest only and when we model the difference between

the region of interest and a reference region. Finally, pointwise confidence intervals of the estimated curves were constructed based on bootstrap samples.

In the most general sense one can view the relationship between the TAC and the input function as a function-on-function regression problem (Scheipl et al., 2015), although it may be more appropriately posed as a model with a historical term relating recent tracer availability to current tracer density. Our proposal focuses on the use of functional data approaches to increase the flexibility in estimating the IRF in comparison to methods that focus the estimation of rate parameters in a compartment model. We do this through the convolution of the IRF and input function which reduces the model to a function-on-scalar regression problem. To this framework, we add scientifically relevant constraints on monotonicity and non-negativity to the usual estimation process. A careful consideration of the input function and TAC from the perspective of a function-on-function regression model would allow one to study the adequacy of the convolution operator, and is an important direction for future work.

Additional extension of our methodology may take several directions. Because of the way we construct the model, additional covariates, including continuous variables, can be incorporated in the model if the IRF is thought to be associated with those covariates. In addition, it would be useful to develop a goodness-of-fit test based on the estimated curves to evaluate how well the standard parametric models are able to describe the observed data. The development of approaches for TACs in multiple regions is conceptually possible in our modeling framework, but suitable and computationally feasible models for the covariance across regions may be challenging. Lastly, the classification of subjects into diagnostic groups based on PET imaging data is of general interest. However, given the overlap among groups in our data as shown in Figure 3.4, accurate classification based on PET data alone may not be successful.

Chapter 4

Nonlinear Mixed-Effects Models for PET Data

4.1 Introduction

Dynamic PET imaging has been widely used in studies of mental and neurological disorders. One very common application of PET imaging involves estimating the distribution of various macromolecules, often proteins, throughout the brain. In dynamic PET studies, a time activity curve (TAC) reflects the sequence of concentrations across time for any given voxel or region and is often used to estimate quantities related to the density of the target protein at each location.

The TAC, denoted as C_T, is conceptualized as the convolution between two functions

$$C_T(t) = (H \otimes C_P)(t) \tag{4.1}$$

where t is time, C_P is the *input function* and H is the the voxel-specific *impulse response function* (IRF). The input function C_P represents the concentration of the tracer in the arterial plasma over time, corrected to account for the radioactive

metabolites of the tracer, and quantifies the amount of tracer molecules that are available to enter the brain at any given time. In practice, the input function requires blood data during the scan. The location-specific IRF may be interpreted as the hypothetical concentration of the tracer over time in the corresponding region if the input function were an instantaneous bolus spike. Because the IRF describes the physiological and pharmacological properties of the system, the analysis of the kinetic behavior of the tracer centers on estimating the IRF in Model (4.1).

The most widely used approach for tracer kinetic analysis is compartment modeling. Under the assumptions of this approach, the IRF has the form

$$H_{\boldsymbol{k}}(t) = \sum_{j=1}^{J} L_j e^{-R_j t}, \tag{4.2}$$

where J is the total number of tissue compartments, and L_j and R_j are functions of the rate parameters $\boldsymbol{k}$. The rate parameters are the key elements to be estimated because they completely characterize the kinetic behavior of the tracer based on the assumed model. Standard quantities of clinical importance, such as volume of distribution (V_T) and binding potential (BP_{ND} & BP_P), are functions of the rate parameters. For example, under the assumption of two-tissue compartment model which has four rate parameters (k_1, k_2, k_3, k_4), the forms of these measures are given by (Innis et al., 2007):

$$\begin{aligned} V_T &= \frac{k_1}{k_2}(1 + \frac{k_3}{k_4}); \\ BP_{ND} &= \frac{k_1}{k_2}; \\ BP_P &= \frac{k_1 k_3}{k_2 k_4}. \end{aligned}$$

A well established and almost universally applied method for estimating the rate parameters is a two-stage approach: in Stage 1, individual estimates of all kinetic parameters are obtained by fitting each individual's data, one subject at a time, using

Nonlinear least squares (NLS). Standard measure of binding, including V_T, BP_{ND}, BP_P, etc., can be calculated. Population level effects are estimated by treating the individual estimates as if they were observed data. In Stage 2, the binding measures are compared across subjects and population-level effects, such as the difference between patients and controls, are examined using standard statistical methods. This two-stage approach has several shortcomings. First, sufficient data from each subject are needed to obtain reliable individual-level estimates. In practice, the lack of sufficient data frequently causes numeric instability, especially for complex multi-tissue compartment models which have many parameters to estimate. Finally, the two-stage approach treats subjects individually rather than as members of a common population with shared features, thereby neglecting useful information.

We propose new methods for the analysis of dynamic PET data that provide a flexible and efficient alternative to the two-stage approach. Specifically, we propose to model all subjects simultaneously rather than one at a time by fitting nonlinear mixed-effect (NLME) models. NLME addresses the inherent instability of subject-level rate parameter estimates in the two-stage approach by jointly modeling all subjects, and produces improved individual estimates. This approach accounts for subject-to-subject variability directly by modeling each subject's rate parameter as coming from a distribution of rate parameters. For instance, assuming one distribution for the patients and another for the controls, with an NLME approach, we are really just analyzing the difference between these two distributions. Under the modeling framework of our proposed approach, the difference is limited to a mean shift as we assume there is a shared variability in random effects of both groups. Meanwhile, in the NLME modeling framework, both individual rate parameters and the effects of some covariates on the rate parameters can be estimated in a single analysis. Also, taking this approach allows for more complex models than could be fit otherwise.

In addition to proposing the NLME modeling approach for compartment modeling

with PET data, we also describe a model building procedure that is important when applying our approach to real data. Two main issues that need to be addressed are selecting which explanatory variable should be included as fixed effects and which parameters should have an associated random effect with non-zero variance. We illustrate this model building procedure through the careful analysis of our motivating clinical PET data.

NLME has been used in previous analysis of PET data. Berges et al. (2013) applies the NLME approach to PET data under the assumption of a PK-receptor occupancy (PK-RO) model. This model, which has only one kinetic parameter to estimate, is less complicated than the two-tissue compartment model that we build in this paper. In addition, Veronese et al. (2013) assumes C_P in (4.1) to be a product of the total tracer activity and a Parent Plasma fraction (PPf) function and applies the NLME approach to estimate PPf. In contrast we focus on estimating the IRF directly.

The rest of the paper is organized as follows. In Section 4.2, we present the NLME modeling framework of PET data and introduce the tests and criteria that can be used to determine the fixed and random effects. The results of simulations comparing our proposed NLME method to the two-stage approach are given in Section 4.3, with particular emphasis on power to detect differences across groups. In Section 4.4, we illustrate the model building procedure by using clinical PET data as an example. Finally, we summarize the main results and present a short discussion in Section 4.5. Code files that are used to generate the results in this paper are provided in the supplementary materials.

4.2 Methodology

As the intravascular activity may have significant contribution to the total concentration of the tracer, the whole blood concentration should be accounted for in the data analysis. Under the assumption of compartment models, Model (4.1) can be reformulated as

$$C_T(t) = (1 - V_b)(H_{\boldsymbol{k}} \otimes C_P)(t) + V_b C_B(t) \tag{4.3}$$

where C_B is the time activity curve in the whole blood and V_b is the fractional blood volume of the tissue. Again, $H_{\boldsymbol{k}}$ here is the IRF, and under the assumptions of the compartment models, it depends on a vector of rate parameters $\boldsymbol{k}$.

We now describe how Model (4.3) can be cast in the general NLME framework. A general expression for the NLME model is given by

$$y_{ij} = f(\boldsymbol{\theta}_i, \boldsymbol{z}_{ij}) + \epsilon_{ij},\ i = 1, \cdots, n,\ j = 1, \cdots, n_i, \tag{4.4}$$

where y_{ij} is the jth observation of the ith subject, n is the number of subjects, and n_i is the number of measurements for subject i. In practice, the continuous-valued function C_T in (4.3) is observed on a discrete grid of time points $\{t_{ij}\}$. Therefore, in the context of Model (4.3), the response y_{ij} is the TAC observations $C_T(t_{ij})$; f is defined by the functions $H_{\boldsymbol{k}}$, C_P and C_B; and the parameter vector $\boldsymbol{\theta}_i$, specific to subject i, consists of rate parameters $\boldsymbol{k}_i$ and the fractional blood volume V_{bi}, i.e. $\boldsymbol{\theta}_i = (\boldsymbol{k}_i^T, V_{bi})^T$. The exact form of f under the assumption of compartment models is given in the appendix.

Within this modeling framework, the kinetic parameters for each subject can be thought of as coming from the distribution of kinetic parameters; conceptually, this distribution characterizes the natural subject-to-subject variability. For example, the value k_{4i} for patient i comes from a normal distribution with a mean and variance shared across all patients, while the value k_{4j} for the control subject j comes from

some other normal distribution. A major goal, then, is to determine whether these two distributions are the same. In Model (4.4), we expand the subject-level parameter vector $\boldsymbol{\theta}_i$ as $\boldsymbol{\theta}_i = \boldsymbol{\beta}^T \boldsymbol{x}_i + \boldsymbol{b}_i$, where $\boldsymbol{\beta}$ is a matrix of the fixed effect coefficients; $\boldsymbol{x}_i$ is the design matrix; and $\boldsymbol{b}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ is the vector of the random effects for the ith subject, which represent the subject-specific deviations from the population averages. This expansion separates population averages from subject-specific deviations, and provides the mechanism through which group-level and subject-level kinetic behaviors can be understood. The assumption that random effects $\boldsymbol{b}_i$ share a distribution arises from the recognition that subjects come from the same population. By modeling all subjects simultaneously, the properties of this distribution (especially the variance $\boldsymbol{\Sigma}$) can be inferred and used to stabilize individual-level estimates. This approach allows us to jointly model all subjects and to directly estimate and test for the significance of effects of covariates, such as diagnosis or age, on the rate parameters.

Taken together, the fixed and random effects are the parameters of interest in NLME representation of Model (4.1). These can be estimated either by maximum likelihood (ML) or by restricted maximum likelihood (REML) using the Lindstrom and Bates (LB) algorithm (Lindstrom and Bates, 1990). For our analyses, we use the implementation of this algorithm in the `nlme` package (Pinheiro et al., 2016) in R. From the fitted model, estimates of the summarized measures, like binding potentials (BP_{ND} and BP_P) and total volume (V_T) can be computed at the population level using the fixed effects estimates.

Key issues that arise when fitting NLME models for PET data include selecting covariates to include as fixed effects and determining which elements of the parameter vector $\boldsymbol{\theta}_i$ should have associated random components. Because covariates can affect each of the rate parameters through the fixed effects specification, a global test can be used to assess the global significance of the covariate effect, for example, whether there exist non-zero differences comparing patients to controls for any of the rate

parameters. Choices for determining the fixed effects include the likelihood ratio test (LRT) (Neyman and Pearson, 1992), alternative likelihood-based tests, such as Wald test (Wald, 1943) and score test (Rao, 1948), and information criterion statistics, such as Akaike information criterion (AIC) (Akaike, 1998) and Bayesian information criterion (BIC) (Schwarz et al., 1978). AIC and BIC can also be used to determine which effects should have associated random components. Alternatively, a non-standard likelihood ratio test can be applied on nested models to test whether one random effect component is zero, i.e., whether a parameter has significant between-subject variability. This non-standard LRT, proposed by Stram and Lee (1994), addresses the issue that the testing value under the null hypothesis is on the boundary of the support of the parameter. Lastly, to obtain inferences directly for BP_{ND} , BP_P and V_T the Delta method (Dorfman, 1938) is used to derive the standard errors based on the estimates and variance of rate parameter fixed effects.

4.3 Simulation

In this section we undertake a simulation exercise to understand the properties of NLME modeling for PET data and to compare the performance of the proposed methods to that of the two-stage approach.

Simulated datasets are designed to mimic our motivating data (PET data with WAY tracer described in Section 4.4) in the following way. We begin by fitting Model (4.3) under a two-tissue compartment model assumption to the observed data with no covariates. From this model fit, we extract estimates of fixed effects $\boldsymbol{\beta}$, the random effect covariance $\boldsymbol{\Sigma}$, and of the error variance σ^2 and these estimates are set to be the "truth" for the purposes of this simulation. To simulate new subject data, we sample observed input functions C_P and whole blood time activity curves C_B from subjects in the observed data with replacement. C_P and C_B in our data have the

same forms as in Parsey et al. (2000). Subject-specific random effects are generated from a multivariate normal distribution with mean $\boldsymbol{\beta}$ and covariance $\boldsymbol{\Sigma}$. The sampled functions and generated random effects are combined with estimated fixed effect parameters to produce simulated time activity curves C_T according to Model (4.3). Simulated errors ϵ_{ij} are generated from truncated Gaussian distributions with mean 0 and variance $\sigma^2 = 0.01$ to ensure the simulated TAC observations are non-negative. Each simulated dataset consists of 90 subjects, with half in each covariate group. Values of the $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$ used to generate individual parameters are provided in the appendix.

4.3.1 Quality of fixed effect estimation

Our first objective is to assess how well NLME modeling estimates the "true" fixed effects. To do so, we generate 1000 datasets under the above design. For each of the simulated datasets, we apply both our proposed approach and two stage approach assuming there exists a group effect on all rate parameters but not blood volume. We arbitrarily choose one of the two groups to be the reference group, analogous to the control group in a medical study. In this simulation, every subject has a different set of rate parameters, drawn from a distribution that differs for the control group and the patient group. Also, every subject has a different blood volume V_{bi} and subjects in both groups have observations drawn from a common distribution. Below we compare the estimated values for both approaches to the "truth", i.e., the values used to generate the data.

Figure 4.1 compares the estimates of both approaches for fixed effects related to each rate parameter as well as the summarized measures V_T, BP_{ND} and BP_P. The top row shows the relative estimation errors $\frac{\hat{\beta}-\beta}{\beta}$ of fixed effects for subjects in the reference group, and the bottom row shows the absolute estimation errors $\hat{\beta} - \beta$

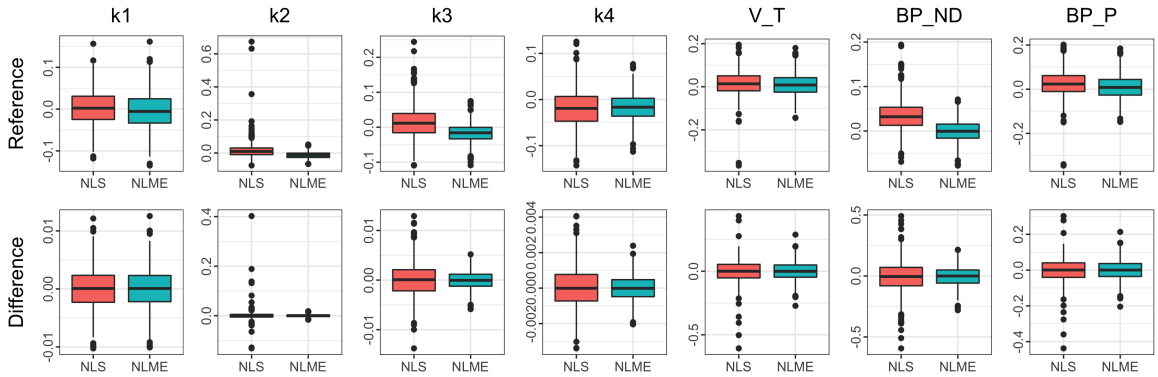

Figure 4.1: Relative estimation errors of fixed effects in the reference group (*top*) and absolute estimation errors of the difference between groups (*bottom*) using both approaches.

of the difference between groups. As expected, the distribution of estimated values for the proposed NLME approach are generally narrower and include fewer outlying values than the corresponding distributions for the two-stage approach. As described in Section 4.2, NLME improves and stabilizes the estimation of rate parameters at the subject level; this, in turn, leads to the observed improvement in estimation for population-level fixed effects.

4.3.2 Comparison of power for detecting group differences

The preceding simulation indicates that the proposed NLME approach is more accurate than two-stage approach for estimating group differences. We now explore how this difference affects the power to detect true differences in rate parameters or in binding measures when testing hypotheses.

We use the simulation design described above, with modifications that allow a careful comparison of power between approaches. Keeping the fixed effects in the reference group as they were, we initially set all group differences to zero. Then,

we gradually increase the group difference for each rate parameter individually while keeping group differences for other rate parameters equal to zero. For each collection of true fixed effects, we apply both approaches to 200 simulated datasets.

First, we compare methods on their ability to detect differences for individual rate parameters. Previous studies focus on testing effects on summarized measurements. However, with our proposed NLME approach, it's possible to test effects on individual rate parameters, e.g., $H_{0\ell} : k_\ell^{\text{Control}} = k_\ell^{\text{Patient}}$, where k_ℓ^{Control} and k_ℓ^{Patient} are k_ℓ's ($\ell \in \{1, 2, 3, 4\}$) for controls and patients, respectively. Here we use standard two-sample t test to examine group differences for the particular rate parameter with a true difference. Results for the NLME approach are obtained directly from the model fitting procedure, while for the two-stage approach we perform a t test on the individual rate parameter estimated from subject-specific NLS fits.

However, in practical settings, these may not be a prior hypothesis about which parameters are affected by covariates. In this case, it is appropriate to use a global test that examines group differences across all rate parameters, e.g., $H_0 : k_1^{\text{Control}} = k_1^{\text{Patient}}, k_2^{\text{Control}} = k_2^{\text{Patient}}, k_3^{\text{Control}} = k_3^{\text{Patient}}, k_4^{\text{Control}} = k_4^{\text{Patient}}$. To test the global hypothesis using NLME, we use a multivariate analysis of variance (MANOVA) test, which is designed to test an effect on several dependent variables. In this case, once we have the individual rate parameters from Stage 1, we can use MANOVA to test whether there is a significant group effect on any of the four rate parameters. To conduct this test using the two-stage approach, we use LRT by fitting two models under different assumptions: none of the rate parameters depend on group and all the rate parameters depend on group. Then the likelihoods of these two models are compared.

Figure 4.2 shows results for true differences in each of the four rate parameters, with power defined as the proportion of rejected null hypotheses across the 200 simulated datasets for each effect size. Unsurprisingly, the parameter-specific test of the

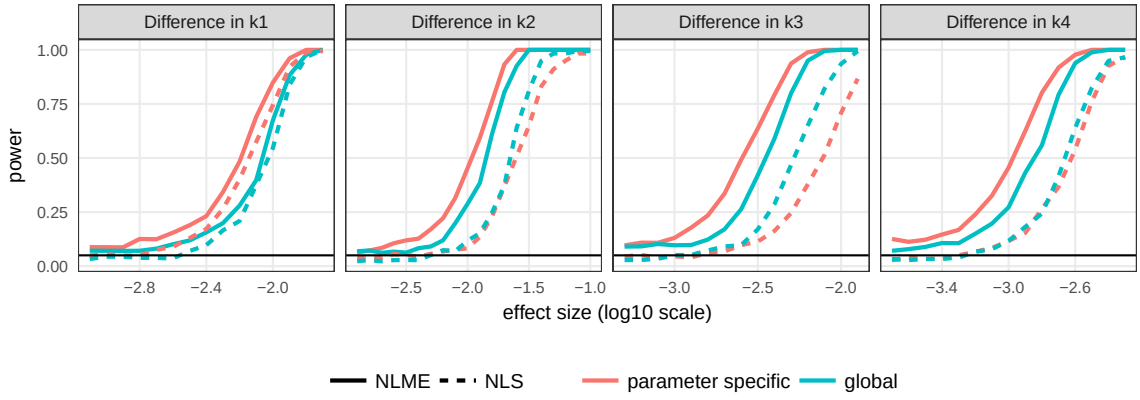

Figure 4.2: Power curves of detecting the group mean difference on rate parameters using four different tests: parameter specific t test of group effect on each rate parameter using NLME model; LRT of overall group effect comparing nested NLME models; parameter specific t test of group effect on the rate parameters based on the two-stage approach; MANOVA test of overall group effect based on the two-stage approach. The black line in each plot represents the 0.05 nominal level.

from $H_{0\ell} : k_\ell^{\text{Control}} = k_\ell^{\text{Patient}}$ is more powerful than the global test in all cases. Importantly, for either test, the NLME approach is more powerful than the two-stage approach, often substantially so. This improvement in power derives from the better estimation of fixed effects observed in Section 4.3.1.

Next, we compare methods on their ability to detect differences in the summary measures V_T, BP_{ND} and BP_P. The simulation design is as before, meaning that differences between groups exist in only one rate parameter at a time. The power to detect resulting differences in summary measures is shown in Figure 4.3. Because k_1 and k_2 do not affect BP_{ND}, group differences in these rate parameters are not detectable through this summary measure. However, V_T and BP_P are affected by such differences and NLME has much greater power than the two-stage approach to detect differences in those measures. Group differences in k_3 and k_4 affect all

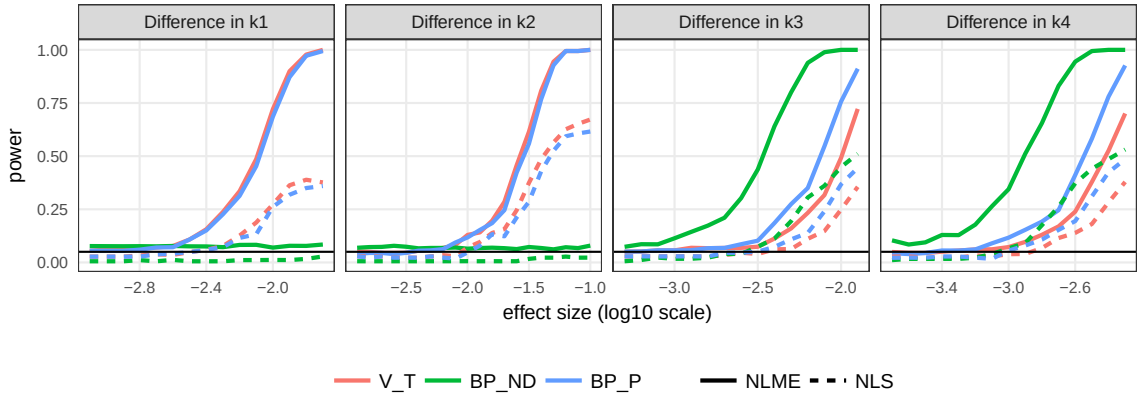

Figure 4.3: Power curves of detecting the group mean difference on the summarized measures using six different tests: t test of group effect in V_T based on NLME model; t test of group effect in BP_{ND} based on NLME model; t test of group effect in BP_P based on NLME model; t test of group effect in V_T based on two-stage approach; t test of group effect in BP_{ND} based on two-stage approach; t test of group effect in BP_P based on two-stage approach. The black line in each plot represents the 0.05 nominal level.

the summary measures, and as these group differences in rate parameters increases so does the power to detect differences in the summary measures. Again, NLME uniformly outperforms the two-stage approach.

4.4 PET data analysis

Recent studies have shown that serotonin 1A receptor (5-HT$_{1A}$) plays a key role in major depressive disorder (MDD) (Parsey et al., 2010) and bipolar disorder (Sullivan et al., 2009). The [^{11}C]WAY tracer has been used widely to quantify 5-HT$_{1A}$ binding and the rate constant parameters when a compartment model is assumed (Parsey et al., 2000).

Our data consist of TACs in the midbrain of 97 subjects who can be divided into three groups based on their prior medication history: MDD subjects who have not recently been on medication (NRM); antidepressant-exposed (AE) MDD subjects and MDD subjects who are on an adequate dose of antidepressant for at least 4 weeks (Parsey et al., 2010); and control subjects. Other covariates that may have effects on the rate parameters include age and gender. Metabolite corrected plasma data and whole blood data are available for all subjects.

In this section, we apply the NLME modeling approach described in Section 4.2 to the PET data under the assumption of a two-tissue compartment model. We use this data as an example to illustrate a model building procedure of NLME models on PET data. A related model building framework for NLME models can be found in Pinheiro et al. (1995). Because the primary interest lies in analyzing the group differences, our starting point is a model that includes this variable. Throughout, we will use global tests under the assumption that covariates may affect the four rate parameters but not the blood volume, and to start we assume that all rate parameters plus V_b have associated random effects.

4.4.1 Testing for random effects

The first question to address is whether all parameters exhibit subject-level variability, i.e., whether a particular parameter is identical for all the subjects with the same fixed effect specification or a parameter-specific random effect is needed. We fit separate models in which each the random effect for each of the parameters is omitted, and compare the results with the initial model using AIC and the LRT described in Section 4.2. The initial model has the smallest AIC among all candidate models, and the p-values from the non-standard LRTs indicate that the random effect on each of the parameter is significant (largest p-value $= 5.667 \times 10^{-4}$). Thus, both

criteria indicate that the model in which all the parameters have associated random components is superior, and we proceed to the selection of covariates for fixed effects.

Results are shown in Table 4.1. Here we refer the model in which only the variance of random effect of k_ℓ is zero as Model ℓ and the model in which only the variance of random effect of V_b is zero as Model 5. For instance, Model 1 represents the model in which the variance of random effect of k_1 is zero and the random effects of other parameters are allowed to have non-zero variances.

Table 4.1: AIC and LRT results for models with different number of random components.

Model	AIC	log likelihood	test	p-value
Model 0	-11617.98	5837.990		
Model 1	-11003.95	5525.974	0 vs 1	6.972×10^{-133}
Model 2	-11596.87	5822.436	0 vs 2	5.917×10^{-6}
Model 3	-11606.99	5827.497	0 vs 3	5.667×10^{-4}
Model 4	-11588.86	5818.430	0 vs 4	1.452×10^{-7}
Model 5	-11579.03	5813.515	0 vs 5	1.435×10^{-9}

4.4.2 Including covariate fixed effects

Covariates such as age and gender may affect rate parameters, and we now consider their addition to our model. We add these variables as fixed effects in a global way by including covariate effects on all rate parameters, and build our model using forward selection with a global hypothesis test. Both main effects and interactions between variables are considered. The results of our model building process are shown in Table 4.2, and indicate that age and gender are significant predictors and none of the two-way interactions are significant. Therefore, we determine that the model with

only the main effects of group, age and gender is our final model. All the estimates of the final model are given in Table 4.3. The estimated correlation matrix of the random effects is

$$\begin{pmatrix} 1 & -0.552 & -0.667 & -0.308 & 0.157 \\ -0.552 & 1 & 0.104 & 0.476 & -0.081 \\ -0.667 & 0.104 & 1 & 0.664 & -0.160 \\ -0.308 & 0.476 & 0.664 & 1 & -0.310 \\ 0.157 & -0.081 & -0.160 & -0.310 & 1 \end{pmatrix}$$

Among the rate parameters, the correlation between k_2 and $k3$, as well as the correlation between k_1 and k_4, is small. And V_b has a weak correlation with all the rate parameters.

Table 4.2: Results of LRT comparing nested models with difference combination of covariates

Model	Fixed effect structure	Test	p-value
1	Group		
2	Group + Gender	1 vs 2	0.0293
3	Group + Age	1 vs 3	0.0013
4	Group + Gender + Age	3 vs 4	0.0154
5	Group + Gender + Age + Gender * Age	4 vs 5	0.4331
6	Group + Gender + Age + Group * Gender	4 vs 6	0.4100
7	Group + Gender + Age + Group * Age	4 vs 7	0.6280

4.4.3 Comparison with the two-stage approach

Next, we compare NLS estimates of the parameters obtained from the two-stage approach by fitting a two-tissue compartment model on each subject to those obtained

from NLME fit of the final model. Figure 4.4 plots the these estimates for all the parameters and includes an identity line for reference. The approaches give similar estimates for k_1, but the impact of assuming a random effects structure is clear for k_2, k_3, k_4 and V_b: the NLME estimates have smaller variances. This "shrinkage" is expected from the NLME approach, and is a reason why the approach is less vulnerable to individual outliers than NLS estimates. That is, it is difficult to obtain accurate and stable estimates for these rate parameters using NLS. In contrast, by simultaneously estimating rate parameters for all subjects and using the random effects distribution, NLME is able to balance subject- and population-level data to improve rate parameter estimation.

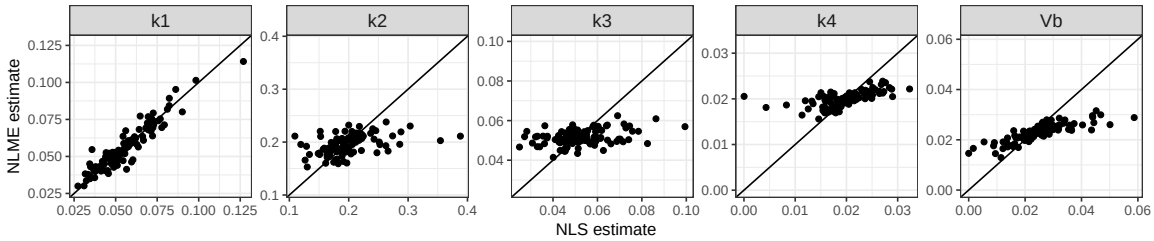

Figure 4.4: Individual NLME estimates vs NLS estimates for the five parameters. The solid line on each panel is the identity line with intercept 0 and slope 1.

The parameter estimates, standard errors and the p-values of the t tests, are given in Table 4.3. According to the p-values of t tests associated with the comparisons of different covariates, our final NLME model identifies gender as a significant factor for k_1, k_2, V_T, BP_{ND} and BP_P; and age is a significant factor for k_3, BP_{ND} and BP_P. We can draw many conclusions based on the results. For example, adjusted for group and age, k_1 of males is 8.232×10^{-3} less than k_1 of females. Also, adjusted for group and gender, as the age increases by 1, k_3 decreases by 1.454×10^{-4}.

Table 4.4 shows the significance level of the overall effects of the covariates in both NLME and the two-stage approaches. Likelihood ratio tests are performed to

assess the global effects on all rate parameters for NLME while the MANOVA F-tests are used for the two-stage approach. Both models identify age as a significant factor, but only the NLME approach detects a significant overall effect of gender. Neither approach suggests a significant overall effect of prior medication history group, although the p-value is somewhat smaller for NLME than for the two-stage approach.

4.5 Conclusion

We proposed a NLME approach for compartment modeling of PET data. The NLME approach addresses known shortcomings of the standard two-stage approach by fitting all subjects simultaneously and estimating covariate effects in a one-step model process. Our simulations indicate that the proposed NLME approach is more accurate and correspondingly more powerful in detecting group differences than the two-stage approach. In real data analyses, the NLME estimates of individual rate parameters often had narrower distributions than estimates derived from two-stage approach, an expected byproduct of the balancing subject and population data to estimate individual effects. We applied a model building procedure for the NLME approach to WAY tracer based on the two-tissue compartment model, and found effects not detected by a two-stage approach.

The instability of NLS for estimating rate parameters is a frequently encountered issue in practice. One way to control outlier rate parameter estimates is to set bounds. However, it is arbitrary and would have to set separately for each tracer. These bounds artificially reduce the range of rate parameters, and introduce a new problem of sensitivity to their specification. Additionally, such bounds still result in individual estimates of rate parameters, and group differences must be assessed in a two-stage approach. In contrast, our NLME approach is based on a statistically principled model technique that uses available data to stabilize individual rate parameter

estimates and assesses covariate effects in a single step.

Our work has focused on the two-tissue compartment model; extending this to a more complicated three-tissue compartment model will introduce additional complexity but which will be important in some applications. Another direction we might take includes developing an NLME modeling approach to model multiple regions simultaneously to account for heterogeneity across regions. Code files that are used for the simulation and data analysis in this paper are provided in the supplementary materials.

Table 4.3: Results of the NLME and two-stage approaches

Parameter	Variable	NLME			Two-stage approach		
		Estimate	Std.Error	p-value	Estimate	Std.Error	p-value
k_1	AE vs Control	-6.183×10^{-3}	4.023×10^{-3}	0.125	-6.821×10^{-3}	4.167×10^{-3}	0.105
k_1	NRM vs Control	5.414×10^{-3}	4.179×10^{-3}	0.195	3.827×10^{-3}	4.235×10^{-3}	0.369
k_1	Gender	-8.232×10^{-3}	3.302×10^{-3}	0.013	-7.028×10^{-3}	3.407×10^{-3}	0.042
k_1	Age	-2.145×10^{-5}	1.200×10^{-4}	0.858	-3.478×10^{-5}	1.256×10^{-4}	0.783
k_2	AE vs Control	7.854×10^{-3}	7.959×10^{-3}	0.324	-1.502×10^{-3}	1.169×10^{-2}	0.898
k_2	NRM vs Control	1.414×10^{-2}	8.700×10^{-3}	0.104	2.240×10^{-3}	1.188×10^{-2}	0.851
k_2	Gender	-1.476×10^{-2}	6.546×10^{-3}	0.024	-1.101×10^{-2}	9.557×10^{-3}	0.252
k_2	Age	1.152×10^{-4}	2.134×10^{-4}	0.589	1.546×10^{-5}	3.525×10^{-4}	0.965
k_3	AE vs Control	-1.912×10^{-3}	2.450×10^{-3}	0.435	-3.728×10^{-3}	3.433×10^{-3}	0.280
k_3	NRM vs Control	2.176×10^{-3}	2.680×10^{-3}	0.417	6.975×10^{-4}	3.489×10^{-3}	0.842
k_3	Gender	-2.429×10^{-3}	2.074×10^{-3}	0.242	-2.084×10^{-3}	2.807×10^{-3}	0.460
k_3	Age	-1.454×10^{-4}	6.482×10^{-5}	0.025	-1.214×10^{-4}	1.035×10^{-4}	0.244
k_4	AE vs Control	-2.224×10^{-5}	9.676×10^{-4}	0.982	-1.301×10^{-4}	1.331×10^{-3}	0.922
k_4	NRM vs Control	-5.445×10^{-4}	9.634×10^{-4}	0.572	-4.758×10^{-4}	1.353×10^{-3}	0.726
k_4	Gender	5.003×10^{-4}	7.774×10^{-4}	0.520	5.360×10^{-4}	1.088×10^{-3}	0.624
k_4	Age	3.072×10^{-5}	2.593×10^{-5}	0.236	6.252×10^{-5}	4.014×10^{-5}	0.123
V_T	AE vs Control	-1.992×10^{-1}	9.498×10^{-2}	0.039	-4.509	5.484	0.413
V_T	NRM vs Control	6.454×10^{-2}	9.654×10^{-2}	0.506	-3.984	5.574	0.477
V_T	Gender	-1.655×10^{-1}	7.765×10^{-2}	0.036	-4.190	4.484	0.352
V_T	Age	-4.699×10^{-3}	2.864×10^{-3}	0.104	6.216×10^{-2}	1.654×10^{-1}	0.708
BP_{ND}	AE vs Control	-1.036×10^{-1}	4.330×10^{-2}	0.019	-21.59	27.29	0.431
BP_{ND}	NRM vs Control	1.809×10^{-1}	4.401×10^{-2}	<0.001	-20.08	27.74	0.471
BP_{ND}	Gender	-1.940×10^{-1}	3.540×10^{-2}	<0.001	-20.25	22.31	0.366
BP_{ND}	Age	-1.167×10^{-2}	1.306×10^{-3}	<0.001	3.200×10^{-1}	8.228×10^{-1}	0.698
BP_P	AE vs Control	-1.531×10^{-1}	6.978×10^{-2}	0.031	-4.468	5.487	0.418
BP_P	NRM vs Control	6.271×10^{-2}	7.093×10^{-2}	0.379	-3.994	5.577	0.476
BP_P	Gender	-1.364×10^{-1}	5.705×10^{-2}	0.019	-4.163	4.486	0.356
BP_P	Age	-4.303×10^{-3}	2.104×10^{-3}	0.044	6.237×10^{-2}	1.655×10^{-1}	0.707

Table 4.4: p-values of overall effects in NLME and two-stage approaches

	NLME	Two-stage approach
Group	0.0798	0.1140
Gender	0.0154	0.1376
Age	0.0006	0.0006

Chapter 5

Conclusions

The overall theme of this thesis focuses on methods for functional regression and nonlinear mixed-effects models with applications to PET data.

Chapter 2 considers the problem of variable selection in regression models with functional responses and scalar predictors. Few methods for variable selection exist for function-on-scalar models, and none account for the inherent correlation of residual curves in such models. By expanding the coefficient functions using a B-spline basis, we pose the function-on-scalar model as a multivariate regression problem. Spline coefficients are grouped within coefficient function, and group-MCP is used for variable selection. We adapt techniques from generalized least squares to account for residual covariance by "pre-whitening" using an estimate of the covariance matrix, and establish theoretical properties for the resulting estimator. We further develop an iterative algorithm that alternately updates the spline coefficients and covariance; simulation results indicate that this iterative algorithm often performs as well as pre-whitening using the true covariance, and substantially outperforms methods that neglect the covariance structure. We apply our method to two-dimensional planar reaching motions in a study of the effects of stroke severity on motor control, and find that our

method provides lower prediction errors than competing methods.

Chapter 3 introduces a functional data analytic approach that models multiple subjects simultaneously, and estimates the IRF nonparametrically. One application of PET, a nuclear imaging technique, in neuroscience involves in vivo estimation of the density of various proteins (often, neuroreceptors) in the brain. PET scanning begins with the injection of a radiolabeled tracer that binds preferentially to the target protein; tracer molecules are then continuously delivered to the brain via the bloodstream. By detecting the radioactive decay of the tracer over time, dynamic PET data are constructed to reflect the concentration of the target protein in the brain at each time. The fundamental problem in the analysis of dynamic PET data involves estimating the IRF, which is necessary for describing the binding behavior of the injected radiotracer. Virtually all existing methods have three common aspects: summarizing the entire IRF with a single scalar measure; modeling each subject separately; and the imposition of parametric restrictions on the IRF. In contrast, we propose a functional data analytic approach that regards each subject's IRF as the basic analysis unit, models multiple subjects simultaneously, and estimates the IRF nonparametrically. We pose our model as a linear mixed effect model in which population level fixed effects and subject-specific random effects are expanded using a B-spline basis. Shrinkage and roughness penalties are incorporated in the model to enforce identifiability and smoothness of the estimated curves, respectively, while monotonicity and non-negativity constraints impose biological information on estimates. We illustrate this approach by applying it to clinical PET data with subjects belonging to three diagnosic groups. We explore differences among groups by means of pointwise confidence intervals of the estimated mean curves based on bootstrap samples.

Chapter 4 discusses a nonlinear mixed-effects modeling approach for PET data analysis. The kinetic behavior of many tracers used in neurological mapping studies

can be approximated using a compartment model. The rate parameters of tracer transferring between compartments are estimated using NLS approach. The NLS estimators of the population parameters are applied in a two-stage analysis, in which individual estimates are obtained by fitting models subject-by-subject and population estimates are subsequently computed by treating individual estimates as observed data. This approach brings instability issue and neglects the variation in rate parameters. We propose to estimate the rate parameters by fitting nonlinear mixed-effects (NLME) models, which addresses both concerns of NLS. In the NLME framework, all the subjects are modeled simultaneously by allowing rate parameters to have random effects and population parameters can be estimated directly from the joint model. Simulations are conducted to compare the power of detecting group effect in both rate parameters and summarized measures of tests based on both NLS and NLME models. The results indicate that the test based on NLME model has greater power compared to its NLS counterpart. We apply our NLME approach to clinical PET data to illustrate the model building procedure including selecting fixed effects and determining random effect.

In future research, we will consider adding roughness penalty in addition to the variable selection to enforce smoothness of the coefficient functions when fitting function-on-scalar regression models. Additionally, it is worthwhile to develop goodness-of-fit testS based on the estimated curves to evaluate how well our nonparametric and parametric methods are able to describe the observed PET data.

Bibliography

Abramovich, F., Bailey, T. C., and Sapatinas, T. (2000). Wavelet analysis and its statistical applications. Journal of the Royal Statistical Society. Series D (The Statistician) **49**:pp. 1–29.

Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In Selected Papers of Hirotugu Akaike, pp. 199–213. Springer.

Barber, R. F., Reimherr, M., and Schill, T. (2015). The Function-on-Scalar LASSO with Applications to Longitudinal GWAS (technical report).

Berges, A., Cunningham, V. J., Gunn, R. N., and Zamuner, S. (2013). Non linear mixed effects analysis in pet pk-receptor occupancy studies. NeuroImage **76**:155–166.

Breheny, P. and Huang, J. (2013). Group descent algorithms for nonconvex penalized linear and logistic regression models with grouped predictors. Statistics and Computing **23**:1–15.

Ciprian Crainiceanu, Philip Reiss, Jeff Goldsmith, Lei Huang, Lan Huo, and Fabian Scheipl (2014). refund: Regression with Functional Data. R package version 0.1-11.

Cunningham, V. J. and Jones, T. (1993). Spectral analysis of dynamic PET studies. Journal of Cerebral Blood Flow & Metabolism **13**:15–23.

de Boor, C. (1978). A practical guide to splines. Appl. Math. Sci. Springer, New York, NY.

Di, C.-Z., Crainiceanu, C. M., Caffo, B. S., and Punjabi, N. M. (2009). Multilevel functional principal component analysis. The Annals of Applied Statistics **3**:458–488.

Donoho, D. L. and Johnstone, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage. Biometrika **81**:425–455.

Dorfman, R. (1938). A note on the delta-method for finding variance formulae. The Biometric Bulletin **1**:92.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. Journal of the American Statistical Association **96**:1348–1360.

Frankle, W. G., Slifstein, M., Gunn, R. N., Huang, Y., Hwang, D.-R., Darr, E. A., Narendran, R., Abi-Dargham, A., and Laruelle, M. (2006). Estimation of serotonin transporter parameters with 11c-dasb in healthy humans: Reproducibility and comparison of methods. Journal of Nuclear Medicine **47**:815–826.

Fugl-meyer, A. R., Jaasko, L., Leyman, I., Olsson, S., and Steglind, S. (1975). The post-stroke hemiplegic patient: a method for evaluation of physical performance. Scandinavian Journal of Rehabilitation Medicine **7**:13–31.

Gertheiss, J., Maity, A., and Staicu, A.-M. (2013). Variable selection in generalized functional linear models. Stat **2**:86–101.

Goldsmith, J. and Kitago, T. (2015). Assessing systematic effects of stroke on motor control using hierarchical function-on-scalar regression. Journal of the Royal Statistical Society: Series C (Accepted).

Gunn, R. N., Gunn, S. R., and Cunningham, V. J. (2001). Positron emission tomography compartmental models. Journal of Cerebral Blood Flow & Metabolism **21**:635–652.

Gunn, R. N., Gunn, S. R., Turkheimer, F. E., Aston, J. A., and Cunningham, V. J. (2002). Positron emission tomography compartmental models: A basis pursuit strategy for kinetic modeling. Journal of Cerebral Blood Flow & Metabolism **22**:1425–1439.

Innis, R., Cunningham, V., Delforge, J., Fujita, M., Gjedde, A., Gunn, R., Holden, J., Houle, S., Huang, S.-C., Ichise, M., Iida, H., Ito, H., Kimura, Y., Koeppe, R., Knudsen, G., Knuuti, J., Lammertsma, A., Laruelle, M., Logan, J., Maguire, R., Mintun, M., Morris, E., Parsey, R., Price, J., Slifstein, M., Sossi, V., Suhara, T., Votaw, J., Wong, D., and Carson, R. (2007). Consensus nomenclature for in vivo imaging of reversibly binding radioligands. Journal of Cerebral Blood Flow & Metabolism **27**:1533–1539.

James, G. M. (2002). Generalized linear models with functional predictors. Journal of the Royal Statistical Society: Series B (Statistical Methodology) **64**:411–432.

Jiang, C.-R., Aston, J. A., and Wang, J.-L. (2015). A functional approach to deconvolve dynamic neuroimaging data. Journal of the American Statistical Association (in press).

Jiang, H. and Ogden, R. T. (2008). Mixture modeling for dynamic pet data. Statistica Sinica **18**:1341.

Lin, Y., Haldar, J. P., Li, Q., Conti, P. S., and Leahy, R. M. (2014). Sparsity constrained mixture modeling for the estimation of kinetic parameters in dynamic pet. Medical Imaging, IEEE Transactions on **33**:173–185.

Lindstrom, M. J. and Bates, D. M. (1990). Nonlinear mixed effects models for repeated measures data. Biometrics pp. 673–687.

Logan, J., Fowler, J. S., Volkow, N. D., Wolf, A. P., Dewey, S. L., Schlyer, D. J., MacGregor, R. R., Hitzemann, R., Bendriem, B., Gatley, S. J., and Christman, D. R. (1990). Graphical analysis of reversible radioligand binding from time-activity measurements applied to [n-11c-methyl]-(-)-cocaine pet studies in human subjects. Journal of Cerebral Blood Flow & Metabolism **10**:740–747.

Malison, R. T., Price, L. H., Berman, R., Van Dyck, C. H., Pelton, G. H., Carpenter, L., Sanacora, G., Owens, M. J., Nemeroff, C. B., Rajeevan, N., et al. (1998). Reduced brain serotonin transporter availability in major depression as measured by [123 i]-2β-carbomethoxy-3β-(4-iodophenyl) tropane and single photon emission computed tomography. Biological psychiatry **44**:1090–1098.

Mallat, S. (1989). A theory for multiresolution signal decomposition: the wavelet representation. Pattern Analysis and Machine Intelligence, IEEE Transactions on **11**:674–693.

McLean, M. W., Hooker, G., Staicu, A.-M., Scheipl, F., and Ruppert, D. (2014). Functional generalized additive models. Journal of Computational and Graphical Statistics **23**:249–269.

Meier, P. and Zierler, K. L. (1954). On the theory of the indicator-dilution method for measurement of blood flow and volume. Journal of Applied Physiology **6**:731–744.

Miller, J. M., Everett, B. A., Oquendo, M. A., Ogden, R. T., Mann, J. J., and Parsey, R. V. (2016). Positron emission tomography quantification of serotonin transporter binding in medication-free bipolar disorder. Synapse **70**:24–32.

Miller, J. M., Hesselgrave, N., Ogden, R. T., Sullivan, G. M., Oquendo, M. A., Mann, J. J., and Parsey, R. V. (2013). Positron emission tomography quantification of serotonin transporter in suicide attempters with major depressive disorder. Biological Psychiatry **74**:287–295.

Morris, J. S. (2015). Functional regression. Annual Review of Statistics and Its Application **2**:321–359.

Müller, H.-G. and Stadtmüller, U. (2005). Generalized functional linear models. The Annals of Statistics **33**:774–805.

Müller, H.-G. and Yao, F. (2008). Functional additive models. Journal of the American Statistical Association **103**:1534–1544.

Neyman, J. and Pearson, E. S. (1992). On the problem of the most efficient tests of statistical hypotheses. In Breakthroughs in statistics, pp. 73–108. Springer.

Ogden, R. T., Ojha, A., Erlandsson, K., Oquendo, M. A., Mann, J. J., and Parsey, R. V. (2007). In vivo quantification of serotonin transporters using 11c-dasb and positron emission tomography in humans: Modeling considerations. Journal of Cerebral Blood Flow & Metabolism **27**:205–217.

O'Sullivan, F., Muzi, M., Spence, A. M., Mankoff, D. M., O'Sullivan, J. N., Fitzgerald, N., Newman, G. C., and Krohn, K. A. (2009). Nonparametric residue analysis of dynamic pet data with application to cerebral fdg studies in normals. Journal of the American Statistical Association **104**:556–571.

Parsey, R. V., Hastings, R. S., Oquendo, M. A., Huang, Y.-y., Simpson, N., Arcement, J., Huang, Y., Ogden, R. T., Van Heertum, R. L., Arango, V., et al. (2006a). Lower serotonin transporter binding potential in the human brain during major depressive episodes. American Journal of Psychiatry **163**:52–58.

Parsey, R. V., Ogden, R. T., Miller, J. M., Tin, A., Hesselgrave, N., Goldstein, E., Mikhno, A., Milak, M., Zanderigo, F., Sullivan, G. M., et al. (2010). Higher serotonin 1a binding in a second major depression cohort: modeling and reference region considerations. Biological psychiatry **68**:170–178.

Parsey, R. V., Ojha, A., Ogden, R. T., Erlandsson, K., Kumar, D., Landgrebe, M., Van Heertum, R., and Mann, J. J. (2006b). Metabolite considerations in the in vivo quantification of serotonin transporters using 11c-dasb and pet in humans. Journal of Nuclear Medicine **47**:1796–1802.

Parsey, R. V., Slifstein, M., Hwang, D.-R., Abi-Dargham, A., Simpson, N., Mawlawi, O., Guo, N.-N., Van Heertum, R., Mann, J. J., and Laruelle, M. (2000). Validation and reproducibility of measurement of 5-ht1a receptor parameters with [carbonyl-11c] way-100635 in humans: comparison of arterial and reference tissue input functions. Journal of Cerebral Blood Flow & Metabolism **20**:1111–1133.

Peng, H. and Lu, Y. (2012). Model selection in linear mixed effect models. Journal of Multivariate Analysis **109**:109 – 129.

Peng, J.-Y., Aston, J. A., Gunn, R. N., Liou, C.-Y., and Ashburner, J. (2008). Dynamic positron emission tomography data-driven analysis using sparse bayesian learning. IEEE Transactions on Medical Imaging **27**:1356–1369.

Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., and R Core Team (2016). nlme: linear and nonlinear mixed effects models. R package version 3.1-128.

Pinheiro, J. C., Bates, D. M., and Lindstrom, M. J. (1995). Model building for nonlinear mixed effects models. Citeseer.

Ramsay, J. O. and Dalzell, C. J. (1991). Some tools for functional data analysis.

Journal of the Royal Statistical Society: Series B (Statistical Methodology) **53**:539–572.

Ramsay, J. O. and Silverman, B. W. (2005). Functional Data Analysis. Springer Series in Statistics. Springer, second edition.

Rao, C. R. (1948). Large sample tests of statistical hypotheses concerning several parameters with applications to problems of estimation. In Mathematical Proceedings of the Cambridge Philosophical Society, volume 44, pp. 50–57. Cambridge Univ Press.

Rao, C. R. (1958). Some statistical methods for comparison of growth curves. Biometrics **14**:pp. 1–17.

Reiss, P., Huang, L., and Mennes, M. (2010). Fast function-on-scalar regression with penalized basis expansions. International Journal of Biostatistics **6**:article–28.

Reiss, P. T. and Ogden, R. T. (2010). Functional generalized linear models with images as predictors. Biometrics **66**:61–69.

Rice, J. A. and Silverman, B. W. (1991). Estimating the mean and covariance structure nonparametrically when the data are curves. Journal of the Royal Statistical Society. Series B (Methodological) **53**:pp. 233–243.

Ruppert, D. (2002). Selecting the number of knots for penalized splines. Journal of Computational and Graphical Statistics **11**:735–757.

Scheipl, F., Staicu, A.-M., and Greven, S. (2015). Functional additive mixed models. Journal of Computational and Graphical Statistics **24**:477–501.

Schwarz, G. et al. (1978). Estimating the dimension of a model. The annals of statistics **6**:461–464.

Silverman, B. W. (1996). Smoothed functional principal components analysis by choice of norm. The Annals of Statistics **24**:1–24.

Slifstein, M. and Laruelle, M. (2001). Models and methods for derivation of in vivo neuroreceptor parameters with pet and spect reversible radiotracers. Nuclear Medicine and Biology **28**:595–608.

Stram, D. O. and Lee, J. W. (1994). Variance components testing in the longitudinal mixed effects model. Biometrics pp. 1171–1177.

Sullivan, G. M., Ogden, R. T., Oquendo, M. A., Kumar, J. D., Simpson, N., Huang, Y.-Y., Mann, J. J., and Parsey, R. V. (2009). Positron emission tomography quantification of serotonin-1a receptor binding in medication-free bipolar depression. Biological Psychiatry **66**:223–230.

Tibshirani, R. (1994). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Statistical Methodology) **58**:267–288.

Tucker, L. (1958). Determination of parameters of a functional relation by factor analysis. Psychometrika **23**:19–23.

Veronese, M., Gunn, R. N., Zamuner, S., and Bertoldo, A. (2013). A non-linear mixed effect modelling approach for metabolite correction of the arterial input function in pet studies. Neuroimage **66**:611–622.

Wald, A. (1943). Tests of statistical hypotheses concerning several parameters when the number of observations is large. Transactions of the American Mathematical society **54**:426–482.

Wang, L., Chen, G., and Li, H. (2007). Group SCAD regression analysis for microarray time course gene expression data. Bioinformatics **23**:1486–1494.

Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. Journal of the Royal Statistical Society (B) **73**:3–36.

Yao, F., MÃijller, H.-G., and Wang, J.-L. (2005a). Functional linear regression analysis for longitudinal data. The Annals of Statistics **33**:2873–2903.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005b). Functional data analysis for sparse longitudinal data. Journal of the American Statistical Association **100**:577–590.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. Journal of the Royal Statistical Society: Series B **68**:49–67.

Zanderigo, F., Parsey, R. V., and Ogden, R. T. (2015). Model-free quantification of dynamic pet data using nonparametric deconvolution. Journal of Cerebral Blood Flow & Metabolism **35**:1368–1379.

Zeng, F. and Goodman, M. M. (2013). Fluorine-18 radiolabeled heterocycles as pet tracers for imaging β-amyloid plaques in alzheimer's disease. Current Topics in Medicinal Chemistry **13**:909–919.

Zeng, L. and Xie, J. (2014). Group variable selection via SCAD-L2. Statistics **48**:49–66.

Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. The Annals of Statistics **38**:894–942.

Zhao, Y., Ogden, R. T., and Reiss, P. T. (2012). Wavelet-based lasso in functional linear regression. Journal of Computational and Graphical Statistics **21**:600–617.

Appendix A

Appendices to: Variable Selection in Function-on-Scalar Regression

A.1 Proof of the theorems

These proofs follow the same general strategy of Zeng and Xie (2014) and Peng and Lu (2012).

For convenience of notation, we denote $\text{vec}(\boldsymbol{Y}^T)$ as $\boldsymbol{W}$, $\boldsymbol{X} \otimes \boldsymbol{\Phi}$ as $\boldsymbol{Z}$, $\text{vec}(\boldsymbol{B}^T)$ as $\boldsymbol{\theta}$ and $\text{vec}(\boldsymbol{E}^T)$ as $\boldsymbol{\epsilon}$. Then Model (2.5) can be rewritten as

$$\boldsymbol{W} = \boldsymbol{Z}\boldsymbol{\theta} + \boldsymbol{\epsilon}$$

where $\boldsymbol{W} = (\boldsymbol{w}_1^T, ..., \boldsymbol{w}_n^T)^T$ is a vector of length nD with $\boldsymbol{w}_i = (w_{i1}, ..., w_{iD})^T, i = 1, ..., n$; $\boldsymbol{Z} = (\boldsymbol{z}_1^T, ..., \boldsymbol{z}_n^T)^T$ is a $nD \times pK$ matrix with $\boldsymbol{z}_i = (\boldsymbol{z}_{i1}, ..., \boldsymbol{z}_{ip}), i = 1, ..., n$, where $\boldsymbol{z}_{ij}$, $j = 1 ..., p$ is a $D \times K$ matrix; and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^T, ..., \boldsymbol{\theta}_p^T)^T$ is a vector of length Kp with $\boldsymbol{\theta}_j = (\theta_{j1}, ..., \theta_{jK})^T, j = 1, ..., p$. Without loss of generality, we assume the first s groups of coefficients, $\boldsymbol{\theta}_+ = (\boldsymbol{\theta}_1^T, ..., \boldsymbol{\theta}_s^T)^T$, are nonzero and the rest $p - s$ groups of coefficients, $\boldsymbol{\theta}_0 = (\boldsymbol{\theta}_{s+1}^T, ..., \boldsymbol{\theta}_p^T)^T$, are zeros Let $\boldsymbol{Z}_+$ denote the design matrix

associated with $\boldsymbol{\theta}_+$ and $\boldsymbol{Z}_0$ denote the one associated with $\boldsymbol{\theta}_0$. Therefore, we have $\boldsymbol{Z} = (\boldsymbol{Z}_+|\boldsymbol{Z}_0)$ and $\boldsymbol{\theta} = (\boldsymbol{\theta}_+^T, \boldsymbol{\theta}_0^T)^T$. Also, $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{V})$ where $\boldsymbol{V}$ is a $nD \times nD$ block diagonal matrix with diagonal elements, the $D \times D$ matrix $\boldsymbol{\Sigma}$, i.e. $\boldsymbol{V} = \boldsymbol{I}_n \otimes \boldsymbol{\Sigma}$. The assumptions required for the proofs are

1. $\lim_{n\to\infty} \frac{1}{nD} \boldsymbol{Z}^T \boldsymbol{Z}$ is a positive definite matrix;
2. $\lambda_n \to 0$ and $\sqrt{n}\lambda_n \to \infty$ as $n \to \infty$;
3. there exists a $\sqrt{n}$-consistent estimate $\hat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$;
4. the tuning parameter γ of the penalty is fixed.

Proof of Theorem 1:

The penalized least square objective function for estimating $\boldsymbol{\theta}$ is

$$\begin{aligned} Q(\boldsymbol{\theta}) &= \frac{1}{2}(\boldsymbol{W} - \boldsymbol{Z}\boldsymbol{\theta})^T \hat{\boldsymbol{V}}^{-1} (\boldsymbol{W} - \boldsymbol{Z}\boldsymbol{\theta}) + nD \sum_{j=1}^{p} p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||) \\ &= L(\boldsymbol{\theta}) + nD \sum_{j=1}^{p} p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||). \end{aligned}$$

Let's consider a ball $B = \{\boldsymbol{\theta} + n^{-1/2}\boldsymbol{u} : ||\boldsymbol{u}|| \le C\}$ where C is a constant. Since B is a compact set and $Q(\boldsymbol{\theta})$ is a continuous function on B, there exists a minimum of $Q(\boldsymbol{\theta})$ on B. If $Q(\boldsymbol{\theta}^*) > Q(\boldsymbol{\theta})$ for every $\boldsymbol{\theta}^*$ on the boundary of B, then there exists a local minimizer inside the ball B.

Therefore it suffices to show that for any given $\epsilon > 0$, there exists a large constant C such that

$$P\left\{ \inf_{||u||=C} Q(\boldsymbol{\theta} + n^{-1/2}\boldsymbol{u}) > Q(\boldsymbol{\theta}) \right\} > 1 - \epsilon. \tag{A.1}$$

This implies with probability at least $1-\epsilon$ that there exists a local minimizer in the ball $\{\boldsymbol{\theta}+n^{-1/2}\boldsymbol{u} : ||\boldsymbol{u}|| \leq C\}$. Equivalently, for any given $\epsilon > 0$, there exists a constant C such that $P\{n^{-1/2}||\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}|| < C\} \geq 1-\epsilon$, where $\hat{\boldsymbol{\theta}}$ is the local minimizer that satisfies $||\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}|| = O_p(n^{-1/2})$.

Since $p_{\lambda_m,\gamma}(0) = 0$, we have

$$\begin{aligned}
Q(\boldsymbol{\theta}+n^{-1/2}\boldsymbol{u}) - Q(\boldsymbol{\theta}) &\geq L(\boldsymbol{\theta}+n^{-1/2}\boldsymbol{u}) - L(\boldsymbol{\theta}) + nD\sum_{j=1}^{s} p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j + n^{-1/2}\boldsymbol{u}_j||) - nD\sum_{j=1}^{s} p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||) \\
&= \frac{1}{2}n^{-1}\boldsymbol{u}^T\boldsymbol{Z}^T\hat{\boldsymbol{V}}^{-1}\boldsymbol{Z}\boldsymbol{u} - n^{-1/2}\boldsymbol{u}^T\boldsymbol{Z}^T\hat{\boldsymbol{V}}^{-1}\boldsymbol{\epsilon} \\
&\quad + nD\sum_{j=1}^{s}\Big[p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j + n^{-1/2}\boldsymbol{u}_j||) - p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)\Big] \\
&\triangleq I_1 + I_2 + I_3.
\end{aligned}$$

Under the assumption $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{V})$, since $\lim_{n\to\infty}\frac{1}{nD}\boldsymbol{Z}^T\boldsymbol{Z}$ is positive definite and $\hat{\boldsymbol{V}} \to \boldsymbol{V}$, we have

$$I_1 = \frac{D}{2}\boldsymbol{u}^T[\frac{1}{nD}\boldsymbol{Z}^T\hat{\boldsymbol{V}}^{-1}\boldsymbol{Z}]\boldsymbol{u} \geq O_p(1)||u||^2 D;$$

and

$$I_2 = -\frac{1}{\sqrt{n}}\boldsymbol{u}^T\boldsymbol{Z}^T\hat{\boldsymbol{V}}^{-1}\boldsymbol{\epsilon} = O_p(1)||u||\sqrt{D}.$$

Hence I_1 dominates I_2 uniformly in $||\boldsymbol{u}|| = C$ by choosing a sufficiently large C.

Moreover,

$$\begin{aligned}
I_3 &= nD\sum_{j=1}^{s}\Big[p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j + n^{-1/2}\boldsymbol{u}_j||) - p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)\Big] \\
&= nD\sum_{j=1}^{s}\Big[n^{-1/2}\frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial\boldsymbol{\theta}_j}^T\boldsymbol{u}_j + \frac{1}{2}n^{-1}\boldsymbol{u}_j^T\frac{\partial^2 p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial\boldsymbol{\theta}_j\partial\boldsymbol{\theta}_j^T}\boldsymbol{u}_j\{1+o(1)\}\Big],
\end{aligned}$$

which is bounded by

$$sD\sqrt{n}\max\{\frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial\boldsymbol{\theta}_j}^T : ||\boldsymbol{\theta}_j|| \neq 0\}||\boldsymbol{u}|| + sD\max\{\frac{\partial^2 p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial\boldsymbol{\theta}_j\partial\boldsymbol{\theta}_j^T} : ||\boldsymbol{\theta}_j|| \neq 0\}||\boldsymbol{u}||^2.$$

Since $\lambda_n \to 0$ as $n \to \infty$, $\max\{\frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_j}^T : ||\boldsymbol{\theta}_j|| \neq 0\} \to 0$ and $\max\{\frac{\partial^2 p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_j \partial \boldsymbol{\theta}_j^T} : ||\boldsymbol{\theta}_j|| \neq 0\} \to 0$ by the definition of the MCP function.

Hence, by choosing a sufficiently large C, I_1 dominates all other terms when n is large enough. Then we have $Q(\boldsymbol{\theta} + n^{-1/2}\boldsymbol{u}) > Q(\boldsymbol{\theta})$ with arbitrary large probability $1 - \epsilon$, i.e., Inequality A.1 holds. Based on the discussions above, there exists a local minimizer $\hat{\boldsymbol{\theta}}$ such that

$$||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}|| = O_p(n^{-1/2}).$$

□

Proof of Theorem 2:

The partial derivative of $Q(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}_j$ is

$$\begin{aligned}\frac{\partial Q(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}_j} &= -\sum_{i=1}^{n} \boldsymbol{z}_{ij}^T \hat{\boldsymbol{\Sigma}}^{-1}(\boldsymbol{w}_i - \boldsymbol{z}_i\boldsymbol{\theta}) + nD\frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_j} \\ &\triangleq D_1 + D_2\end{aligned}$$

Since $\boldsymbol{\epsilon}_i = \boldsymbol{w}_i - \boldsymbol{z}_i\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and $\hat{\boldsymbol{\Sigma}} \to \boldsymbol{\Sigma}$, we have $D_1 = O_p(\sqrt{n})$.

Furthermore,

$$D_2 = nD(\frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \theta_{j1}}, ..., \frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \theta_{jK}})^T = nD(p'_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)\frac{\theta_{j1}}{||\boldsymbol{\theta}_j||}, ..., p'_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)\frac{\theta_{jK}}{||\boldsymbol{\theta}_j||})^T.$$

Since $\sqrt{n}\lambda_n \to \infty$ as $n \to \infty$, there exists an δ such that $C\sqrt{n} \leq \delta < \lambda_n$ for sufficient large n. When $||\boldsymbol{\theta}_j|| < \delta$, we have $p'_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||) = \lambda_n$. Hence, D_2 dominates D_1 and determines the sign the each element of $\frac{\partial Q(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}_j}$ when n is large enough.

This means as $n \to \infty$, with probability tending to 1, for any $\boldsymbol{\theta}^*$ satisfying $||\boldsymbol{\theta}^* - \boldsymbol{\theta}|| \leq C\sqrt{n}$ and constant C, there exists $\delta \geq C\sqrt{n} > 0$ such that

$$\frac{\partial Q(\boldsymbol{\theta})}{\partial \theta_{jk}} > 0, \text{ for } 0 < \theta_{jk} \leq ||\boldsymbol{\theta}_j|| < \delta,$$

$$\frac{\partial Q(\boldsymbol{\theta})}{\partial \theta_{jk}} < 0, \text{ for } -\delta < -||\boldsymbol{\theta}_j|| \leq \theta_{jk} < 0,$$

which implies

$$Q((\boldsymbol{\theta}_+^T, \mathbf{0})^T) = \inf_{\boldsymbol{\theta}^*: ||\boldsymbol{\theta}^* - \boldsymbol{\theta}|| \leq C\sqrt{n}} Q((\boldsymbol{\theta}_+^T, \boldsymbol{\theta}_0^T)^T).$$

In other words, with probability approaching 1, $Q(\boldsymbol{\theta})$ reaches its minimum when $\boldsymbol{\theta}_j = \mathbf{0}, j > s$. In the proof of Theorem 1, we have shown that there exists a constant C such that $||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}|| < C\sqrt{n}$. Therefore, $\hat{\boldsymbol{\theta}}_0 = \mathbf{0}$ with probability approaching 1.

Next we will prove the asymptotic normality of $\hat{\boldsymbol{\theta}}_+$. $\hat{\boldsymbol{\theta}}_+$ is the minimizer of $Q(\boldsymbol{\theta})|_{\boldsymbol{\theta}_0 = \mathbf{0}}$, i.e. $\frac{\partial Q(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}_+}|_{\boldsymbol{\theta} = (\hat{\boldsymbol{\theta}}_+^T, \mathbf{0})^T} = 0$.

By expanding the equation above, we have

$$\begin{aligned} 0 &= -\boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1}(\boldsymbol{W} - \boldsymbol{Z}_+ \hat{\boldsymbol{\theta}}_+) + nD \sum_{j=1}^{s} \frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_+}|_{\boldsymbol{\theta}_+ = \hat{\boldsymbol{\theta}}_+} \\ &= -\boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1}(\boldsymbol{W} - \boldsymbol{Z}_+ \boldsymbol{\theta}_+) + \boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1} \boldsymbol{Z}_+ (\hat{\boldsymbol{\theta}}_+ - \boldsymbol{\theta}_+) + nD \sum_{j=1}^{s} \frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_+}|_{\boldsymbol{\theta}_+ = \hat{\boldsymbol{\theta}}_+} \end{aligned}$$

where

$$nD \sum_{j=1}^{s} \frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_+}|_{\boldsymbol{\theta}_+ = \hat{\boldsymbol{\theta}}_+} = nD \left(\sum_{j=1}^{s} \frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_+}|_{\boldsymbol{\theta}_+ = \hat{\boldsymbol{\theta}}_+} + \sum_{j=1}^{s} \frac{\partial^2 p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_+ \partial \boldsymbol{\theta}_+^T}|_{\boldsymbol{\theta}_+ = \hat{\boldsymbol{\theta}}_+} [\hat{\boldsymbol{\theta}}_+ - \boldsymbol{\theta}_+] + o_p(1) \right).$$

Since $\lambda_n \to 0$ as $n \to \infty$, we have $||\boldsymbol{\theta}_j|| > \gamma\lambda_n$, $j < s$ when n is large enough. Then for sufficiently large n, $p'_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||) = 0$ and $p''_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||) = 0$, which indicates $nD \sum_{j=1}^{s} \frac{\partial p_{\lambda_n,\gamma}(||\boldsymbol{\theta}_j||)}{\partial \boldsymbol{\theta}_+}|_{\boldsymbol{\theta}_+ = \hat{\boldsymbol{\theta}}_+}$ is negligible.

Therefore, rearranging the equation above, we have

$$\begin{aligned} \frac{1}{\sqrt{nD}} \boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1} \boldsymbol{Z}_+ (\hat{\boldsymbol{\theta}}_+ - \boldsymbol{\theta}_+) &= \frac{1}{\sqrt{nD}} \boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1} (\boldsymbol{W} - \boldsymbol{Z}_+ \boldsymbol{\theta}_+), \\ (\frac{1}{nD} \boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1} \boldsymbol{Z}_+) \sqrt{m} (\hat{\boldsymbol{\theta}}_+ - \boldsymbol{\theta}_+) &= \frac{1}{\sqrt{nD}} \boldsymbol{Z}_+^T \hat{\boldsymbol{V}}^{-1} (\boldsymbol{W} - \boldsymbol{Z}_+ \boldsymbol{\theta}_+). \end{aligned}$$

Since $\hat{\boldsymbol{V}} \to \boldsymbol{V}$, by central limit theorem and Slutsky's theorem, we have

$$\hat{\boldsymbol{V}}^{-1} (\boldsymbol{W} - \boldsymbol{Z}_+ \boldsymbol{\theta}_+) \xrightarrow{D} \mathcal{N}(\mathbf{0}, \boldsymbol{V}^{-1}).$$

Additionally, since $\lim_{n\to\infty} \frac{1}{nD}\boldsymbol{Z}^T\boldsymbol{Z}$ is positive definite,

$$\frac{1}{\sqrt{nD}}\boldsymbol{Z}_+^T\hat{\boldsymbol{V}}^{-1}(\boldsymbol{W}-\boldsymbol{Z}_+\boldsymbol{\theta}_+) \xrightarrow{D} \mathcal{N}\left(\mathbf{0}, \lim_{n\to\infty}\frac{1}{nD}\boldsymbol{Z}_+^T\boldsymbol{V}^{-1}\boldsymbol{Z}_+\right).$$

Hence,

$$\sqrt{nD}(\hat{\boldsymbol{\theta}}_+ - \boldsymbol{\theta}_+) \xrightarrow{D} \mathcal{N}\left(\mathbf{0}, \left(\lim_{n\to\infty}\frac{1}{nD}\boldsymbol{Z}_+^T\boldsymbol{V}^{-1}\boldsymbol{Z}_+\right)^{-1}\right),$$

i.e.,

$$\sqrt{n}(\hat{\boldsymbol{\theta}}_+ - \boldsymbol{\theta}_+) \xrightarrow{D} \mathcal{N}\left(\mathbf{0}, \left(\lim_{n\to\infty}\frac{1}{n}\boldsymbol{Z}_+^T\boldsymbol{V}^{-1}\boldsymbol{Z}_+\right)^{-1}\right).$$

□

A.2 Simulation results for data with i.i.d. errors

Similar simulation studies were conducted for the case that within-function errors are uncorrelated. Datasets were constructed in the same way as what we did for the correlated case in Section 2.3 except that the errors were generated from a Gaussian distribution with covariance $\boldsymbol{\Sigma} = \boldsymbol{I}$. Likewise, we implemented four of our proposed method: one-step approaches using raw and FPCA-based covariance matrix estimates, and iterative approaches using raw and FPCA-based covariance matrix estimates. Two different values for PVE, 0.5 and 0.99, were used in the approaches involving FPCA-based covariance matrix estimate. In addition, the approach that pre-whitens using true covariance matrix, as well as ordinary least squares, the variational Bayes approach, FS-LASSO and a method that assumes uncorrelated error curves, were included for comparison.

Table A.1 reports the true positive (TP) and true negative (TN) rates of the estimates of both zero and non-zero coefficient functions. Our iterative approaches using FPCA-based covariance matrix estimate outperform all competing approaches in terms of correctly identifying the zero functions. In this case, there is no sign of a substantial decline in performance for the approaches using PVE=0.5 compared with the ones using PVE=0.99. Most methods are capable of identifying $\beta_1(t)$, $\beta_2(t)$ and $\beta_3(t)$ as non-zero functions.

Estimates of zero and non-zero coefficient functions obtained using the iterative algorithm with FPCA-based covariance matrix estimate using PVE=0.99, as well as the mean squared error and squared bias are shown in Figure A.1. As indicated in Table A.1, none of the estimates of $\beta_1(\cdot)$, $\beta_2(\cdot)$ and $\beta_3(\cdot)$ are set equal to zero. Driven by the sinusoidal shape of the coefficient function itself and by the shrinkage to zero, the squared bias curve of $\beta_1(\cdot)$ presents a sinusoidal shape. On the contrary, the

estimate of $\beta_2(\cdot)$ and $\beta_3(\cdot)$ are approximately unbiased owing to the structure of the penalty since their coefficients are relatively large. Due to the increased variability of curves at both ends of the distribution, large MSE is observed at both ends of the curves for all three non-zero coefficient functions.

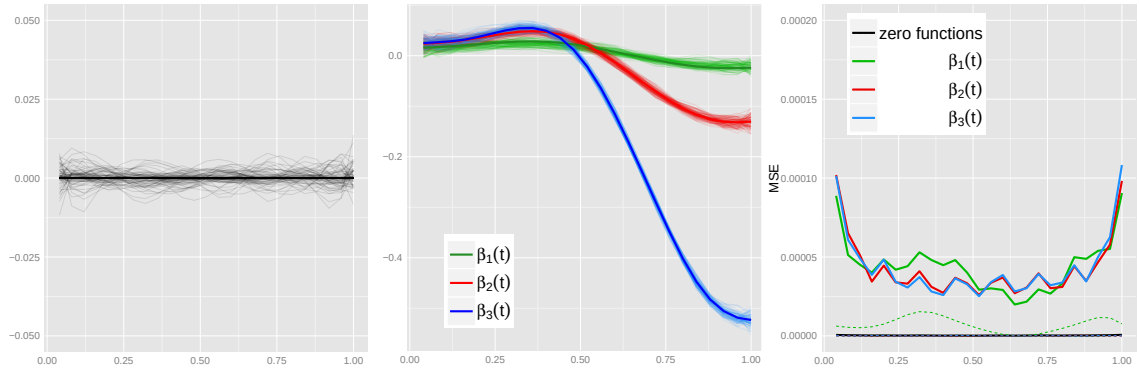

Figure A.1: Estimates of zero functions (*left*) and non-zero functions (*middle*) obtained using the iterative approach with FPCA-based covariance matrix estimate using PVE=0.99 across all simulated datasets. The true functions are overlaid (bold curves). The right panel shows the both MSE (solid) and squared bias (dashed) as functions of time for all the coefficient functions.

RMISE for estimated zero and non-zero functions, as well as the prediction errors on the test sample are presented in Figure A.2. Comparisons based on PVE=0.99 and PVE=0.5 are shown in the top and bottom rows, respectively. Our iterative algorithm with FPCA-based covariance matrix estimate, in particular when PVE=0.99, compares favorably to other approaches, reinforcing the results from Table A.1. In terms of both RMISE and prediction error, it is comparable to the method assuming independent error and the one that pre-whitens using true $\mathbf{\Sigma}$.

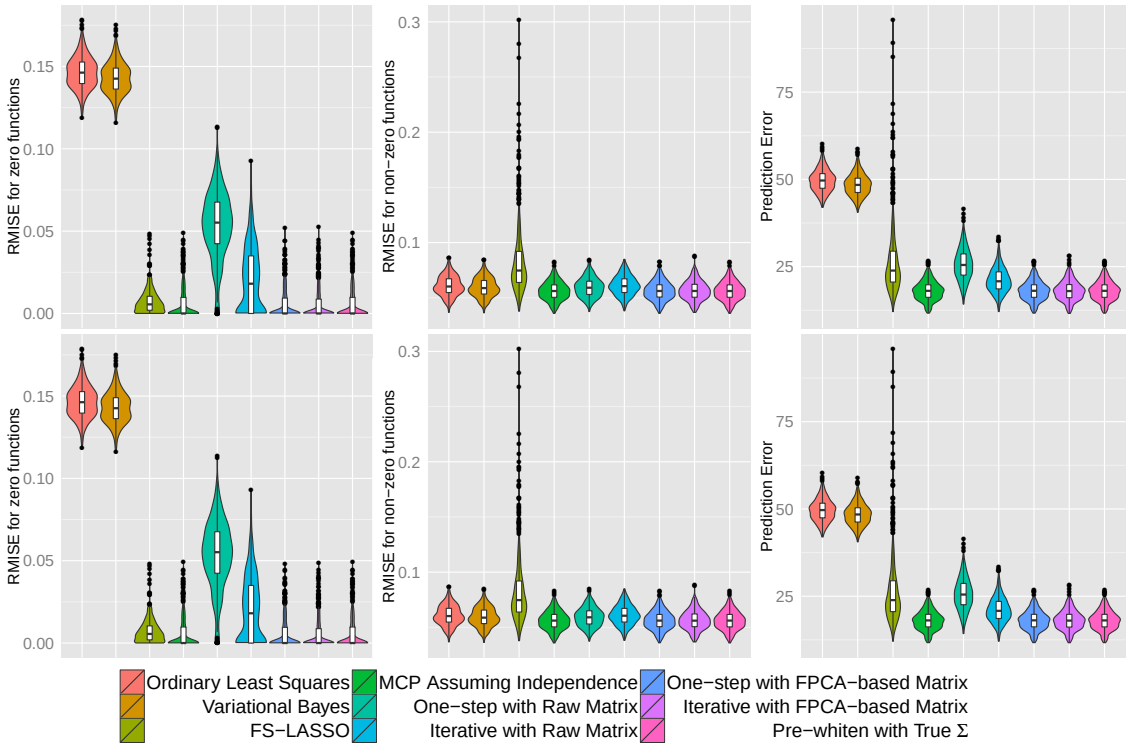

Figure A.2: The top row shows the comparison among the algorithms when PVE = 0.99 while the second row shows the comparison when PVE = 0.5. The three columns show RMISE for zero functions (*left*) and non-zero functions (*middle*); and prediction error (*right*).

Table A.1: true positive (TP) and true negative (TN) rates of estimated coefficient functions, where FN is false negative and FP is false positive. They are estimated across all the training samples.

	$\frac{TN}{FP+TN}$	$\frac{TP}{TP+FN}(\beta_1)$	$\frac{TP}{TP+FN}(\beta_2)$	$\frac{TP}{TP+FN}(\beta_3)$
FS-LASSO	0.908	0.990	1.000	1.000
MCP Assuming Independent Errors	0.943	1.000	1.000	1.000
One-step with Raw Matrix	0.371	1.000	1.000	1.000
Iterative with Raw Matrix	0.822	1.000	1.000	1.000
One-step with FPCA-based Matrix (PVE = 0.5)	0.944	1.000	1.000	1.000
Iterative with FPCA-based Matrix (PVE = 0.5)	0.950	1.000	1.000	1.000
One-step with FPCA-based Matrix (PVE = 0.99)	0.945	1.000	1.000	1.000
Iterative with FPCA-based Matrix (PVE = 0.99)	0.951	1.000	1.000	1.000
Pre-whiten with True $\mathbf{\Sigma}$	0.943	1.000	1.000	1.000

Appendix B

Appendices to: Nonlinear Mixed-Effects Models for PET Data

B.1 Forms of nonlinear models

The exact forms of f in Model 4.4 are shown as follows. The input function C_P and the whole blood function C_B have the same forms as in Parsey et al. (2000). $\mathbb{1}$ is the indicator function. f is based on an analytic convolution of the functions, while other approaches just involve in numerical convolutions.

B.1.1 One-tissue compartment (1TC) model

$$
\begin{aligned}
C_P(t) &= \mathbb{1}_{[0,\,t_c)}(t)bt + \mathbb{1}_{[t_c,\,\infty)}(t)\sum_{i=1}^{3} A_i \exp(-\lambda_i t) \\
C_B(t) &= \mathbb{1}_{[0,\,t'_c)}(t)b't + \mathbb{1}_{[t'_c,\,\infty)}(t)\sum_{i=1}^{3} A'_i \exp(-\lambda'_i t) \\
R_1 &= k_2 \\
L_1 &= k_1 \\
f(t) &= (1-V_b)\Bigg[\mathbb{1}_{[0,\,t_c)}(t)\left(bL_1\left(\frac{\exp(-R_1 t)}{R_1^2} + \frac{t}{R_1} - \frac{1}{R_1^2}\right)\right) \\
&\quad + \mathbb{1}_{[t_c,\,\infty)}(t)\Bigg(bL_1\exp(-R_1 t)\left(\frac{t_c}{R_1}\exp(R_1 t_c) - \frac{1}{R_1^2}\exp(R_1 t_c) + \frac{1}{R_1^2}\right) \\
&\quad + \sum_{i=1}^{3} L_1 A_i\left(\frac{\exp(-\lambda_1 t)}{R_1-\lambda_i} - \frac{R_1 t_c - R_1 t - \lambda_i t_c}{R_1-\lambda_i}\right)\Bigg)\Bigg] \\
&\quad + V_b\left(\mathbb{1}_{[0,\,t'_c)}(t)b't + \mathbb{1}_{[t'_c,\,\infty)}(t)\sum_{i=1}^{3} A'_i \exp(-\lambda'_i t)\right)
\end{aligned}
$$

B.1.2 Two-tissue compartment (2TC) model

$$
\begin{aligned}
C_P(t) &= \mathbb{1}_{[0,\,t_c)}(t)bt + \mathbb{1}_{[t_c,\,\infty)}(t)\sum_{i=1}^{3} A_i \exp(-\lambda_i t) \\
C_B(t) &= \mathbb{1}_{[0,\,t'_c)}(t)b't + \mathbb{1}_{[t'_c,\,\infty)}(t)\sum_{i=1}^{3} A'_i \exp(-\lambda'_i t) \\
R_1 &= \frac{1}{2}(k_2 + k_3 + k_4 + \sqrt{(k_2+k_3+k_4)^2 - 4k_2k_4}) \\
R_2 &= \frac{1}{2}(k_2 + k_3 + k_4 - \sqrt{(k_2+k_3+k_4)^2 - 4k_2k_4}) \\
L_1 &= \frac{k_1(R_1 - k_3 - k_4)}{R_1 - R_2} \\
L_2 &= \frac{k_1(k_3 + k_4 - R_2)}{R_1 - R_2} \\
f(t) &= (1 - V_b)\Bigg[\mathbb{1}_{[0,\,t_c)}(t)\left(bL_1\left(\frac{\exp(-R_1t)}{R_1^2} + \frac{t}{R_1} - \frac{1}{R_1^2}\right) + bL_2\left(\frac{\exp(-R_2t)}{R_2^2} + \frac{t}{R_2} - \frac{1}{R_2^2}\right)\right) \\
&\quad + \mathbb{1}_{[t_c,\,\infty)}(t)\Bigg(bL_1\exp(-R_1t)\left(\frac{t_c}{R_1}\exp(R_1t_c) - \frac{1}{R_1^2}\exp(R_1t_c) + \frac{1}{R_1^2}\right) \\
&\quad + bL_2\exp(-R_2t)\left(\frac{t_c}{R_2}\exp(R_2t_c) - \frac{1}{R_2^2}\exp(R_2t_c) + \frac{1}{R_2^2}\right) \\
&\quad + \sum_{i=1}^{3} L_1A_i\left(\frac{\exp(-\lambda_1t)}{R_1 - \lambda_i} - \frac{R_1t_c - R_1t - \lambda_it_c}{R_1 - \lambda_i}\right) \\
&\quad + \sum_{i=1}^{3} L_2A_i\left(\frac{\exp(-\lambda_1t)}{R_2 - \lambda_i} - \frac{R_2t_c - R_2t - \lambda_it_c}{R_2 - \lambda_i}\right)\Bigg)\Bigg] \\
&\quad + V_b\left(\mathbb{1}_{[0,\,t'_c)}(t)b't + \mathbb{1}_{[t'_c,\,\infty)}(t)\sum_{i=1}^{3} A'_i\exp(-\lambda'_it)\right)
\end{aligned}
$$

B.2 Parameter values used to simulate data

In Section 4.3, the individual parameters are generated from the following multivariate normal distribution:

$$\begin{pmatrix} k_1 \\ k_2 \\ k_3 \\ k_4 \\ V_b \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0.0565 \\ 0.1935 \\ 0.0510 \\ 0.0200 \\ 0.0234 \end{pmatrix}, \begin{pmatrix} 2.438 \times 10^{-4} & -1.078 \times 10^{-4} & -2.855 \times 10^{-5} & -1.278 \times 10^{-5} & 2.048 \times 10^{-5} \\ -1.078 \times 10^{-4} & 2.963 \times 10^{-4} & -2.155 \times 10^{-5} & 9.736 \times 10^{-6} & -3.130 \times 10^{-6} \\ -2.855 \times 10^{-5} & -2.155 \times 10^{-5} & 1.036 \times 10^{-5} & 1.921 \times 10^{-6} & 3.104 \times 10^{-6} \\ -1.278 \times 10^{-5} & 9.736 \times 10^{-6} & 1.921 \times 10^{-6} & 4.921 \times 10^{-6} & -4.755 \times 10^{-6} \\ 2.048 \times 10^{-5} & -3.130 \times 10^{-6} & 3.104 \times 10^{-6} & -4.755 \times 10^{-6} & 3.037 \times 10^{-5} \end{pmatrix} \right).$$