



Bayesian analysis of longitudinal and multidimensional functional data

JOHN SHAMSHOIAN, DAMLA ŞENTÜRK

Department of Biostatistics, University of California, Los Angeles, CA, USA

SHAFALI JESTE

Department of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles, CA, USA

DONATELLO TELESCA*

Department of Biostatistics, University of California, Los Angeles, CA, USA

donatello.telesca@gmail.com

SUMMARY

Multi-dimensional functional data arises in numerous modern scientific experimental and observational studies. In this article, we focus on longitudinal functional data, a structured form of multidimensional functional data. Operating within a longitudinal functional framework we aim to capture low dimensional interpretable features. We propose a computationally efficient nonparametric Bayesian method to simultaneously smooth observed data, estimate conditional functional means and functional covariance surfaces. Statistical inference is based on Monte Carlo samples from the posterior measure through adaptive blocked Gibbs sampling. Several operative characteristics associated with the proposed modeling framework are assessed comparatively in a simulated environment. We illustrate the application of our work in two case studies. The first case study involves age-specific fertility collected over time for various countries. The second case study is an implicit learning experiment in children with autism spectrum disorder.

Keywords: Factor analysis; Functional data analysis; Gaussian process; Longitudinal mixed model; Marginal covariance; Rank regularization; Tensor spline.

1. INTRODUCTION

We investigate Bayesian modeling and inference for longitudinal functional data, conceptualized as functional data observed repeatedly over a dense set of longitudinal time points. A typical dataset would contain n patients observed over the course of multiple visit times, with each visit contributing a functional datum. Thus, for patient i we would record the outcome $y_i(s, t)$, where s is the visit time and t is the functional argument. In this setting, it is reasonable to expect non-trivial correlations between functions from one visit time to another. Therefore, appropriate modeling of this dependence pattern would be critical for

*To whom correspondence should be addressed.

the validity of statistical inference. This manuscript outlines a flexible Bayesian framework for the estimation of the functional mean structure, possibly dependent on a set of time-stable covariates, as well as an adaptive regularization framework for the estimation of the covariance operator of $y_i(s, t)$ and its eigenstructure.

The frequentist analysis of longitudinal functional data is a mature field. In particular, semiparametric modeling strategies, depending on the mixed effects modeling framework, have been proposed by [Di and others \(2009\)](#) in the context of hierarchical functional data, and [Greven and others \(2010\)](#) for longitudinal functional data. Important generalizations have been introduced by [Chen and Müller \(2012\)](#), through structured functional principal components analysis (FPCA), with more parsimonious representations introduced by [Park and Staicu \(2015\)](#), and [Chen and others \(2017\)](#) in the more general context of function-valued stochastic processes. The appealing nature and flexibility of structured FPCA modeling strategies has seen the application and extension of these methods to challenging scientific problems ranging from functional brain imaging ([Hasenstab and others, 2017](#); [Scheffler and others, 2020](#)), to the exploration of complex data from wearable devices ([Goldsmith and others, 2016](#)). The vast majority of approaches based on FPCA, generally focus on point estimation from a frequentist perspective, and do not provide reliable uncertainty quantification without bootstrapping. The very application of the bootstrap methodology to structured functional data has not been the subject of rigorous investigation. The literature, in fact, is ambiguous on the handling of the many tuning parameters, typical of structured FPCA models. Although there are some consistency results regarding the bootstrap for functional data ([Cuevas and others, 2006](#); [Ferraty and others, 2010](#)), the procedure is relatively underdeveloped for hierarchical data ([Ren and others, 2010](#)). Bayesian methods in functional data analysis define a straightforward mechanism for uncertainty quantification. This appealing inferential structure comes, however, at the cost of having to specify a full probability model, and priors with broad support on high-dimensional spaces ([Shi and Choi, 2011](#); [Yang and others, 2016, 2017](#)). In hierarchical and multi-dimensional functional data settings, starting from the seminal work of [Morris and others \(2003\)](#), and recent extensions in [Lee and others \(2019\)](#), the prevalent strategy has been to work within the framework of basis function transforms, defining flexible mixed effect models at the level of the basis coefficients ([Morris and others, 2003](#); [Baladandayuthapani and others, 2008](#); [Zhang and others, 2016](#)). The resulting functional mixed effects models, like their finite dimensional counterpart, require a certain degree of subject matter expertise in the definition of random effects and their covariance structure ([Morris and others, 2011](#); [Morris and Carroll, 2006](#)).

This manuscript aims to merge the appealing characterization of longitudinal functional data through FPCA decompositions ([Chen and Müller, 2012](#); [Park and Staicu, 2015](#); [Chen and others, 2017](#)), with flexible probabilistic representations of the classical Karhunen–Loève expansion of square integrable random functions. Our work builds on the ideas of [Suarez and Ghosal \(2017\)](#) and [Montagna and others \(2012\)](#), who adapted the regularized product Gamma prior for infinite factor models of [Bhattacharya and Dunson \(2011\)](#), to the analysis of random functions. Extensions of this framework to the longitudinal functional setting are discussed in Section 2. In Section 3, we discuss prior distributions and ensuing implications for the covariance operator. A comprehensive framework for posterior inference is discussed in Section 4. Section 5 contains a comparative simulation study. Finally, in Section 6 we discuss the application of our proposed methodology to two case studies. The first case study explores age-specific fertility dynamics in the global demographic study conducted by the Max Plank Institute and the Vienna Institute of Demography ([HFD, 2019](#)). While purely illustrative, these data allow for a direct comparison with the original analysis of [Chen and others \(2017\)](#). The second case study, involves the analysis of electroencephalogram (EEG) data from an investigation of implicit learning in children with autism spectrum disorder (ASD) ([Jeste and others, 2015](#)). The main interest in both case studies is modeling and interpreting the longitudinal component.

2. A PROBABILITY MODEL FOR LONGITUDINAL FUNCTIONAL DATA

Let $y_i(s, t)$ denote the response for subject i , ($i = 1, \dots, n$), at longitudinal time $s \in \mathcal{S}$ and functional time $t \in \mathcal{T}$, where \mathcal{S} and \mathcal{T} are compact subspaces of \mathbb{R} . Furthermore, for each subject, assume we observe a time-stable d -dimensional covariate $\mathbf{x}_i \in \mathbb{R}^d$. In practice, we only obtain observations $y_i(s_j, t_k)$ at discrete sampling locations $(s_j, t_k) \in \mathcal{S} \times \mathcal{T}$, $j = 1, \dots, n_j^s$, $k = 1, \dots, n_k^t$. However, in subsequent developments, we maintain the lighter notation $y_i(s, t)$ without loss of generality.

Let $f_i(s, t)$ be a Gaussian Process with mean $E\{f_i(s, t)\} = \mu(\mathbf{x}_i, s, t)$ and covariance kernel $\text{Cov}\{f_i(s, t), f_i(s', t')\} = K\{(s, t), (s', t')\}$. A familiar sampling model for $y_i(s, t)$ assumes:

$$y_i(s, t) = f_i(s, t) + \epsilon_i(s, t), \quad \epsilon_i(s, t) \stackrel{iid}{\sim} N(0, \varphi^2); \quad (2.1)$$

where $\varphi^2 > 0$ is the overall residual variance. Given a set of suitable basis functions $b_m^{(1)}(s) : \mathcal{S} \rightarrow \mathbb{R}$, ($m = 1, 2, \dots, p_1$), and $b_\ell^{(2)}(t) : \mathcal{T} \rightarrow \mathbb{R}$, ($\ell = 1, 2, \dots, p_2$), and a set of random coefficients $\theta_{im\ell}$, the prior for the underlying signal $f_i(s, t)$ is constructed through a random tensor product expansion, so that

$$f_i(s, t) = \sum_{m=1}^{p_1} \sum_{\ell=1}^{p_2} b_m^{(1)}(s) b_\ell^{(2)}(t) \theta_{im\ell}. \quad (2.2)$$

Since the truncation values p_1 and p_2 may be large to ensure small bias in the estimation of the true $f_i(s, t)$, we follow [Bhattacharya and Dunson \(2011\)](#) and project the basis coefficients on a lower-dimensional space.

Let $\Theta_i = \{\theta_{im\ell}\} \in \mathbb{R}^{p_1 \times p_2}$ be the matrix of basis coefficients for subject i . After defining loading matrices $\Lambda \in \mathbb{R}^{p_1 \times q_1}$, ($q_1 \ll p_1$), and $\Gamma \in \mathbb{R}^{p_2 \times q_2}$, ($q_2 \ll p_2$), and a latent matrix of random scores $\boldsymbol{\eta}_i \in \mathbb{R}^{q_1 \times q_2}$, we assume

$$\Theta_i = \Lambda \boldsymbol{\eta}_i \Gamma^\top + \boldsymbol{\zeta}_i, \quad \text{vec}(\boldsymbol{\zeta}_i) \sim \mathcal{N}(0, \Sigma), \quad (2.3)$$

where Σ is taken to be diagonal. The foregoing construction has connections with factor analysis. In fact, vectorizing Θ_i we obtain

$$\text{vec}(\Theta_i) = (\Gamma \otimes \Lambda) \text{vec}(\boldsymbol{\eta}_i) + \text{vec}(\boldsymbol{\zeta}_i);$$

which resembles the familiar $(q_1 \times q_2)$ latent factor model, with loading matrix $\Gamma \otimes \Lambda$ and latent factors $\text{vec}(\boldsymbol{\eta}_i)$. Differently from standard latent factor models, our use of a Kronecker product representation for the loading matrix introduces additional structural assumptions about $\text{Cov}(\Theta_i)$, and the ensuing form of the covariance kernel $K\{(s, t), (s', t')\}$.

More precisely, assuming $\text{Cov}(\boldsymbol{\eta}_i) = H$, the marginal covariance of Θ_i takes the form

$$\text{Cov}(\Theta_i) = (\Gamma \otimes \Lambda) H (\Gamma \otimes \Lambda)^\top + \Sigma = \Omega. \quad (2.4)$$

Furthermore, defining $B_1(s) = \{b_1^{(1)}(s), \dots, b_{p_1}^{(1)}(s)\}^\top$ and $B_2(t) = \{b_1^{(2)}(t), \dots, b_{p_2}^{(2)}(t)\}^\top$, the model in (2.3) induces the following representation for the covariance kernel $K\{(s, t), (s', t')\}$, s.t.

$$K\{(s, t), (s', t')\} = \{B_1(s) \otimes B_2(t)\} \Omega \{B_1(s') \otimes B_2(t')\}^\top. \quad (2.5)$$

The low-rank structure of Ω in (2.4), depends on the number of latent factors q_1 and q_2 in the quadratic form $(\Gamma \otimes \Lambda)^\top H (\Gamma \otimes \Lambda)$. Rather than selecting the number of factors a priori, in Section 3, we introduce prior distributions encoding rank restrictions through continuous stochastic regularization of the loading

coefficient’s magnitude. Additional structural restrictions may ensue from specific assumptions about the latent factors’ covariance H . Specifically, setting $H = I_{q_1 q_2}$ leads to strong covariance separability of the longitudinal and functional dimensions. A more flexible covariance model hinges on the notion of weak-separability (Lynch and Chen, 2018). This is achieved by assuming $H = \text{diag}(h_1, \dots, h_{q_1 q_2}) > 0$.

Finally, let \mathbf{x}_i be a d -dimensional time-stable covariate for subject i . Dependence of the longitudinal functional outcome $y_i(s, t)$ on this set of predictors is conveniently introduced through the prior expectation of $\boldsymbol{\eta}_i$. More precisely, let $\boldsymbol{\beta}$ be a $d \times q_1 q_2$ matrix of regression coefficients, we assume

$$\text{vec}(\boldsymbol{\eta}_i) \sim N(\boldsymbol{\beta}^\top \mathbf{x}_i, H),$$

which implies the following marginal mean structure for $y_i(s, t)$,

$$E\{y_i(s, t)\} = \mu(\mathbf{x}_i, s, t) = \{B_1(s)\Gamma \otimes B_2(t)\Lambda\} \boldsymbol{\beta}^\top \mathbf{x}_i. \tag{2.6}$$

The model in (2.1), together with the sandwich factor construction in (2.3) defines a probabilistic representation of the product FPCA decomposition in Chen and others (2017). An intuitive parallel is introduced in Section 3, and a technical discussion is provided in the accompanying web-based supplementary document (Appendix A of the supplementary material available at *Biostatistics* online). Differently from Chen and others (2017), we propose model-based inference through regularized estimation based on the posterior measure.

3. RANK REGULARIZATION AND PRIOR DISTRIBUTIONS

The selection of prior distributions for all parameters introduced in Section 2 is guided by the following considerations. Let γ_{lj} and λ_{mk} be specific entries in the loading matrices Γ and Λ in (2.3), respectively. Defining $\psi_j(s) = \sum_{l=1}^{p_1} \gamma_{lj} b_l^{(1)}(s)$ and $\phi_k(t) = \sum_{m=1}^{p_2} \lambda_{mk} b_m^{(2)}(t)$, we may expand $f_i(s, t)$ as follows:

$$f_i(s, t) = \sum_{j=1}^{q_1} \sum_{k=1}^{q_2} \psi_j(s) \phi_k(t) \eta_{ijk} + r_i(s, t),$$

$$r_i(s, t) = \sum_{j=1}^{p_1} \sum_{k=1}^{p_2} b_j^{(1)}(s) b_k^{(2)}(t) \zeta_{ijk}.$$

The first component in the expression for $f_i(s, t)$ describes a mechanism of random functional variability which depends on the tensor combination of q_1 and q_2 data-adaptive basis functions $\psi_j(s)$ and $\phi_k(t)$, respectively, and $q_1 \times q_2$ basis coefficients η_{ijk} . Given q_1 and q_2 , any residual variability is represented in the random function $r_i(s, t)$. When $\psi_j(s)$ and $\phi_k(t)$ are chosen to be eigenfunctions of the marginal covariance kernels in s and t , this representation is essentially equivalent to the product FPCA construction of Chen and others (2017).

Statistical inference for FPCA constructions, commonly selects a small number of eigenfunctions on the basis of empirical considerations. Here, we take an adaptive regularization approach, choose q_1 and q_2 relatively large, and assume the variance components in the priors for Λ and Γ to follow a modified multiplicative gamma process prior (MGPP) Bhattacharya and Dunson (2011) Montagna and others (2012).

Let λ_{mk} denote the (m, k) entry of Λ . The modified MGPP is defined by setting

$$\begin{aligned}\lambda_{mk} &\sim N(0, \rho_{1mk}^{-1} \tau_{1k}^{-1}), \quad \rho_{1mk}^{-1} \sim \text{Gamma}(v_1/2, v_1/2), \\ \tau_{1k} &= \prod_{v=1}^k \delta_{1v}, \quad \delta_{11} \sim \text{Gamma}(a_{11}, 1), \\ \delta_{1v} &\sim \text{Gamma}(a_{12}, 1) \mathbb{1}(\delta_{1v} > 1), \text{ for } v \geq 2; \quad k = 1, 2, \dots, q_1.\end{aligned}\tag{3.7}$$

Using the ‘‘rate’’ parametrization for Gamma distributions (i.e., if $a \sim \text{Gamma}(b, c)$, then $E(a) = bc$), this prior is designed to encourage small loadings in Λ as the column index increases. In the original formulation of [Bhattacharya and Dunson \(2011\)](#) and [Montagna and others \(2012\)](#), choosing $a_{12} > 1$, insures stochastic ordering of the prior precision, in the sense that $E(\tau_{1k}) < E(\tau_{1(k+1)})$, for any $k = 1, 2, \dots, (q_1 - 1)$. In our setting, we require the more stringent probabilistic ordering $Pr(\tau_{1k} < \tau_{1(k+1)}) = 1$, by assuming $\delta_{1v} > 1$, which results in a more stable and efficient Gibbs sampling scheme. Analogous regularization over the columns of Γ is achieved by setting:

$$\begin{aligned}\gamma_{lj} &\sim N(0, \rho_{2lj}^{-1} \tau_{2l}^{-1}), \quad \rho_{2lj}^{-1} \sim \text{Gamma}(v_2/2, v_2/2) \\ \tau_{2l} &= \prod_{v=1}^l \delta_{2v}, \quad \delta_{21} \sim \text{Gamma}(a_{21}, 1), \\ \delta_{2v} &\sim \text{Gamma}(a_{22}, 1) \mathbb{1}(\delta_{2v} > 1), \text{ for } v \geq 2; \quad l = 1, 2, \dots, q_2.\end{aligned}\tag{3.8}$$

Adaptive shrinkage is induced by placing hyper-priors on a_{11}, a_{12}, a_{21} , and a_{22} , such that

$$a_{11}, a_{21} \overset{\text{ind}}{\sim} \text{Gamma}(r_1, 1), \quad a_{12}, a_{22} \overset{\text{ind}}{\sim} \text{Gamma}(r_2, 1).$$

The model is completed with priors on residual variance components and regression coefficients. Specifically, conditionally conjugate priors are placed on the diagonal elements of Σ and H , respectively, as well as the residual variance φ , such that:

$$\sigma_j^{-1} \sim \text{Gamma}(a_\sigma, b_\sigma), \quad h_j^{-1} \sim \text{Gamma}(a_h, b_h), \quad \varphi^{-1} \sim \text{Gamma}(a_\varphi, b_\varphi).$$

Finally, we induce a Cauchy prior for the regression coefficients matrix β as in [Montagna and others \(2012\)](#). Denoting with $\beta_{j\ell}$ the (j, ℓ) entry of β , we assume

$$\beta_{j\ell} \sim N(0, \omega_{j\ell}), \quad \omega_{j\ell}^{-1} \sim \text{Gamma}(1/2, 1/2); \quad \ell = 1, \dots, q_1 q_2, \quad j = 1, \dots, d.$$

In summary, our approach starts with the projection of the observed data onto a set of known basis functions in (2.2). This initial projection is similar to the interpolation or smoothing step commonly used in functional data analysis ([Chen and others, 2017](#); [Morris and Carroll, 2006](#)). The basis coefficients Θ_i are assumed to arise from the latent factor model in (2.3), resulting in the weakly separable covariance model in (2.4) and (2.5). Finally, the MGPP priors in (3.7) and (3.8), allow for adaptive regularization of the covariance operator. The mean structure is made dependent on a set of time stable covariates through a varying coefficient model in (2.6).

4. POSTERIOR INFERENCE

Posterior simulation through Markov chain Monte Carlo is relatively straightforward, after selection of an appropriate basis transform and truncation of Γ and Λ to include $q_1 \ll p_1$ and $q_2 \ll p_2$ columns respectively. The use of conditionally conjugate priors allows for simple Gibbs transitions for all parameters, with the exception of the shrinkage parameters a_{11}, a_{12}, a_{21} , and a_{22} , which are updated via a Metropolis–Hastings step. A detailed description of the proposed algorithm is reported in the web-based supplement (Appendix B of the supplementary material available at *Biostatistics* online).

We note that the decomposition of $\text{Cov}(\Theta_i)$ in (2.4) may not be unique. However, from a Bayesian perspective, one does not require identifiability of the loading elements for the purpose of covariance estimation. Direct inference for $K\{(s, t), (s', t')\}$ and its functionals may be achieved by post-processing Monte Carlo draws from the posterior $p(\Omega \mid \mathbf{y})$ and evaluating the covariance function over arbitrarily dense points $\mathbf{t}^* := (t_1^*, \dots, t_{w1}^*)^\top \in \mathcal{T}$ and $\mathbf{s}^* := (s_1^*, \dots, s_{w2}^*)^\top \in \mathcal{S}$ using (2.5). Analogously, given samples from $p(\boldsymbol{\beta} \mid \mathbf{y})$, inference about the mean structure is achieved by evaluating $\mu(\mathbf{x}_i, s, t)$ over \mathbf{s}^* and \mathbf{t}^* using the expansion in (2.6).

Some useful posterior summaries may be obtained through marginalization. We define marginal covariance functions $K_{\mathcal{T}}(t, t')$ and $K_{\mathcal{S}}(s, s')$ as follows:

$$K_{\mathcal{T}}(t, t') = \int_{\mathcal{S}} K\{(s, t)(s, t')\}ds, \quad K_{\mathcal{S}}(s, s') = \int_{\mathcal{T}} K\{(s, t)(s', t)\}dt. \tag{4.9}$$

Intuitively, $K_{\mathcal{S}}(\cdot)$ and $K_{\mathcal{T}}(\cdot)$ summarize patterns of functional co-variation along a specific coordinate, and their lower-dimensional posterior summaries may be obtained through functional eigenanalysis as in [Chen and others \(2017\)](#). We outline details on extracting lower-dimensional summaries of the marginal covariance functions without computing $K\{(s, t), (s', t')\}$, $K_{\mathcal{T}}(t, t')$, or $K_{\mathcal{S}}(s, s')$ in Appendix F of the of the supplementary material available at *Biostatistics* online. Simultaneous credible intervals for all functions of interest are easily obtained from Monte Carlo samples, by applying the methodology discussed in [Crainiceanu and others \(2007\)](#) and [Baladandayuthapani and others \(2005\)](#).

Specifically, M Monte Carlo draws from a posterior function of interest, say $g(\tau)$, are used to estimate the posterior mean $\widehat{g}(\tau)$, and standard deviation $\sqrt{\widehat{\text{var}}\{g(\tau)\}}$. Assuming approximate normality of the posterior distribution, we derive the $(1 - \alpha)$ quantile c_α of the pivotal quantity

$$\max_{\tau} \left| \frac{g^{(i)}(\tau) - \widehat{g}(\tau)}{\sqrt{\widehat{\text{var}}\{g(\tau)\}}} \right|, \quad i = 1, \dots, M.$$

An approximate simultaneous $(1 - \alpha)$ posterior band can then be constructed as a hyperrectangular region over τ : $\left[\widehat{g}(\tau) \pm c_\alpha \cdot \sqrt{\widehat{\text{var}}\{g(\tau)\}} \right]$. More general simultaneous bands have been proposed by [Krivobokova and others \(2010\)](#), but are not implemented in this manuscript.

The proposed modeling framework relies on a specific basis transform strategy. While the literature has suggested the use of zero-loss transforms as a default option ([Morris and others, 2003](#); [Lee and others, 2019](#)), we find that it is not uncommon to observe some sensitivity to the number of basis functions used in the initial projection. Furthermore, the choice of more parsimonious designs, when warranted by the application, may lead to important gains in computational and estimation efficiency.

Model flexibility is governed by choice of (p_1, p_2) , as the number of smoothing basis functions, and (q_1, q_2) as the number of latent factors. Due to the adaptive rank regularizing prior, q_1 and q_2 should be chosen as large as possible. In practice (q_1, q_2) are chosen as fraction of (p_1, p_2) .

Our simulation studies (Appendix C of the supplementary material available at *Biostatistics* online) demonstrate that point estimates and uncertainty of mean and covariance functions are generally insensitive

to choice of p_1 and p_2 , provided q_1 and q_2 are large. Some sensitivity is, however, observed in the posterior estimate of the residual error φ^2 . An alternative method is to simply rely on the minimization of information criteria. In this article, we consider simple versions of the deviance information criterion (DIC), and two versions of the Bayesian information criteria (BIC and BIC_h [Delattre and others \(2014\)](#)). Our simulation studies (Appendix C of the supplementary material available at *Biostatistics* online) indicate that the proposed information criteria perform well in selecting an adequate number of basis functions.

From a computational perspective, the most time consuming steps in the Gibbs sampling algorithm are the Cholesky decompositions used in updating Θ_i and η_i , requiring $O(p_1^3 p_2^3)$ and $O(q_1^3 q_2^3)$ floating point operations, respectively. Therefore, scalability of naïve Gibbs sampling is a potential issue for very large samples and/or very long longitudinal or functional evaluation domains. In these cases, adapting the estimation approach of [Morris and Carroll \(2006\)](#), is easily implemented, by treating the estimation of Θ_i as a pre-processing step, and considering 2.3 as the sampling model. For big data applications, other analytical approximations to the posterior measure are also accessible, e.g., Integrated Nested Laplace Approximations (INLA) ([Rue and others, 2009](#)).

5. A MONTE CARLO STUDY OF OPERATING CHARACTERISTICS

We performed a series of numerical experiments aimed at evaluating the estimation performance for both the functional mean and covariance. We study three simulation scenarios, including two weakly separable kernels (Cases 1 and 2) and one non-separable covariance function (Case 3). Specifically, for $s \in [0, 1]$ and $t \in [0, 1]$, we take:

- (1) $K_S(s, s') = \sum_{j=1}^2 \lambda_j \psi_j(s) \psi_j(s')$, with eigenvalues $\lambda_j = \frac{1}{j^2 \pi^2}$ and eigenfunctions $\psi_j(s) = \sqrt{2} \sin(j\pi s)$, $K_T(t, t') = \sigma^2 \left(1 + \frac{\sqrt{3}|t-t'|}{\rho}\right) \exp\left(-\frac{\sqrt{3}|t-t'|}{\rho}\right)$, in the Matérn class, and mean $\mu(s, t) = \sqrt{\frac{1}{5\sqrt{s+1}}} \sin(5t)$.
- (2) $K_S(s, s') = \sum_{j=1}^2 \lambda_j \psi_j(s) \psi_j(s')$, with eigenvalues $\lambda_j = \frac{1}{(j-1/2)^2 \pi^2}$ and eigenfunctions $\psi_j(s) = \sqrt{2} \sin((j-1/2)\pi s)$, $K_T(t, t') = \sum_{k=1}^{50} \lambda_k \phi_k(t) \phi_k(t')$, with $\lambda_k = k^{-2\alpha}$ and $\phi_k(t) = \cos(k\pi t)$, and mean $\mu(s, t) = 5\sqrt{1 - (s-.5)^2 - (t-.5)^2}$.
- (3) $K((s, t), (s', t')) = \frac{1}{(t-t')^2 + 1} \exp\left\{-\frac{(s-s')^2}{(t-t')^2 + 1}\right\}$, stationary non-separable ([Gneiting, 2002](#)), and mean $\mu(s, t) = \sqrt{1 + \sin(\pi s) + \cos(\pi t)}$.

Scenario 1 combines a simple Matérn class pattern on the time-domain t with a more complex oscillatory dependence pattern for the functional domain s . Covariance 2 includes an oscillatory pattern in both s and t . Finally, Covariance 3, while defining simple parametric dependence in both longitudinal and functional times, is not weakly separable, allowing for comparisons on misspecified models.

We consider estimation of the mean, covariance, marginal covariance functions, and the associated two principal eigenfunctions. Each simulation includes 1000 Monte Carlo experiments. For each experiment, posterior estimates are based on 10 000 iterations of four independent Markov chains, after discarding 2500 draws for burn-in. We compare estimation of covariance, marginal covariance functions, and associated two principal eigenfunctions to the respective estimates provided by the product FPCA ([Chen and others, 2017](#)), as well as finite-dimensional empirical estimates of the mean and covariance defined as by their vectorized sample counterparts. Estimates obtained with the product FPCA have data type set to sparse and fraction of variance explained threshold set to 0.9999.

Table 1. Mean and covariance relative errors under the three simulation cases described in section 5. Bayes refers to the proposed method in this paper, product refers to the product decomposition Chen and others (2017), and empirical refers to point-wise empirical estimation. Each case is repeated 1000 times for sample sizes of $n = 30$ and $n = 60$. We report the 50th percentile of the relative error, with the numbers in the parentheses denoting the 10th and 90th percentiles of the relative error.

Case 1		Bayes	Product	Empirical
$n = 30$	$\mu(s, t)$	0.014 (0.005, 0.038)	0.019 (0.010, 0.044)	0.019 (0.010, 0.044)
	$K\{(s, t), (s', t')\}$	0.062 (0.023, 0.224)	0.085 (0.047, 0.200)	0.151 (0.097, 0.297)
$n = 60$	$\mu(s, t)$	0.007 (0.003, 0.019)	0.010 (0.005, 0.021)	0.010 (0.005, 0.021)
	$K\{(s, t), (s', t')\}$	0.030 (0.010, 0.097)	0.057 (0.038, 0.128)	0.076 (0.050, 0.151)
Case 2				
$n = 30$	$\mu(s, t)$	0.024 (0.007, 0.101)	0.031 (0.013, 0.118)	0.031 (0.013, 0.118)
	$K\{(s, t), (s', t')\}$	0.039 (0.011, 0.184)	0.050 (0.012, 0.202)	0.067 (0.030, 0.228)
$n = 60$	$\mu(s, t)$	0.014 (0.004, 0.054)	0.017 (0.007, 0.062)	0.017 (0.007, 0.062)
	$K\{(s, t), (s', t')\}$	0.019 (0.005, 0.091)	0.024 (0.007, 0.093)	0.032 (0.014, 0.106)
Case 3				
$n = 30$	$\mu(s, t)$	0.155 (0.046, 0.389)	0.160 (0.051, 0.393)	0.160 (0.051, 0.393)
	$K\{(s, t), (s', t')\}$	0.051 (0.016, 0.187)	0.051 (0.014, 0.183)	0.067 (0.023, 0.200)
$n = 60$	$\mu(s, t)$	0.073 (0.019, 0.216)	0.076 (0.021, 0.219)	0.076 (0.021, 0.219)
	$K\{(s, t), (s', t')\}$	0.028 (0.008, 0.091)	0.027 (0.007, 0.089)	0.034 (0.011, 0.099)

All comparisons are based on the relative mean integrated squared error. For a function f with domain D and estimator \hat{f} , we define $\text{RE}(\hat{f}, f) = \int_D \{\hat{f}(u) - f(u)\}^2 du / \int_D f(u)^2 du$. Note that D can be multi-dimensional and in practice the integral is replaced with a sum. Table 1 compares mean $\mu(s, t)$ and covariance $K\{(s, t), (s', t')\}$ estimation under the three settings listed above. We find that estimates from each method improve in accuracy with increasing sample size ($n = 30, 60$), with the posterior and product FPCA showing greater accuracy than empirical approach in terms of covariance estimation. Similar findings characterize the estimation performance of all marginal covariance functions (K_S, K_T), and the associated two principal eigenfunctions ($\psi_i(s), i = 1, 2$), and ($\phi_i(t), i = 1, 2$). Detailed numerical results and extended simulations are reported in the web-based supplement (Appendix C of the supplementary material available at *Biostatistics* online).

In summary, we observe that posterior estimates are associated with similar, and potentially improved accuracy in the estimation of the mean and covariance functions, when compared with product FPCA. This similarity in estimation performance, provides some empirical assurances that the chosen probabilistic representation of structured covariance functions, and estimation based on adaptive shrinkage, maintains a data-adaptive behavior with good operating characteristics.

6. CASE STUDIES

We illustrate the application of the proposed modeling frameworks in two case studies. The first dataset concerns fertility rate and age of mothers by country. The second case study focuses on functional brain imaging through EEG in the context of implicit learning in children with ASD.

6.1. Fertility rates

The Human Fertility Database (HFD, 2019) compiles vital statistics to facilitate research on fertility in the past twentieth century and in the modern era. Age-specific fertility rates are available for 32 countries over different time periods. The age-specific fertility rate $ASFR(s, t)$ is defined as

$$ASFR(s, t) = \frac{\text{Births during year } s \text{ given } t \text{ by women aged } t}{\text{Person-years lived during year } s \text{ by women aged } t}.$$

The dataset was previously analyzed and interpreted in a longitudinal functional framework using the product FPCA (Chen and others, 2017). This section focuses on a comparative analysis of product FPCA and the proposed probability model.

We follow Chen and others (2017), and consider $n = 17$ countries, with complete data for the time period 1951 to 2006, 44 functional time points (ages 12–55), and 56 longitudinal time points (years 1951 to 2006). Since these rates are population measurements, we expect the data to contain very little noise.

We use cubic b-splines as our basis functions since the data look smooth with no sharp changes in fertility rate over year or age of mother (Figure S1 of the supplementary material available at *Biostatistics* online) and consider $(p_1, p_2) = (22, 28)$ splines and $(q_1, q_2) = (11, 10)$ latent factors, selected by minimizing the DIC.

Longitudinal and aging dynamics are largely determined by their associated marginal covariance functions $K_S(s, s')$ and $K_T(t, t')$. Figure 1 displays the first three marginal eigenfunctions for age and calendar year. We include the 95% simultaneous credible bands (Crainiceanu and others, 2007) as well as estimates obtained via product FPCA. We note that Bayesian posterior mean eigenfunctions are qualitatively similar to the inferred product FPCA estimates, therefore warranting similar interpretations to the one originally offered by Chen and others (2017).

In particular, the first marginal eigenfunction for age (Figure 1, left panel) can be interpreted as the indexing variability in young fertility before the age of 25, with the second marginal eigenfunction for age (Figure 1, central panel) indexing variability in fertility for mature age, between the ages of 20 and 40. As our modeling framework allows for rigorous uncertainty quantification in these posterior summaries, we note that the credible bands for the first and second eigenfunction are relatively wide, indicating that specific patterns should be interpreted with care. Examining directions of variance in fertility through the years, we note that the first marginal eigenfunction for year (Figure 1, left panel) is relatively constant and can be interpreted as representing an overall “size-component” of fertility from 1951 to 2006. The second eigenfunction (Figure 1, central panel) defines a contrast of fertility in countries before and after 1975. For both the year and age coordinates, the third marginal eigenfunctions capture a smaller fraction of the total variance and index higher patterns of dispersion at and around age 25 and at or around the year 1975.

We investigate sensitivity to the number of basis and latent factors considering four different models: model 1: $(p_1, p_2) = (44, 50)$, $(q_1, q_2) = (20, 20)$; model 2 $(p_1, p_2) = (44, 50)$, $(q_1, q_2) = (6, 6)$; model 3: $(p_1, p_2) = (16, 20)$, $(q_1, q_2) = (12, 12)$; and model 4: $(p_1, p_2) = (16, 20)$, $(q_1, q_2) = (6, 6)$. We also estimate the marginal covariance function with product FPCA using both the dense and sparse settings. Point estimate for $K_T(t, t')$ are reported in Figure 2. Comparing estimates within column (left and center panels), we assess sensitivity to a drastic reduction in the number of latent factors. Comparing estimates within row (left and center panels), we instead assess sensitivity to a drastic reduction in the number of basis functions. We note that the marginal age covariance function is relatively stable in all four settings. We contrast this relative robustness with estimates based on the product FPCA. In particular, sparse estimation using 10-fold cross-validation results in meaningfully diminished local features. A possible reason for the instability is due to the small sample size ($n = 17$). In this example, Bayesian estimation is perhaps preferable, as adaptive penalization allows for stable estimates within a broad class of model specifications.

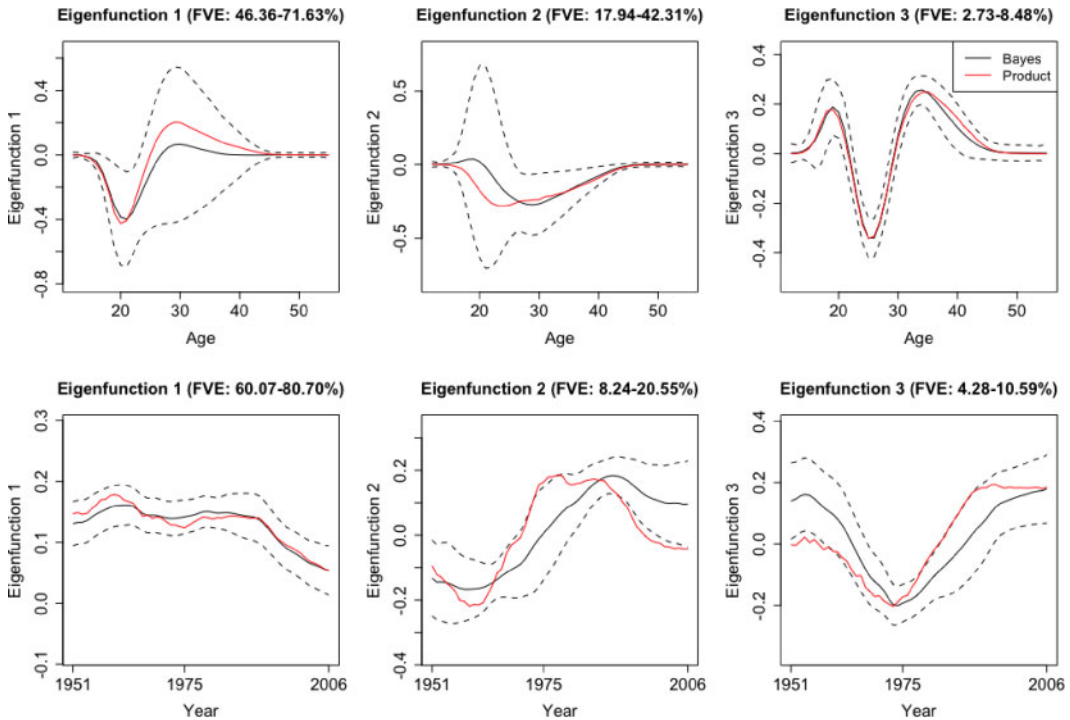


Fig. 1. Age and calendar year marginal eigenfunctions. The above plots include the Bayesian posterior means, 95% credible bands, and the product FPCA marginal eigenfunctions.

6.2. An EEG study on implicit learning in children with ASD

This analysis is motivated by a functional brain imaging study of implicit learning in young children with autism spectrum disorder (ASD), a developmental condition that affects an individual's communication and social interactions (Lord and others, 2000). Implicit learning is defined as learning without the intention to learn or without the conscious awareness of the knowledge that has been acquired. We consider functional brain imaging through EEG, an important and highly prevalent imaging paradigm aimed at studying macroscopic neural oscillations projected onto the scalp in the form of electrophysiological signals.

This study, carried out by our collaborators in the Jeste laboratory at UCLA, targets the neural correlates of implicit learning in the setting of an event-related shape learning paradigm (Jeste and others, 2015). Children aged 2–6 years old with ASD were recruited through the UCLA Early Childhood Partial Hospitalization Program (ECPHP). Each participant had an official diagnosis of ASD prior to enrollment. Age-matched typically developing (TD) children from the greater Los Angeles area were recruited as controls.

Six colored shapes (turquoise diamond, blue cross, yellow circle, pink square, green triangle, and red octagon) were presented one at a time in a continuous “stream” in the center of a computer monitor. There were three shape pairings randomized to each child. For instance, a pink square may always be followed by a blue cross. After the blue cross would come a new shape pair. Within a shape pair would constitute an “expected” transition and between shape pairs would constitute an “unexpected” transition. Each child would wear a 128-electrode Geodesic Sensor Net and observe the stream of shapes on the

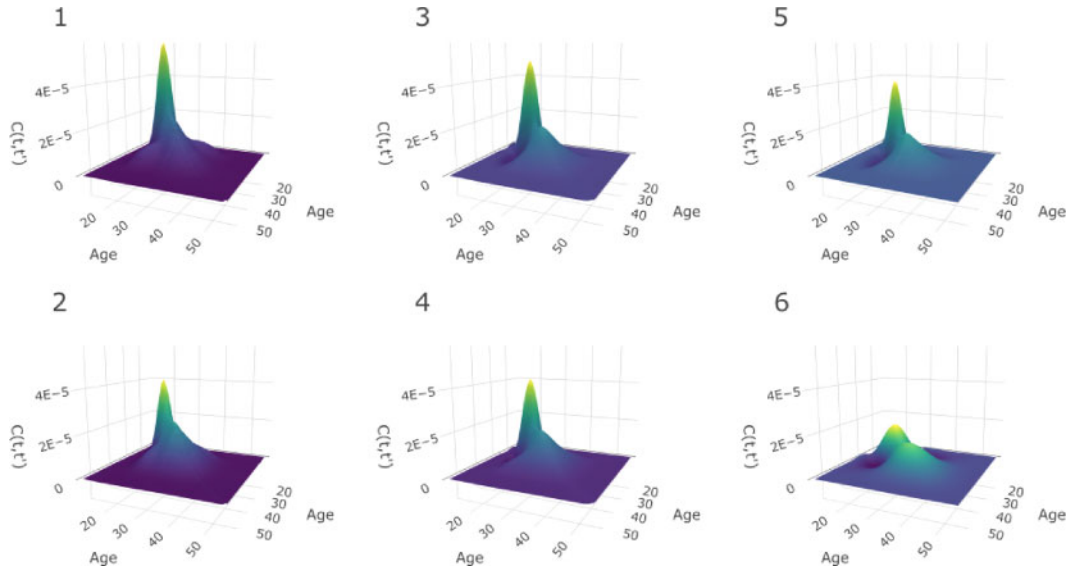


Fig. 2. Sensitivity analysis for the marginal covariance function $K_{\mathcal{T}}(t, t')$ (HFD study). Panels (1,2,3,4) refer to posterior mean estimates obtained under different projections and numbers of latent factors (Specific details are provided in Section 6.1). Panels (5, 6) refer to product FPCA estimates obtained under dense (5) or sparse (6) settings.

computer monitor. Each stimulus, or presentation of a single shape, is referred to as a trial, and can result in frequency-specific changes to ongoing EEG oscillations, which are measured as event related potentials (ERPs).

Each waveform contains a phasic component called the P300 peak which represents attention to salient information. This phasic component is typically studied in EEG experiments and is thought to be related to cognitive processes and early category recognition (Jeste and others, 2015). We use the same post-processed data as in Hasenstab and others (2017). Namely, we consider 37 ASD patients and 34 TD patients using data from trials 5 through 60 and averaging ERPs in a 30 trial sliding window (Hasenstab and others, 2015). The sliding window enhances the signal to noise ratio at which the P300 peak locations can be identified for each waveform. Each waveform is sampled at 250 Hz resulting in 250 within-trial time points over 1000 ms. Following Hasenstab and others (2017), we reduce each waveform to a 140 ms window around each P300 peak. This 140 ms window results in 37 within-trial time points. We do not apply warping techniques because each within-trial curve is centered about the P300 peak. Our analysis focuses on condition differentiation, formally defined as the difference between the expected and subsequent unexpected condition. Modeling condition differentiation for waveforms within a narrow window about the P300 peak over trials may give insights into learning rates for the ASD and TD groups. Thus, the main interest in this study is changes in condition differentiation over trials, and a longitudinal functional framework is required for statistical inference in this setting. Our analysis is based on the condition differentiation, averaged within subject over the four electrodes in the right frontal region of the brain. In summary, for each subject we consider $n_s = 56$ observations within trial, and $n_t = 37$ total trials.

We model the ASD and TD data cohorts separately, in order to estimate ERP time and trial covariance functions within group. All inference is based on a model with $p_1 = 20$, $p_2 = 56$, $q_1 = 10$, $q_2 = 28$, selected minimizing DIC. A comprehensive analysis is reported in the web-based supplementary material available at *Biostatistics* online. Statistical inference is based on 50K MCMC posterior draws, after 20K




Fig. 3. Posterior expected mean condition differentiation along trial and ERP time for the ASD (a) and the TD (b) cohorts.

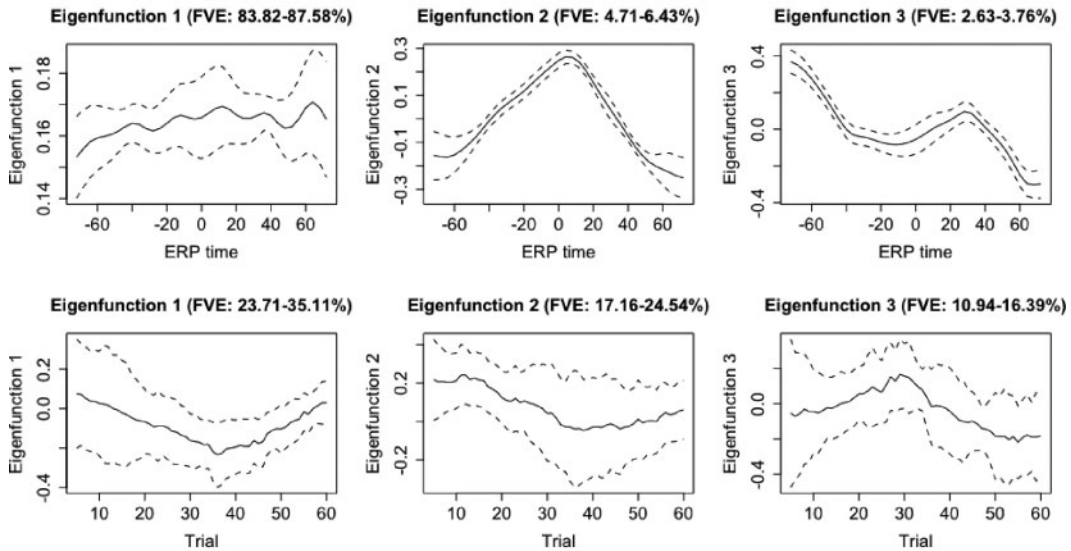
burn-in. We considered relatively diffuse priors: $a_\sigma = b_\sigma = 0.5$, $a_h = b_h = 1$, $v_1 = v_2 = 1$, $r_1 = r_2 = 1$, and $a_\varphi = b_\varphi = 0.0001$. Results are relatively insensitive to these hyperparameter settings.

The estimated mean surfaces for the two groups are plotted in Figure 3. The ASD group tends to have positive condition differentiation between trials 30 and 55, whereas the TD group tends to have positive condition differentiation in earlier trials. Positive condition differentiation is thought to be indicative of learning, so these results suggest that the TD group is learning at a faster rate than the ASD group. However, even though qualitatively the surfaces look very different, there is a substantial amount of heterogeneity in the subject-level data, resulting in broad confidence bands around the mean, and perhaps suggesting that differential patterns of condition differentiation between ASD and TD groups are best explored considering both the mean and the covariance structure.

Next we conduct an eigen-analysis of the covariance structure for both cohorts separately. Figure 4 plots eigenfunctions of the marginal covariances over ERP time and trials. Credible intervals are calculated following [Crainiceanu and others \(2007\)](#). We start by analyzing summaries indexing variability in ERP time. For both the TD and ASD cohorts, the first eigenfunction explains the vast majority of the marginal covariance (84–88% in ASD, and 86–90% in TD). In both groups, this first eigenfunction is relatively flat and can be interpreted as representing variability in the overall level of condition differentiation within a trial. The magnitude and shape of variation is comparable between TD and ASD children. Finer differences may be detected in the second and third eigenfunction, which further characterize variability in the shape of the ERP waveforms about the P300 peak. For both cohorts, however, these summaries represent only a small percentage of the variance in ERP waveform within trial.

Perhaps more interesting is an analysis of the marginal covariance across trial, as probabilistic learning patterns are likely to unfold with prolonged exposure to expected vs. unexpected shape pairings. For the ASD group, the first eigenfunction dips in an approximately quadratic fashion, suggesting enhanced variability in condition differentiation at around trial 35. Similarly, for the TD group, the first trial eigenfunction has a slight peak around trial 25. A possible interpretation of these covariance components relates to implicit learning, with higher variance in differentiation occurring earlier for TD than for ASD children. For both TD and ASD, the second eigenfunction across trials is interpreted as a contrast between high condition differentiation at early trials and low condition differentiation at later trials. Finally for the ASD

ASD



TD

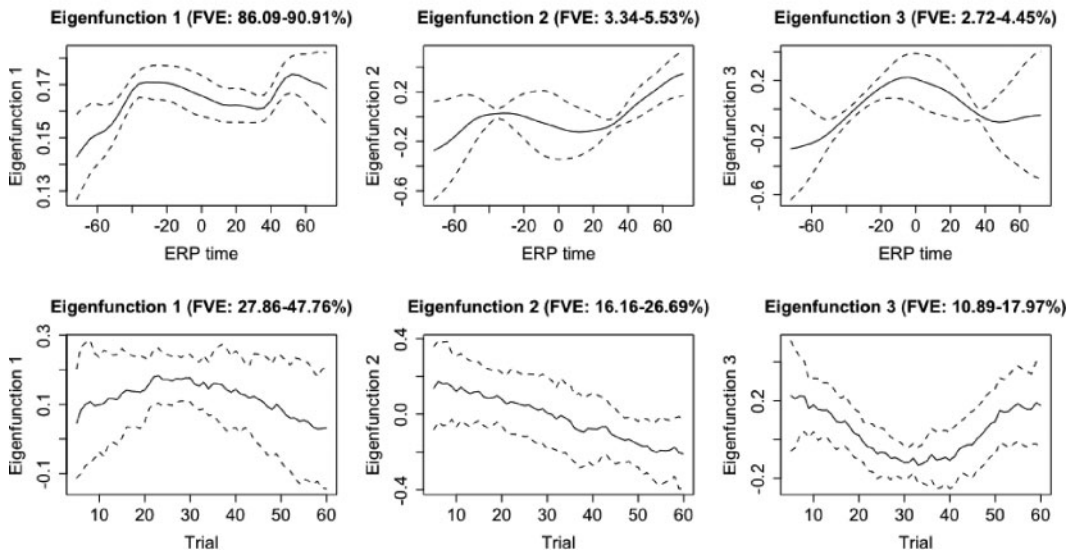


Fig. 4. Marginal eigenfunctions with associated uncertainty for the ASD and TD groups. Solid black lines represent posterior means and dotted lines represent 95% simultaneous credible bands.

cohort, the third eigenfunction exhibits a peak around trial 30. A possible interpretation would identify heterogeneity in the timing of learning, with some of the trajectories inducing variation in condition differentiation around trial 30, as opposed to the first eigenfunction identifying increased variance at around trial 35. Similarly for the TD group, the third trial eigenfunction has a dip around trial 35, indexing delayed increased variability in condition differentiation around trial 35.

7. DISCUSSION

In this article, we provide a probabilistic characterization of longitudinal functional data. As part of our work, we propose a joint framework for the estimation of the mean or the regression function, and a flexible prior for covariance operators. Regularized estimation relies crucially on the projection of a set of basis coefficients onto a latent subspace, with adaptive shrinkage achieved via a broadly supported class of product Gamma priors. While we have not established theoretical results on posterior consistency, we have shown that the proposed framework exhibits competitive operating characteristics, when compared with alternative modeling strategies.

Importantly, uncertainty quantification, is achieved without having to rely on the asymptotic performance of bootstrap methods. From an applied perspective, analysts are charged with choosing the appropriate projection space. However, we see this as a feature rather than a problem, as different data scenarios may require and motivate the use of alternative basis systems. Because regularization is achieved jointly with estimation, inference is straightforward and does not need to account separately for the estimation of nuisance parameters or the choice of a finite number of eigenfunctions to use in a truncated version of the model, as is the case for FPCA-based methods.

Crucially, the level of flexibility afforded by the proposed method is most important when inference centers on both the mean and covariance structure. Simpler modeling strategies, e.g., the Integrated Nested Laplace Approximations (LFPCA) approach of [Grevén and others \(2010\)](#), are likely to be more appropriate when the number of longitudinal observations is small, or if inference centers mostly on the mean structure. In a small simulation study (Appendix C of the supplementary material available at *Biostatistics* online), we found that the proposed approach performs similarly to LFPCA, even when data are generated from the latter scheme.

We have shown that posterior inference using MCMC is implemented in a relatively straightforward fashion and need not rely on complicated posterior sampling strategies. When dealing with large data-sets, this naïve inferential strategy may not be appropriate. For example, in updating the basis coefficients Θ_i , the number of floating point operations grow at a cubic rate with respect to the dimensions of the spline bases. When naïve Gibbs is not scalable (e.g., for very large samples or long evaluation domains), potentially promising acceleration strategies include the zero-loss projection approach of [Morris and Carroll \(2006\)](#), and adaptations of the INLA framework for approximate inference ([Rue and others, 2009](#)).

From a modeling perspective, our probabilistic characterization of the longitudinal functional covariance function is essentially equivalent to the weakly separable model of [Chen and others \(2017\)](#). While more general than a strictly separable model, this strategy makes strong assumptions about the structure of a high-dimensional covariance operator. Testing strategies have been developed in the literature ([Lynch and Chen, 2018](#)). However, we find that a more natural approach to the problem is one of regularized estimation. In this setting, a possible extension of our modeling framework could include an embedding strategy for the regularization of a nonseparable covariance operator towards a weakly separable one.

8. SOFTWARE

Software in the form of an R package including complete documentation and a sample data set is available from <https://github.com/jshamsho/LFBayes>.

SUPPLEMENTARY MATERIAL

Supplementary material is available online at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

<https://www.overleaf.com/project/5f31eabda15f6d000112076d>. *Conflict of Interest*: None declared.

FUNDING

The National Institute of Mental Health (R01 MH122428-01 to D.S., D.T.).

REFERENCES

- BALADANDAYUTHAPANI, V., MALLICK, B. K. AND CARROLL, R. J. (2005). Spatially adaptive Bayesian regression splines. *Journal of Computational and Graphical Statistics* **14**, 378–394.
- BALADANDAYUTHAPANI, V., MALLICK, B. K., YOUNG HONG, M., LUPTON, J. R., TURNER, N. D. AND CARROLL, R. J. (2008). Bayesian hierarchical spatially correlated functional data analysis with application to colon carcinogenesis. *Biometrics* **64**, 64–73.
- BHATTACHARYA, A. AND DUNSON, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika* **98**, 291–306.
- CHEN, K., DELICADO, P. AND MÜLLER, H.G. (2017). Modelling function-valued stochastic processes, with applications to fertility dynamics. *Journal of the Royal Statistical Society, Series B* **79**, 177–196.
- CHEN, K. AND MÜLLER, H.G. (2012). Modeling repeated functional observations. *Journal of the American Statistical Association* **107**, 1599–1609.
- CRAINICEANU, C.M., RUPPERT, D., CARROLL, R.J., JOSHI, A. AND GOODNER, B. (2007). Spatially adaptive Bayesian penalized splines with heteroscedastic errors. *Journal of Computational and Graphical Statistics* **16**, 265–288.
- CUEVAS, A., FEBRERO, M. AND FRAIMAN, R. (2006). On the use of the bootstrap for estimating functions with functional data. *Computational Statistics & Data Analysis* **51**, 1063 – 1074.
- DELATTRE, M., LAVIELLE, M. AND POURSAT, M. (2014). A note on BIC in mixed-effects models. *Electronic Journal of Statistics* **8**, 456–475.
- DI, C-Z., CRAINICEANU, C. M., CAFFO, B. S. AND PUNJABI, N. M. (2009). Multilevel functional principal component analysis. *Annals of Applied Statistics* **3**, 458–488.
- FERRATY, F., VAN KEILEGOM, I. AND VIEU, P. (2010). On the validity of the bootstrap in non-parametric functional regression. *Scandinavian Journal of Statistics* **37**, 286–306.
- GNEITING, T. (2002). Nonseparable, stationary covariance functions for space-time data. *Journal of the American Statistical Association* **97**, 590–600.
- GOLDSMITH, J., LIU, X., JACOBSON, J. S. AND RUNDLE, A. (2016). New insights into activity patterns in children, found using functional data analyses. *Medicine and Science in Sports and Exercise* **48**, 1723–1729.
- GREVEN, S., CRAINICEANU, C., CAFFO, B. AND REICH, D. (2010). Longitudinal functional principal component analysis. *Electronic Journal of Statistics* **4**, 1022–1054.
- HASENSTAB, K., SCHEFFLER, A., TELESKA, D., SUGAR, C. A., JESTE, S., DISTEFANO, C. AND ŞENTÜRK, D. (2017). A multi-dimensional functional principal components analysis of EEG data. *Biometrics* **73**, 999–1009.
- HASENSTAB, K., SUGAR, C. A., TELESKA, D., MCEVOY, K., JESTE, S. AND ŞENTÜRK, D. (2015). Identifying longitudinal trends within EEG experiments. *Biometrics* **71**, 1090–1100.
- HFD. (2019). *Human Fertility Database*. Max Planck Institute for Demographic Research (Germany) and Vienna Institute of Demography (Austria).
- JESTE, S. S., KIRKHAM, N., SENTURK, D., HASENSTAB, K., SUGAR, C., KUPELIAN, C., BAKER, E., SANDERS, A. J., SHIMIZU, C., NORONA, A., PAPARELLA, T., FREEMAN, S. F. N. and others. (2015). Electrophysiological evidence of heterogeneity in visual statistical learning in young children with ASD. *Developmental Science* **18**, 90–105.
- KRIVOBOKOVA, T., KNEIB, T. AND CLAESKENS, G. (2010). Simultaneous confidence bands for penalized spline estimators. *Journal of the American Statistical Association* **105**, 852–863.

- LEE, W., MIRANDA, M. F., RAUSCH, P., BALADANDAYUTHAPANI, V., FAZIO, M., DOWNS, J. C. AND MORRIS, J. S. (2019). Bayesian semiparametric functional mixed models for serially correlated functional data, with application to glaucoma data. *Journal of the American Statistical Association* **114**, 495–513.
- LORD, C., RISI, S., LAMBRECHT, L., COOK, E. H., LEVENTHAL, B. L., DILAVORE, P. C., PICKLES, A. AND RUTTER, M. (2000). The autism diagnostic observation schedule—generic: a standard measure of social and communication deficits associated with the spectrum of autism. *Journal of Autism and Developmental Disorders* **30**, 205–223.
- LYNCH, B. AND CHEN, K. (2018). A test of weak separability for multi-way functional data, with application to brain connectivity studies. *Biometrika* **105**, 815–831.
- MONTAGNA, S., TOKDAR, S., NEELON, B. AND DUNSON, D. B. (2012). Bayesian latent factor regression for functional and longitudinal data. *Biometrics* **68**, 1064–1073.
- MORRIS, J. S., BALADANDAYUTHAPANI, V., HERRICK, R. C., SANNA, P. AND GUTSTEIN, H. (2011). Automated analysis of quantitative image data using isomorphic functional mixed models, with application to proteomics data. *The Annals of Applied Statistics* **5**, 894–923.
- MORRIS, J. S. AND CARROLL, R. J. (2006). Wavelet-based functional mixed models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**, 179–199.
- MORRIS, J. S., VANNUCCI, M., BROWN, P. J. AND CARROLL, R. J. (2003). Wavelet-based nonparametric modeling of hierarchical functions in colon carcinogenesis. *Journal of the American Statistical Association* **98**, 573–583.
- PARK, S.Y. AND STAIU, A. (2015). Longitudinal functional data analysis. *Stat* **4**, 212–226.
- REN, S., LAI, H., TONG, W., AMINZADEH, M., HOU, X. AND LAI, S. (2010). Nonparametric bootstrapping for hierarchical data. *Journal of Applied Statistics* **37**, 1487–1498.
- RUE, H., MARTINO, S. AND CHOPIN, N. (2009). Approximate Bayesian inference for latent gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **71**, 319–392.
- SCHEFFLER, A., TELESKA, D., LI, Q., SUGAR, C., DISTEFANO, C., JESTE, S. AND ŞENTÜRK, D. (2020). Hybrid principal components analysis for region-referenced longitudinal functional EEG data. *Biostatistics* **21**, 139–157.
- SHI, J. Q. AND CHOI, T. (2011). *Gaussian Process Regression Analysis for Functional Data*. New York: Chapman and Hall/CRC.
- SUAREZ, A. J. AND GHOSAL, S. (2017). Bayesian estimation of principal components for functional data. *Bayesian Analysis* **12**, 311–333.
- YANG, J., COX, D.D., LEE, J.S. AND CHOI, P. REN T. (2017). Efficient Bayesian hierarchical functional data analysis with basis function approximations using Gaussian-Wishart processes. *Biometrics* **73**, 1082–1091.
- YANG, J., ZHU, H., CHOI, T. AND COX, D.D. (2016). Smoothing and mean–covariance estimation of functional data with a Bayesian hierarchical model. *Bayesian Analysis* **11**, 649–670.
- ZHANG, L., BALADANDAYUTHAPANI, V., ZHU, H., BAGGERLY, K. A., MAJEWSKI, T., CZERNIAK, B. A. AND MORRIS, J. S. (2016). Functional car models for large spatially correlated functional datasets. *Journal of the American Statistical Association* **111**, 772–786.

[Received September 19, 2019; revised August 31, 2020; accepted for publication September 4, 2020]