

Smoothing and mean-covariance estimation of functional data with a Bayesian hierarchical model

Jingjing Yang^{*,¶} and Hongxiao Zhu[†] and Taeryon Choi^{‡,||} and Dennis D. Cox^{§,¶}

Abstract. Functional data, with basic observational units being functions (e.g., curves, surfaces) varying over a continuum, are frequently encountered in various applications. While many statistical tools have been developed for functional data analysis, the issue of smoothing all functional observations simultaneously is less studied. Existing methods often focus on smoothing each individual function separately, at the risk of removing important systematic patterns common across functions. We propose a nonparametric Bayesian approach to smooth all functional observations simultaneously and nonparametrically. In the proposed approach, we assume that the functional observations are independent Gaussian processes subject to a common level of measurement errors, enabling the borrowing of strength across all observations. Unlike most Gaussian process regression models that rely on pre-specified structures for the covariance kernel, we adopt a hierarchical framework by assuming a Gaussian process prior for the mean function and an Inverse-Wishart process prior for the covariance function. These prior assumptions induce an automatic mean-covariance estimation in the posterior inference in addition to the simultaneous smoothing of all observations. Such a hierarchical framework is flexible enough to incorporate functional data with different characteristics, including data measured on either common or uncommon grids, and data with either stationary or nonstationary covariance structures. Simulations and real data analysis demonstrate that, in comparison with alternative methods, the proposed Bayesian approach achieves better smoothing accuracy and comparable mean-covariance estimation results. Furthermore, it can successfully retain the systematic patterns in the functional observations that are usually neglected by the existing functional data analyses based on individual-curve smoothing.

Keywords: functional data, smoothing, Bayesian hierarchical model, Gaussian process, Matérn covariance function, empirical Bayes.

1 Introduction

As more digital data are being collected in modern experiments, considerable efforts have been made to process and analyze *functional data* — data that are realizations

^{*}Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109, USA, yjingj@umich.edu

[†] Department of Statistics, Virginia Tech, Blacksburg, VA 24061, USA, hongxiao@vt.edu

[‡] Department of Statistics, Korea University, Seoul 136-701, Republic of Korea, trchoi@korea.ac.kr

[§] Department of Statistics, Rice University, Houston, TX 77005, USA, dcox@rice.edu

[¶]Supported by NIH grant PO1-CA-082710.

^{||}Supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education, Science and Technology (No. 2012R1A1A2041370).

of random functions varying over a continuum such as a region of time or a range of wavelengths (Ramsay and Silverman, 2002). Since Ramsay and Dalzell (1991) first coined the term “functional data analysis” (FDA), numerous papers have been published on the FDA theory, methods and applications, making it one of the most active research areas in statistics. One salient feature of functional data is that, although the underlying functions are often continuous and smooth, data can only be collected discretely, which often produces measurement errors. It is possible to treat functional data as fully observed when the measurements are dense and the noise level is low (Hall et al., 2001; Cardot et al., 2003; Zhu and Cox, 2009; Zhu et al., 2010). However, in most circumstances, it is necessary to convert the discrete measurements to functions that can be evaluated on any subset of the domain, referred to as smoothing. Smoothing reduces the effect of measurement errors, and better smoothing leads to less biased results in further analyses (see, e.g., Hitchcock et al. (2006)).

Existing approaches for smoothing functional data can be divided into three overlapping categories (see, e.g., Ramsay and Silverman (2005)): (1) methods based on representations with basis functions such as B-splines, Fourier series, and wavelets; (2) nonparametric smoothing with kernels; and (3) roughness-penalty-based methods such as smoothing spline. Although these existing smoothing methods are computationally tractable and generally available in standard statistical softwares, they mainly focus on smoothing each individual curve separately rather than smoothing all curves at the same time, thus failing to borrow strength across all observations. Such a failure may cause the following problems: (1) The order of smoothness of each curve may vary across all functional observations even if they are assumed to be independent and identically distributed (i.i.d.). This problem is severe when functional data are sparse and measured on uncommon grids. (2) The systematic patterns shared by all observations could be blurred or wiped out. This problem becomes more severe when the signal-to-noise-ratio (SNR) is low.

In these regards, we propose a model-based approach for smoothing all functional observations simultaneously and nonparametrically. In particular, we assume that all functional observations are i.i.d. Gaussian processes (GPs) subject to a common level of additive measurement errors, which facilitates the borrowing of strength across all data and the simultaneous smoothing based on the hierarchical GP structures. The GP prior is a nonparametric prior widely used in Bayesian nonparametric regression, spatial modeling, and machine learning (see, e.g., Rasmussen and Williams (2006), Banerjee et al. (2008), Banerjee et al. (2013), and Banerjee et al. (2014)). Recently, GP prior has also been increasingly employed in Bayesian FDA. For instance, GP has been used as the prior for functional batch effects (Kaufman et al., 2010) and as the base measure for a Dirichlet process prior (Nguyen and Gelfand, 2014) in functional ANOVA. GP regression has also been used in FDA (see, e.g., Shi and Choi (2011), Shi et al. (2012), and Wang and Shi (2014)), providing a nonparametric linkage between a functional predictor and a functional response. Despite the wide applicability of GP, most existing approaches assume pre-specified parametric structures for the covariance kernel such as squared exponential or Matérn, and the parameters in the kernel function often need to be pre-determined, which causes a lack of flexibility and introduces potential biases.

We contrast our proposed approach with existing GP methods for exploiting a Bayesian hierarchical model with a GP prior for the mean curve and an Inverse-Wishart process (IWP) prior for the covariance surface. By assuming the IWP prior, our approach avoids possible biases caused by the mis-specification of the covariance kernel as commonly used in the existing methods. Furthermore, our method has the flexibility of allowing the scale parameter of IWP to adjust a smooth covariance surface to be either stationary or nonstationary, which provides more flexible posterior covariance estimation as well as more accurate smoothing of functional observations

In addition, through borrowing strength across all curves, our approach facilitates the smoothing of sparse functional data measured on uncommon grids. An individual curve contains much less information about the functional observation when it is observed at very few grid points. Smoothing individual curves one-at-a-time is extremely difficult in this case. With the assumption that the functional observations share the common mean and covariance, our Bayesian approach can simultaneously smooth all curves on a pooled, dense grid, by treating functional values at unobserved grid points as parameters (or missing values) and updating them during a Markov Chain Monte Carlo (MCMC) procedure.

Finally, our proposed Bayesian inferential procedure also produces mean and covariance estimation as by-products. Functional mean-covariance estimation has already been extensively investigated in the literature. Existing approaches include the sample estimation (Ramsay and Silverman, 2005), the roughness-penalty-based method (Rice and Silverman, 1991, page 22), and the Principal Analysis by Conditional Expectation (PACE) method with local polynomial smoothing (Yao et al., 2005a). Regarding the mean-covariance estimation, our proposed method can be considered as a Bayesian counterpart of PACE. PACE aims to estimate eigenvalues and eigenfunctions of the covariance surface with irregularly spaced longitudinal data (Yao et al., 2005a,b; Müller, 2005; Leng and Müller, 2006). It also produces mean-covariance estimates as by-products. The main idea of PACE is to first formulate a “raw” covariance using pooled sparse longitudinal measurements and then apply a two-dimensional local polynomial smoother to estimate the covariance. Smoothing based on the pooled raw data has the effect of borrowing strength from all data; performing local polynomial smoothing has a similar effect as performing Bayesian smoothing with a GP prior. Therefore, the mean-covariance estimates produced by our approach should be comparable with the ones produced by PACE. In comparison with PACE, however, our Bayesian approach offers the benefit of full posterior inference for all latent processes and parameters, such as credible intervals and simultaneous credible bands for the mean curve and the covariance surface. These uncertainties are hard to quantify in a frequentist approach such as PACE. Furthermore, our method avoids the selection of optimal bandwidths for local polynomial smoothing as required in PACE, which can be sensitive to different selection methods.

Accordingly, the main contributions of our proposed approach can be summarized as follows: (1) It facilitates simultaneous smoothing of all functional observations. (2) It preserves systematic patterns that are common across all curves. (3) It avoids the potential bias caused by the mis-specification of the covariance function. (4) It is

flexible enough to accommodate functional data with common or uncommon grids, dense or sparse measurements, as well as stationarity or nonstationarity. (5) It results in mean-covariance estimates as by-products, with full Bayesian characterization of the uncertainties.

To demonstrate the performance of our proposed method, we considered simulation studies with synthetic data as well as two case studies with real spectroscopy and metabolic data. The results are compared with alternative methods based on single-curve smoothing (e.g., smoothing spline and kernel smoothing) as well as PACE. In simulation studies with noisy functional data based on either stationary or nonstationary covariances, we find that our proposed approach recovers the true signals more accurately than those single-curve smoothing methods. In real case studies, we find that our method retains systematic patterns that were wiped out by the single-curve smoothing methods. Regarding the mean-covariance estimation, our proposed approach is comparable with the PACE method as expected.

The remainder of this article is organized as follows: In Section 2, we outline the details of the proposed Bayesian hierarchical model as well as the heuristic empirical Bayes approach used to determine hyperparameters. In Section 3, we describe the MCMC algorithm for posterior inference. Results of simulations and real case studies are presented in Sections 4 and 5, respectively. Discussion is provided in Section 6.

2 Unified Bayesian hierarchical model

2.1 Model description

Suppose that the functional data contain n independent trajectories, denoted by $\{Y_i(\cdot); i = 1, 2, \dots, n\}$, and the i^{th} trajectory has p_i measurements on grid $\mathbf{t}_i = \{t_{i1}, \dots, t_{ip_i}\}$. Further assume that the i^{th} trajectory $Y_i(\cdot)$ depends on an underlying GP $Z_i(t)$, $t \in \mathcal{T}$, through the following model:

$$Y_i(t_{ij}) = Z_i(t_{ij}) + \epsilon_{ij}, \quad t_{ij} \in \mathcal{T}, \quad i = 1, \dots, n, \quad j = 1, \dots, p_i, \quad (2.1)$$

where $\{Z_i(\cdot)\}$ are i.i.d. GPs with mean function $\mu(\cdot)$ and covariance kernel $\Sigma(\cdot, \cdot)$, denoted by $Z_i \sim GP(\mu, \Sigma)$. The covariance kernel satisfies $\Sigma(s, t) = E[(Z_i(s) - \mu(s))(Z_i(t) - \mu(t))]$, $\forall s, t \in \mathcal{T}$. The error terms $\{\epsilon_{ij}\}$ are assumed to be i.i.d. random normal variables, i.e., $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$, and independent of $\{Z_i(\cdot)\}$. We assume the following priors for the model parameters σ_ϵ^2 , $\mu(\cdot)$, and $\Sigma(\cdot, \cdot)$:

$$\sigma_\epsilon^2 \sim \text{Inverse-Gamma}(a_\epsilon, b_\epsilon), \quad (\mu \mid \Sigma) \sim GP\left(\mu_0, \frac{1}{c}\Sigma\right), \quad \Sigma \sim \text{IWP}(\delta, \Psi), \quad (2.2)$$

where $\text{Inverse-Gamma}(a_\epsilon, b_\epsilon)$ denotes an Inverse Gamma distribution with shape parameter a_ϵ and scale parameter b_ϵ . The scaling parameter c in the GP prior satisfies $c > 0$. Here $\text{IWP}(\delta, \Psi)$ denotes the IWP with shape parameter δ and scaling parameter Ψ . An IWP is defined such that, on any finite grid $\mathbf{t} = \{t_1, t_2, \dots, t_p\}$ with p points, the projection $\Sigma(\mathbf{t}, \mathbf{t})$ is Inverse-Wishart distributed, i.e., $\Sigma(\mathbf{t}, \mathbf{t}) \sim \text{IW}(\delta, \Psi(\mathbf{t}, \mathbf{t}))$. Here the

Inverse-Wishart distribution follows the parametrization of Dawid (1981). In particular, a symmetric and positive definite matrix $\Sigma(\mathbf{t}, \mathbf{t})$ is said to be $IW(\delta, \Psi(\mathbf{t}, \mathbf{t}))$ distributed if $\mathbf{K} = \Sigma(\mathbf{t}, \mathbf{t})^{-1}$ is Wishart distributed with degrees of freedom $\delta + p - 1$ and scale matrix $\Psi(\mathbf{t}, \mathbf{t})^{-1}$.

The advantage of adopting the parameterization of Dawid (1981) is two-fold: the parameter δ does not vary with the dimension of \mathbf{t} , and the resulting Inverse-Wishart distribution is consistent under marginalization. These properties are essential in showing that the IWP prior for $\Sigma(\cdot, \cdot)$ is well-defined when the number of grid points p approaches infinity. This is summarized in Proposition 2.1. Proposition 2.1 essentially states that $IWP(\delta, \Psi)$ is a well-defined (infinite-dimensional) probability measure and its finite-dimensional projection on grid $\mathbf{t} \times \mathbf{t}$ coincides with the Inverse-Wishart distribution $IW(\delta, \Psi(\mathbf{t}, \mathbf{t}))$. The proof can be derived through following similar arguments as in the proof of Lemma 2 in the Appendix of Zhu et al. (2014).

Proposition 2.1. *Let $\mathcal{T} \subseteq \mathbb{R}$ be a compact set and $\delta > 4$ be a positive integer. Suppose that $\Psi : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ is a symmetric and positive semi-definite mapping, i.e., any evaluation of Ψ on a finite grid $\mathbf{t} \times \mathbf{t} \subseteq \mathcal{T} \times \mathcal{T}$ gives a symmetric and positive semi-definite matrix. Then there exists a unique probability measure λ on $(\mathbb{R}^{\mathcal{T} \times \mathcal{T}}, \mathcal{B}(\mathbb{R}^{\mathcal{T} \times \mathcal{T}}))$, such that for any finite discretization \mathbf{t} , $\lambda_{\mathbf{t} \times \mathbf{t}} = IW(\delta, \Psi(\mathbf{t}, \mathbf{t}))$. We denote $\lambda = IWP(\delta, \Psi)$.*

In the prior distribution of $\Sigma \sim IWP(\delta, \Psi)$, a smaller value of δ leads to a less informative prior; the parameter Ψ controls the prior covariance structure. To encourage a smoother estimation of the covariance, we set $\Psi(\cdot, \cdot) = \sigma_s^2 A(\cdot, \cdot)$, where σ_s^2 is a positive scaling parameter and $A(\cdot, \cdot)$ is a smooth correlation/covariance kernel which can be either stationary or nonstationary. The structure of $A(\cdot, \cdot)$ may be nonparametric or of parametric form with hyperparameters. In this paper, we use the Matérn parameterization as an example of the stationary correlation structure, i.e., $\Psi(t_i, t_j) = \sigma_s^2 \text{Matern}_{cor}(|t_i - t_j|; \rho, \nu)$, where the Matérn correlation kernel $\text{Matern}_{cor}(d; \rho, \nu)$ is defined as

$$\text{Matern}_{cor}(d; \rho, \nu) = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\sqrt{2\nu} \frac{d}{\rho} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{d}{\rho} \right), \quad d \geq 0, \rho > 0, \nu > 0.$$

In the above expression, d denotes the distance between two measurement points; ρ is the scale parameter; ν is the order; and $K_\nu(\cdot)$ is the modified Bessel function of the second kind. One of the important properties of the Matérn covariance kernel is its positive definiteness (Stein, 1999). Both ρ and ν can influence the smoothness of the signal estimates. We choose $\nu > 2$ so that the resulting GP is $\lfloor \nu - 1 \rfloor$ times differentiable, which ensures the smoothness of the underlying data. It is also convenient to take ν to be an integer plus 0.5, in which case the Matérn correlation kernel has a closed-form expression. For example, if $\nu = 2.5$, the correlation kernel $A(\cdot, \cdot)$ takes the form

$$A(t_i, t_j) = \left(1 + \frac{\sqrt{5}|t_i - t_j|}{\rho} + \frac{5(|t_i - t_j|)^2}{3\rho^2} \right) \exp \left(-\frac{\sqrt{5}|t_i - t_j|}{\rho} \right). \quad (2.3)$$

Notice that the Matérn structure is sensitive to both ρ and ν , while the estimation of these parameters can be unstable. Zhang (2004) has pointed out in a geostatistical setup

that not all three parameters (σ_s^2, ρ, ν) in the Matérn class can be estimated consistently. To ensure a stable covariance estimation, we set ρ and ν at fixed values obtained from empirical correlation estimation, by minimizing the mean square error between an empirical correlation estimate and the $\text{Matern}_{cor}(\rho, \nu)$ function. The hyperprior for σ_s^2 is set to be $\sigma_s^2 \sim \text{Gamma}(a_s, b_s)$, i.e., a Gamma distribution with shape parameter a_s and inverse scale parameter b_s .

The above Bayesian hierarchical model is constructed based on infinite-dimensional GPs. However, practical posterior calculations can only be conducted in a finite manner. Since we assume latent GPs in our model, posterior inference can be performed similarly as in GP regression. In particular, the latent processes $\{Z_i(\cdot)\}$ and the parameters $\mu(\cdot)$, $\Sigma(\cdot, \cdot)$ will be inferred on a finite grid, while the inference on non-grid points, if needed, can be obtained by posterior prediction. Therefore, we will represent the likelihood for model (2.1) and the priors (2.2) in multivariate forms through evaluating the functions on finite grids. Denoting $Y_i(\mathbf{t}_i)$ by $\mathbf{Y}_{\mathbf{t}_i}$, $Z_i(\mathbf{t}_i)$ by $\mathbf{Z}_{\mathbf{t}_i}$, model (2.1) implies that

$$\mathbf{Y}_{\mathbf{t}_i} | \mathbf{Z}_{\mathbf{t}_i}, \sigma_\epsilon^2 \sim N(\mathbf{Z}_{\mathbf{t}_i}, \sigma_\epsilon^2 \mathbf{I}), \quad i = 1, \dots, n, \quad (2.4)$$

$$\mathbf{Z}_{\mathbf{t}_i} | \mu(\mathbf{t}_i), \Sigma(\mathbf{t}_i, \mathbf{t}_i) \sim N(\mu(\mathbf{t}_i), \Sigma(\mathbf{t}_i, \mathbf{t}_i)), \quad (2.5)$$

where \mathbf{I} is a $p_i \times p_i$ identity matrix. Since the grids $\{\mathbf{t}_i; i = 1, 2, \dots, n\}$ are not required to be common, we evaluate the GP and IWP prior distributions in (2.2) on the pooled grid $\mathbf{t} = \cup_{i=1}^n \mathbf{t}_i$. Denote $\mu(\mathbf{t})$ by $\boldsymbol{\mu}$, $\mu_0(\mathbf{t})$ by $\boldsymbol{\mu}_0$, $\Sigma(\mathbf{t}, \mathbf{t})$ by $\boldsymbol{\Sigma}$ and $\Psi(\mathbf{t}, \mathbf{t})$ by $\boldsymbol{\Psi}$, then

$$\boldsymbol{\mu} | \boldsymbol{\Sigma} \sim N\left(\boldsymbol{\mu}_0, \frac{1}{c} \boldsymbol{\Sigma}\right), \quad \boldsymbol{\Sigma} \sim \text{IW}(\delta, \boldsymbol{\Psi}). \quad (2.6)$$

The multivariate representations in (2.4)–(2.6) enable us to write the joint posterior distribution of $(\{Z_i(\mathbf{t})\}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2, \sigma_s^2)$, based on which posterior sampling can be performed by MCMC; details are presented in Section 3. The posterior means of $\{Z_i(\mathbf{t})\}$ will be treated as smoothed functional signals.

2.2 Prior parameter setup

The proposed Bayesian hierarchical model described in Section 2.1 contains several hyper-parameters: $(c, \boldsymbol{\mu}_0, \nu, \rho, a_\epsilon, b_\epsilon, a_s, b_s)$. Their values are determined using empirical methods. In particular, we set $c = 1$, which implies that the functional mean has Gaussian prior with the same covariance kernel as the data. We set $\boldsymbol{\mu}_0$ to be the smoothed sample mean of $\{\mathbf{Y}_{\mathbf{t}_i}\}$ by existing methods. To facilitate stable covariance estimation, we take ν and ρ as their empirical estimates $\hat{\nu}$ and $\hat{\rho}$ respectively. The empirical estimates can be obtained by minimizing the mean square error between an empirical correlation estimate and a $\text{Matern}_{cor}(\mathbf{D}; \rho, \nu)$ correlation function with distance matrix \mathbf{D} , subject to $\rho > 0$ and $\nu \geq 2.5$.

The values of $(a_\epsilon, b_\epsilon, a_s, b_s)$ are determined using a heuristic empirical Bayes approach with empirical estimates of $\{\sigma_\epsilon^2, \sigma_s^2\}$, which is described as follows:

- The value of $\hat{\sigma}_\epsilon^2$ can be easily obtained using a differencing technique (Von Neumann, 1941)

$$\hat{\sigma}_\epsilon^2 = \frac{1}{2 \sum_{i=1}^n (p_i - 1)} \sum_{i=1}^n \sum_{j=1}^{p_i-1} (Y_i(t_{i(j+1)}) - Y_i(t_{ij}))^2. \quad (2.7)$$

- A moment estimator of σ_s^2 can be derived by taking expectation (with respect to the prior distribution of Σ) and applying a trace operator to both sides of the equation $\text{Cov}(Y(\mathbf{t})) = \Sigma + \sigma_\epsilon^2 \mathbf{I}$:

$$\begin{aligned} \text{trace}(E\{\text{Cov}(Y(\mathbf{t}))\}) &= \frac{\text{trace}(\Psi)}{\delta - 2} + \sigma_\epsilon^2 \text{trace}(\mathbf{I}) = \frac{\sigma_s^2 \text{trace}(\mathbf{A})}{\delta - 2} + p\sigma_\epsilon^2, \\ \hat{\sigma}_s^2 &\approx \frac{\text{trace}(E\{\text{Cov}(Y(\mathbf{t}))\}) - p\hat{\sigma}_\epsilon^2}{\text{trace}(\mathbf{A})/(\delta - 2)}, \end{aligned} \quad (2.8)$$

where p is the length of the pooled grid \mathbf{t} ; $\hat{\sigma}_\epsilon^2$ is given by (2.7); and $E\{\text{Cov}(Y(\mathbf{t}))\}$ can be estimated by an empirical method, e.g. the covariance estimate by PACE (Yao et al., 2005b).

- The values of $(a_\epsilon, b_\epsilon, a_s, b_s)$ can be estimated by the method of moments with the prior distributions of $\sigma_\epsilon^2, \sigma_s^2$ and the empirical estimates in (2.7) and (2.8).

3 Posterior inference with MCMC

Based on the multivariate representations (2.4)–(2.6), we derive the joint posterior distribution and develop a MCMC algorithm for posterior inference in this section. Denote the observed data by $\mathbf{Y} = \{\mathbf{Y}_{\mathbf{t}_1}, \mathbf{Y}_{\mathbf{t}_2}, \dots, \mathbf{Y}_{\mathbf{t}_n}\}$; denote the underlying GP evaluations on the observational grids by $\mathbf{Z} = \{\mathbf{Z}_{\mathbf{t}_1}, \mathbf{Z}_{\mathbf{t}_2}, \dots, \mathbf{Z}_{\mathbf{t}_n}\}$; and denote the evaluations on the pooled grid by $\tilde{\mathbf{Z}} = \{Z_1(\mathbf{t}), Z_2(\mathbf{t}), \dots, Z_n(\mathbf{t})\}$. Recall that $\mathbf{t} = \cup_i \mathbf{t}_i$ is the pooled grid. Let $\mathbf{Z}^* = \tilde{\mathbf{Z}} \setminus \mathbf{Z}$, i.e. $\mathbf{Z}^* = \{\mathbf{Z}_{\mathbf{t}_1^*}, \mathbf{Z}_{\mathbf{t}_2^*}, \dots, \mathbf{Z}_{\mathbf{t}_n^*}\}$, where $\mathbf{Z}_{\mathbf{t}_i^*} = Z_i(\mathbf{t}_i^*)$ and $\mathbf{t}_i^* = \mathbf{t} \setminus \mathbf{t}_i$ is the set of grid points for the i th trajectory with missing observations. The joint posterior density of all parameters can be written as

$$\begin{aligned} f(\tilde{\mathbf{Z}}, \boldsymbol{\mu}, \Sigma, \sigma_\epsilon^2, \sigma_s^2 | \mathbf{Y}) &\propto f(\mathbf{Y} | \tilde{\mathbf{Z}}, \sigma_\epsilon^2) f(\tilde{\mathbf{Z}} | \boldsymbol{\mu}, \Sigma) f(\sigma_\epsilon^2) f(\boldsymbol{\mu} | \Sigma) f(\Sigma | \sigma_s^2) f(\sigma_s^2) \\ &\propto f(\mathbf{Y} | \mathbf{Z}, \sigma_\epsilon^2) f(\mathbf{Z}^* | \mathbf{Z}, \boldsymbol{\mu}, \Sigma) f(\mathbf{Z} | \boldsymbol{\mu}, \Sigma) f(\sigma_\epsilon^2) f(\boldsymbol{\mu} | \Sigma) f(\Sigma | \sigma_s^2) f(\sigma_s^2). \end{aligned} \quad (3.1)$$

Here we have factored the joint prior of $\tilde{\mathbf{Z}}$ as $f(\tilde{\mathbf{Z}} | \boldsymbol{\mu}, \Sigma) = f(\mathbf{Z}^* | \mathbf{Z}, \boldsymbol{\mu}, \Sigma) f(\mathbf{Z} | \boldsymbol{\mu}, \Sigma)$, which enables the updates of \mathbf{Z} and \mathbf{Z}^* by a Gibbs sampler. To design a MCMC algorithm of the Gibbs sampler, we need to derive conditional distributions for the latent variables \mathbf{Z} , \mathbf{Z}^* and the model parameters $\boldsymbol{\mu}, \Sigma, \sigma_s^2, \sigma_\epsilon^2$. For brevity, we only present the conditional posterior distributions of \mathbf{Z} and \mathbf{Z}^* in Section 3.1. The derivation of the conditional posteriors for the remain parameters are trivial due to the conjugacy of their priors.

3.1 Conditional posteriors

In case that all functional data are observed on a common grid, i.e. $\{\mathbf{t}_i \equiv \mathbf{t}; i = 1, 2, \dots, n\}$, then \mathbf{Z}^* vanishes and the conditional posterior distribution of $\mathbf{Z}_{\mathbf{t}_i}$ can be derived from

$$f(\mathbf{Z}_{\mathbf{t}_i} | \mathbf{Y}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2) \propto f(\mathbf{Y}_{\mathbf{t}_i} | \mathbf{Z}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2) f(\mathbf{Z}_{\mathbf{t}_i} | \boldsymbol{\mu}, \boldsymbol{\Sigma}),$$

which gives

$$\begin{aligned} (\mathbf{Z}_{\mathbf{t}_i} | \mathbf{Y}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2) &\sim N(\widetilde{\boldsymbol{\mu}}_i, \widetilde{\mathbf{V}}_i), \\ \widetilde{\mathbf{V}}_i &= ((1/\sigma_\epsilon^2)\mathbf{I} + \boldsymbol{\Sigma}^{-1})^{-1}; \\ \widetilde{\boldsymbol{\mu}}_i &= \widetilde{\mathbf{V}}_i ((1/\sigma_\epsilon^2)\mathbf{Y}_{\mathbf{t}_i} + \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}(\mathbf{t}_i)), \text{ where } \boldsymbol{\Sigma} = \boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i). \end{aligned} \quad (3.2)$$

In case that functional data are collected on uncommon grids, we will first update $\mathbf{Z}_{\mathbf{t}_i^*}$ from

$$\begin{aligned} (\mathbf{Z}_{\mathbf{t}_i^*} | \mathbf{Z}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) &\sim N(\boldsymbol{\mu}_i^*, \mathbf{V}_i^*), \\ \boldsymbol{\mu}_i^* &= \boldsymbol{\mu}(\mathbf{t}_i^*) + \boldsymbol{\Sigma}(\mathbf{t}_i^*, \mathbf{t}_i)\boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i)^{-1}(\mathbf{Z}_{\mathbf{t}_i} - \boldsymbol{\mu}(\mathbf{t}_i)) = \mathbf{B}_i\mathbf{Z}_{\mathbf{t}_i} - \mathbf{u}_i, \\ &\text{where } \mathbf{B}_i = \boldsymbol{\Sigma}(\mathbf{t}_i^*, \mathbf{t}_i)\boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i)^{-1}, \quad \mathbf{u}_i = \mathbf{B}_i\boldsymbol{\mu}(\mathbf{t}_i) - \boldsymbol{\mu}(\mathbf{t}_i^*); \\ \mathbf{V}_i^* &= \boldsymbol{\Sigma}(\mathbf{t}_i^*, \mathbf{t}_i^*) - \boldsymbol{\Sigma}(\mathbf{t}_i^*, \mathbf{t}_i)\boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i)^{-1}\boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i^*). \end{aligned} \quad (3.3)$$

Then update $\mathbf{Z}_{\mathbf{t}_i}$ from the conditional posterior distribution

$$f(\mathbf{Z}_{\mathbf{t}_i} | \mathbf{Y}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2, \mathbf{Z}_{\mathbf{t}_i^*}) \propto f(\mathbf{Y}_{\mathbf{t}_i} | \mathbf{Z}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2) f(\mathbf{Z}_{\mathbf{t}_i} | \boldsymbol{\mu}, \boldsymbol{\Sigma}) f(\mathbf{Z}_{\mathbf{t}_i^*} | \mathbf{Z}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (3.4)$$

which gives

$$\begin{aligned} (\mathbf{Z}_{\mathbf{t}_i} | \mathbf{Y}_{\mathbf{t}_i}, \mathbf{Z}_{\mathbf{t}_i^*}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2) &\sim N(\widetilde{\boldsymbol{\mu}}_i, \widetilde{\mathbf{V}}_i), \\ \widetilde{\mathbf{V}}_i &= ((1/\sigma_\epsilon^2)\mathbf{I} + \boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i)^{-1} + \mathbf{B}_i^T(\mathbf{V}_i^*)^{-1}\mathbf{B}_i)^{-1}; \\ \widetilde{\boldsymbol{\mu}}_i &= \widetilde{\mathbf{V}}_i ((1/\sigma_\epsilon^2)\mathbf{Y}_{\mathbf{t}_i} + \boldsymbol{\Sigma}(\mathbf{t}_i, \mathbf{t}_i)^{-1}\boldsymbol{\mu}(\mathbf{t}_i) + \mathbf{B}_i^T(\mathbf{V}_i^*)^{-1}(\mathbf{u}_i + \mathbf{Z}_{\mathbf{t}_i^*})). \end{aligned} \quad (3.5)$$

When \mathbf{t}^* and \mathbf{t} are both dense, the conditional variance \mathbf{V}_i^* in (3.3) can be very close to a zero matrix, which could cause numerical instability when inverting \mathbf{V}_i^* . In this case, we suggest using the conditional mean in (3.3) as a sample for $\mathbf{Z}_{\mathbf{t}_i^*}$ in the MCMC algorithm and removing the conditional prior $f(\mathbf{Z}_{\mathbf{t}_i^*} | \mathbf{Z}_{\mathbf{t}_i}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ from (3.4) when updating $\mathbf{Z}_{\mathbf{t}_i}$ (note that now (3.5) collapses to (3.2)).

3.2 MCMC algorithm

Based on the joint posterior distributions derived from (3.1), we design a MCMC algorithm for posterior sampling as follows:

Step 0: Set initial values. Set $(\boldsymbol{\mu}, \sigma_\epsilon^2)$ to be the empirical estimates, \mathbf{Z} to be the raw data \mathbf{Y} , and $\boldsymbol{\Sigma}$ to be an identity matrix. The prior parameters $(c, \boldsymbol{\mu}_0, \nu, \rho, a_\epsilon, b_\epsilon, a_s, b_s)$ are set as described in Section 2.2.

Step 1: Conditional on \mathbf{Y} and current values of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, update \mathbf{Z}^* and \mathbf{Z} . In the general case where all data are observed on uncommon grids, update \mathbf{Z}^* and \mathbf{Z} from (3.3) and (3.5) alternatively. In the case of common grids, update $\tilde{\mathbf{Z}}$ (identical to \mathbf{Z}) from (3.2).

Step 2: Conditional on \mathbf{Y} and current values of \mathbf{Z} , update the noise variance σ_ϵ^2 by

$$(\sigma_\epsilon^2 | \mathbf{Y}, \mathbf{Z}) \sim \text{Inverse-Gamma} \left(a_\epsilon + \frac{\sum_{i=1}^n p_i}{2}, \quad b_\epsilon + \frac{1}{2} \sum_{i=1}^n [(\mathbf{Y}_{t_i} - \mathbf{Z}_{t_i})^T (\mathbf{Y}_{t_i} - \mathbf{Z}_{t_i})] \right).$$

Step 3: Conditional on current values of $\tilde{\mathbf{Z}} = \mathbf{Z} \cup \mathbf{Z}^*$ and $\boldsymbol{\Sigma}$, update $\boldsymbol{\mu}$ from

$$(\boldsymbol{\mu} | \tilde{\mathbf{Z}}, \boldsymbol{\Sigma}) \sim N \left(\frac{1}{n+c} \left(\sum_{i=1}^n Z_i(t) + c\boldsymbol{\mu}_0 \right), \quad \frac{1}{n+c} \boldsymbol{\Sigma} \right).$$

Step 4: Conditional on current values of $\tilde{\mathbf{Z}} = \mathbf{Z} \cup \mathbf{Z}^*$ and $\boldsymbol{\mu}$, update $\boldsymbol{\Sigma}$ from

$$(\boldsymbol{\Sigma} | \tilde{\mathbf{Z}}, \boldsymbol{\mu}) \sim IW(n + \delta + 1, \mathbf{Q}), \quad \mathbf{Q} = (\tilde{\mathbf{Z}} - \boldsymbol{\mu}\mathbf{J})(\tilde{\mathbf{Z}} - \boldsymbol{\mu}\mathbf{J})^T + c(\boldsymbol{\mu} - \boldsymbol{\mu}_0)(\boldsymbol{\mu} - \boldsymbol{\mu}_0)^T + \sigma_s^2 \mathbf{A},$$

where $\tilde{\mathbf{Z}}$ denotes a matrix with columns $\{Z_i(\mathbf{t})\}$; $\mathbf{J} = (1, \dots, 1)$ is a vector of ones with length p ; and matrix \mathbf{A} is given by equation (2.3).

Step 5: Given current values of $\boldsymbol{\Sigma}$, update σ_s^2 from

$$(\sigma_s^2 | \boldsymbol{\Sigma}) \sim \text{Gamma} \left(a_s + \frac{(\delta + p - 1)p}{2}, \quad b_s + \frac{1}{2} \text{trace}(\mathbf{A}\boldsymbol{\Sigma}^{-1}) \right).$$

Repeat Steps 1 ~ 5 for a large number of iterations until convergence is achieved.

For the simulations and real case studies in Sections 4 and 5, we ran 10,000 MCMC iterations after a burn-in period of 2,000 iterations. The convergence of the MCMC chains were diagnosed by the Gelman and Rubin diagnostic method (Gelman and Rubin, 1992; Särkkä and Aki, 2014).

4 Simulation studies

Two simulation studies were performed for functional data with stationary and nonstationary covariance, respectively. Within each study, we analyzed functional observations on both common and uncommon grids.

4.1 Simulation 1: functional data with stationary covariance

Two types of functional data with stationary covariance were generated with common and uncommon grids. The underlying smooth functional data were generated from a GP with mean $\mu(t) = 3 \sin(4t)$ and covariance kernel $\Sigma(s, t) = 5 \text{Matern}_{cor}(|s - t|; \rho, \nu)$, where $\rho = 0.5$, $\nu = 3.5$, and $s, t \in [0, \pi/2]$. Each dataset contained $n = 50$ observations.

The common grid contained $p = 80$ equally spaced grid points on $[0, \pi/2]$. Noise terms $\{\epsilon_{ij}; i = 1, \dots, n, j = 1, \dots, p\}$ were independently generated from $N(0, \sigma_\epsilon^2)$ with $\sigma_\epsilon = \sqrt{5}/2$, and then added to the smooth functional data, resulting the final dataset with SNR equal to 2. In the uncommon grid case, we randomly retained data on 60% of the common grid points, which resulted in moderately sparse functional data on uncommon grids. Two sample curves for the common and uncommon grid cases are shown in Figure 1(a, b), where the gray lines/line-segments denote the raw data.

Based on the data generated above, we applied the MCMC algorithm in Section 3.2 to obtain posterior samples. Taking the common grid case as an example, we set $\delta = 5$, $\hat{\rho} = 0.503$, $\hat{\nu} = 3.459$, $\hat{\sigma}_\epsilon^2 = 1.239$, and $\hat{\sigma}_s^2 = 10.750$ using empirical estimation methods. With $\hat{\sigma}_\epsilon^2$ and $\hat{\sigma}_s^2$, we obtained values for the hyperprior parameters $\{a_\epsilon = 0.807, b_\epsilon = 1, a_s = 214.990, b_s = 20\}$. The posterior means of the parameters were calculated by averaging 10,000 posterior samples. The parameter σ_ϵ^2 , with true value 1.250, had posterior mean 1.244 and 95% credible interval $[1.187, 1.302]$. The parameter σ_s^2 , with unidentifiable true value, had posterior mean 73.443 and 95% credible interval $[69.700, 77.620]$.

In Figure 1(a, b), we display the smoothed curves (black solid lines), together with the raw data (gray lines/line segments), the 95% pointwise credible intervals (black dash-dot lines), and the underlying true smooth curves (blue dots) for the common (frame (a)) and uncommon grid (frame (b)) cases. The heat maps of our Bayesian correlation estimates are shown in Figure 1(c, d) for common (frame (c)) and uncommon (frame (d)) grid cases. In contrast, we also plotted the heat map of the sample correlation in the common grid case in Figure 1(e), and the heat map of the underlying true correlation in Figure 1(f). These plots show that our method can recover the true smooth signals accurately for both common and uncommon grid cases, while providing a much smoother correlation estimate than the sample estimate in Figure 1(e). Even though around 40% of the grid points are missing in each trajectory, the results of smoothing and mean-covariance estimation as in Figure 1(a, d) are almost as good as the common grid results as in Figure 1(b, c).

To further assess the performance of the proposed method, we compared the results with four alternative smoothing methods: *best possible least square* (BLS) estimation, smoothing spline with cubic splines applied to each individual curve (Spline), kernel smoothing with local polynomials applied to each individual curve (Kernel), and PACE. The BLS method served as an ‘‘oracle’’ model, in which the smoothed curves were estimated by the conditional mean of the latent GPs, while assuming the true mean and covariance of the latent processes were known. In the common grid case, the conditional mean and covariance for each signal can be written as

$$\begin{aligned} E\{\mathbf{Z}_i | \mathbf{Y}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}\} &= \boldsymbol{\mu} + \boldsymbol{\Sigma}(\boldsymbol{\Sigma} + \sigma_\epsilon^2 \mathbf{I})^{-1}(\mathbf{Y}_i - \boldsymbol{\mu}), \\ \text{Cov}(\mathbf{Z}_i | \mathbf{Y}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) &= \boldsymbol{\Sigma} - \boldsymbol{\Sigma}(\boldsymbol{\Sigma} + \sigma_\epsilon^2 \mathbf{I})^{-1}\boldsymbol{\Sigma}, \end{aligned}$$

which can be easily derived from the joint Gaussian distribution of \mathbf{Y}_i and \mathbf{Z}_i . We obtained the Spline and Kernel estimates using R (R Core Team, 2013) functions `smooth.spline` and `locpoly` from the `KernSmooth` library. The smoothing parameter in the former function was determined by generalized cross-validation (GCV) and




Figure 1: Simulation 1: (a) Two sample curves for functional observations in the common grid case (gray lines), superimposed by the Bayesian smoothed curves (black solid lines), along with the 95% pointwise credible intervals (black dashed lines) and the true signals (blue dots); (b) plots (same as in (a)) for the uncommon grid case; (c) heat map of our Bayesian correlation estimate in the common grid case; (d) heat map (same as in (c)) in the uncommon grid case; (e) heat map of the sample correlation in the common grid case; (f) heat map of the true underlying (Martérn) correlation.

the smoothing parameter in the latter function was selected by a direct plug-in approach using the `dpill` function. We chose `degree = 1` in the `locpoly` function (equivalent to local linear smoothing). For the Spline and Kernel methods, we estimated the mean curve by averaging individually smoothed curves. The PACE method was implemented with a Matlab package developed by Yao et al. (2005b) (<http://www.stat.ucdavis.edu/PACE/>).

To quantitatively measure the goodness of smoothing and mean-covariance estimation, we repeated the above simulation 100 times for the common grid case. We calculated the root integrated mean square errors (RIMSEs), defined by $(\int_0^{\pi/2} (\widehat{Z}(t) - Z(t))^2 dt)^{1/2}$ for signal estimates and by $(\int_0^{\pi/2} \int_0^{\pi/2} (\widehat{\Sigma}(s, t) - \Sigma(s, t))^2 ds dt)^{1/2}$ for surface estimates, where the integrations were approximated by the trapezoidal rule. We compared the mean RIMSEs of our Bayesian estimates with the ones from alternative

Table 1: Simulation results in the common grid case: mean RIMSEs and corresponding standard errors (in parentheses) of $Z_i(\mathbf{t})$ and $\mu(\mathbf{t})$, produced by Spline, Kernel, PACE, our proposed methods (Bayesian), and BLS. The mean RIMSEs of the two best estimates are bold.

		Spline	Kernel	PACE	Bayesian	BLS
Stationary	$Z_i(\mathbf{t})$	0.4272 (0.0021)	0.4591 (0.0016)	0.4072 (0.0017)	0.3783 (0.0014)	0.3649 (0.0014)
	$\mu(\mathbf{t})$	0.3819 (0.0144)	0.3923 (0.0150)	0.3902 (0.0150)	0.3781 (0.0144)	0.3643 (0.0146)
Nonstationary	$Z_i(\mathbf{t})$	0.5176 (0.0021)	0.5453 (0.0016)	0.4614 (0.0023)	0.4599 (0.0021)	0.4137 (0.0016)
	$\mu(\mathbf{t})$	0.5616 (0.0263)	0.5738 (0.0264)	0.5802 (0.0261)	0.5539 (0.0264)	0.5418 (0.0267)

Table 2: Simulation results in the common grid case: mean RIMSEs and corresponding standard errors (in parentheses) for the estimates of $\Sigma(\mathbf{t}, \mathbf{t})$ and $Cor(\mathbf{t}, \mathbf{t})$, produced by sample estimation with raw curves (Sample), PACE method, the proposed method (Bayesian), and sample estimation with Bayesian-smoothed curves (SE-Bayesian). The mean RIMSEs of the two best estimates are bold.

	Data	Sample	PACE	Bayesian	SE-Bayesian
Stationary	$\Sigma(\mathbf{t}, \mathbf{t})$	1.5479 (0.0416)	1.2683 (0.0464)	2.2758 (0.1803)	1.2523 (0.0461)
	$Cor(\mathbf{t}, \mathbf{t})$	0.2616 (0.0052)	0.1493 (0.0065)	0.1410 (0.0066)	0.1506 (0.0068)
Nonstationary	$\Sigma(\mathbf{t}, \mathbf{t})$	2.5908 (0.0758)	2.3143 (0.0787)	2.7586 (0.0951)	2.3456 (0.0790)
	$Cor(\mathbf{t}, \mathbf{t})$	0.2625 (0.0043)	0.1661 (0.0054)	0.1695 (0.0056)	0.1716 (0.0056)

methods in the “Stationary” sections in Tables 1 and 2.

We expect the BLS method provides the smallest mean RIMSEs and standard errors of $Z_i(\mathbf{t})$ and $\mu(\mathbf{t})$ as shown in Table 1, because BLS is an “oracle” method using the true underlying distribution. The resulting mean RIMSEs and the standard errors from BLS should be treated as lower bounds that any statistical methods could achieve. We can also observe that our Bayesian approach achieves evident improvement on smoothing than the Spline/Kernel/PACE methods (with mean RIMSE 0.3783 vs. 0.4272/0.4591/0.4072). The mean curve estimate of the Bayesian method is also slightly better than the Spline/PACE/PACE methods (0.3781 vs. 0.3819/0.3923/0.3902).

In Table 2, we compared the mean RIMSEs from the Bayesian method with the ones from PACE, sample estimation with raw data (Sample), and sample estimation with Bayesian-smoothed curves (SE-Bayesian). We can observe that the SE-Bayesian

Table 3: Coverage probabilities for the 95% pointwise credible intervals of $Z_i(\mathbf{t})$, $\mu(\mathbf{t})$, and $\Sigma(\mathbf{t}, \mathbf{t})$.

Data	$Z_i(\mathbf{t})$	$\mu(\mathbf{t})$	$\Sigma(\mathbf{t}, \mathbf{t})$
Stationary (common-grid)	0.9373	1.0000	0.9983
Stationary (sparse)	0.9685	0.9875	0.9192
Nonstationary (common-grid)	0.9320	0.8750	0.9730
Nonstationary (sparse)	0.9345	0.8625	0.9788

estimate of covariance gives the lowest mean RIMSE, which is much smaller than the RIMSE of the direct Bayesian estimate (1.2523 vs. 2.2758). Here the relatively large mean RIMSE of the direct estimate from the Bayesian method may due to the non-identifiability between σ_s^2 and Ψ (the scale covariance kernel in the IWP prior of Σ). However, the direct Bayesian estimate of the correlation surface has the smallest mean RIMSE, even slightly better than the PACE estimate (0.1410 vs. 0.1493). In addition, the SE-Bayesian covariance estimate is comparable with the PACE estimate (1.2523 vs. 1.2683). This is not a surprise to us because both the proposed Bayesian method and PACE estimate a smooth covariance kernel from noisy (sparse) functional data.

In addition, we calculated the coverage probabilities for the Bayesian 95% credible intervals of $Z_i(\mathbf{t})$, $\mu(\mathbf{t})$, and $\Sigma(\mathbf{t}, \mathbf{t})$, as shown in Table 3. We can see that the coverage probabilities are above 90% for the simulation study with stationary covariance. From simulation studies, we find that the coverage probability for $\mu(\mathbf{t})$ varies along the value of c in (2.2), which controls the magnitude of the prior variance of $\mu(\cdot)$.

4.2 Simulation 2: functional data with nonstationary covariance

We generated functional data with nonstationary covariance by imposing a nonlinear transformation on the true underlying GPs in Section 4.1. Given the stationary GP $\tilde{X}_i(t)$ in Section 4.1, a nonstationary GP can be obtained by taking $X_i(t) = h(t)\tilde{X}_i(\xi(t))$ with $h(t) = t + 1/2$ and $\xi(t) = (t)^{2/3}$. As a result, $X_i(t)$ is a GP with mean $\mu(t) = 3h(t)\sin(4\xi(t))$ and covariance $\Sigma(s, t) = 5h(s)h(t)\text{Matern}_{cor}(|\xi(s) - \xi(t)|; \rho = 0.5, \nu = 3.5)$. Similarly as in Simulation 1, 50 nonstationary functional trajectories were generated on the same common grid, and noises were added to the true smoothed curves. Functional data in uncommon grid (sparse) case were generated in the same way as in Simulation 4.1.

Two sample curves for the common and uncommon grid cases are plotted in Figure 2(a, b). We followed the same method of determining the hyperparameters and calculating the posterior means of the mean curve and covariance surface as in Section 4.1. Particularly, for nonstationary data, it is desirable to choose Ψ with a more flexible nonstationary structure. In this case study, we let $\Psi(\cdot, \cdot) = \sigma_s^2 A(\cdot, \cdot)$, and set $A(\cdot, \cdot)$ to be the covariance estimate obtained from PACE. Note that $A(\cdot, \cdot)$ has to be symmetric and positive definite.

Examples of the Bayesian estimates are shown in Figure 2(a, b, c and d). We also




Figure 2: Simulation 2: (a) two sample curves in the common grid case (gray lines), superimposed by the Bayesian estimates (black solid lines), along with the 95% pointwise credible intervals (black dash-dot lines) and the true signals (blue dots); (b) curve plot (same as (a)) for the uncommon grid case; (c) heat map of our Bayesian correlation estimate in the common grid case; (d) heat map (same as (c)) in the uncommon grid case; (e) heat map of the sample correlation estimate (common grid case); (f) heat map of the true underlying (nonstationary) correlation surface.

plotted the heat map of the sample correlation in the common grid case in Figure 2(e) and the heat map of the true underlying correlation in Figure 2(f). The mean RIMSEs (from 100 repeated simulations of the common grid case) for $\{Z_i(\mathbf{t})\}$, $\mu(\mathbf{t})$, and $\Sigma(\mathbf{t}, \mathbf{t})$ are displayed in the “Nonstationary” sections of Tables 1 and 2. In addition, the coverage probabilities of the 95% Bayesian credible intervals for $\{Z_i(\mathbf{t})\}$, $\mu(\mathbf{t})$, and $\Sigma(\mathbf{t}, \mathbf{t})$ are displayed in the “Nonstationary” section of Table 3. The nonstationary results in Tables 1, 2 and 3 show similar patterns as in the stationary case studies. In particular, the proposed Bayesian approach produces the best smooth signal and mean estimates with the lowest mean RIMSEs.

5 Real case studies

We use two real datasets with different levels of noise to show that our proposed Bayesian method produces accurate smooth estimates for the functional signals and mean-covariance.

5.1 Spectroscopy data with low levels of noise

In this application, we analyzed the spectroscopy data that were produced in a cervical pre-cancer study. The goal of the study was to diagnose early-stage cervical cancer with spectroscopy measurements (Yamal et al., 2012). This data contain $n = 462$ spectroscopy measurements produced by a multi-spectral digital colposcopy (MDC) device (Buys et al., 2012). When taking the measurements, an operator puts a probe in contact with the cervical tissues; the device then ejects a beam of light through the probe onto the tissues and records the spectrum intensities of the reflected light through a white-light filter in the device. Each measurement contains intensity values (\log_{10} transformed) ranging over emission wavelengths from 410 to 700 nanometer (nm). The spectroscopy data are collected on dense emission wavelengths (a case of common grid). We took intensity values at 1/3 of the equally spaced emission wavelengths as our observations (training data), and treated the values at the remaining 2/3 wavelengths as the validation data for prediction purposes. Figure 3(a) shows three raw curves on their original scale (before the \log_{10} transformation).

We then applied our Bayesian method on the training data (\log_{10} transformed), with $\delta = 5$ and hyperpriors set by the heuristic empirical Bayes method in Section 2.2. The values of the prior parameters were taken as: $a_\epsilon = 1117.759$, $b_\epsilon = 1$, $a_s = 0.905$, $b_s = 5$, $\hat{\rho} = 369.716$, $\nu = 3.190$. The Bayesian estimates for σ_ϵ^2 and σ_s^2 were 9.339×10^{-5} with 95% credible interval $(9.199 \times 10^{-5}, 9.488 \times 10^{-5})$, and 196.655 with 95% confidence interval $(173.679, 220.693)$, respectively.

We plotted the smoothed curves by the Kernel method in Figure 3(b), the ones by the Spline method in Figure 3(c), and the ones by our Bayesian method in Figure 3(d) together with 95% pointwise credible intervals. One can observe that all local features in the raw data are completely wiped out by the kernel estimates in Figure 3(b), which may due to the over-estimation of the smoothing parameter. In contrast, both the Spline and the Bayesian methods successfully preserved these features. However, because of the inherent nature of independently smoothing in the Spline method, the amount of smoothness varies across curves. For example, in Figure 3(c), the red and green curves appear to be slightly smoother than the blue curve, while our Bayesian method produces signal estimates with about the same amount of smoothness. This real data analysis demonstrates that our proposed Bayesian method is capable of preserving shared features by borrowing strength across all observations.

Due to low levels of noise, the influence of smoothing on the covariance estimation with this spectroscopy data is almost negligible. In such a situation one would expect a good correlation estimate to have a similar structure as the sample estimate. Figure 4 shows the heat maps of the correlation estimates given by sample estimate with raw




Figure 3: Case study on spectroscopy data: (a) three raw spectroscopy curves; (b) smooth estimates by the Kernel method; (c) smooth estimates by the Spline method; (d) smooth estimates by our Bayesian method along with 95% pointwise credible intervals (black dashed curve).

curves in frame (a), Bayesian method in frame (b), PACE in frame (c) and sample estimate with Bayesian smoothed curves in frame (d). The Bayesian estimate (Figure 4(b)) and the sample estimate with Bayesian smoothed curves (Figure 4(d)) are very close to the sample estimate based on the raw data (Figure 4(a)). Although the PACE estimate in Figure 4(c) shows a similar main structure, it appears to have different details around emission wavelength 500nm (marked by red circles in Figure 4), from the other three estimates.

We used the remaining 2/3 validation data to assess the prediction performance of the Bayesian method, compared with three alternative methods – Spline, Kernel and PACE. For the Spline method, the values on the validation grid points could be easily predicted with the estimated splines coefficients. For the Kernel/PACE/Bayesian methods, predicted values at the validation points were obtained using linear interpolation. Due to low noise levels in the spectroscopy data, we simply treated the observed values as true data and calculated the RMSEs as in Table 4. Here, the Bayesian method achieves the smallest RMSE, demonstrating the benefit of




Figure 4: Heat maps of the correlation estimates by the sample estimation with raw data in (a), Bayesian method in (b), PACE in (c), and sample estimation with Bayesian smoothed data in (d).

simultaneous smoothing on prediction.

Table 4: RMSEs for predictions on the validation grids.

Spline	Kernel	PACE	Bayesian
0.020	0.039	0.025	0.017

5.2 Metabolic data with high levels of noise

We also studied a metabolic dataset from a study of obese children. The data were collected in the Children’s Nutrition Research Center (CNRC) at the Baylor College of Medicine. We considered a subset of the data that contained energy expenditure (EE) measurements of 44 obese children. These measurements were collected every minute over a 100-minute period during their sleeping time. Therefore, each trajectory can be treated as the evaluation of an energy function over time. These raw curves




Figure 5: (a): Bayesian signal estimates (black solid) with 95% credible intervals (black dashed) and raw observations (gray); (b): signal estimates by the Spline method; (c): signal estimates by the Kernel method.

appear to fluctuate frequently, implying potentially high levels of noise; three sample raw trajectories are plotted with the color gray in Figure 5(a).

We applied the proposed Bayesian method to this dataset, using the common grid model ($n < p$). We set $\delta = 5$ and set $A(\cdot, \cdot)$ as a Matérn correlation kernel, which led to signal estimates with similar level of variations (the black solid curves in Figure 5(a)). In comparison with the Bayesian method, the Spline method produces excessive local variations in all three curves (Figure 5(b)), and the Kernel smoothing method shows clear over-smoothing for all three curves (Figure 5(c)).

6 Discussion

We have proposed a model-based nonparametric Bayesian method to smooth all functional observations simultaneously and estimate the mean and covariance functions based on GPs. Specifically, we adopted a hierarchical framework by assuming a GP prior for the mean function and an IWP prior for the covariance function. The proposed method facilitates borrowing strength across all functional observations in simultaneous smoothing and an automatic mean-covariance estimation in the posterior inference.

Simulation studies show that our method produces smoothed signal estimates that are the closest to the “oracle” BLS estimates among all methods compared, and the resulting smoothed data lead to improved covariance estimation. In the real case study with spectroscopy data, our method demonstrates the most accurate prediction of function values at the validation wavelengths among all methods compared. The

real case study with metabolic data shows that the smoothed signal estimates by our approach retain consistent patterns, while the alternative methods lost systematic trends common across all curves.

In our approach, posterior inference is performed by evaluating GPs on finite grids. To evaluate them on a new grid in posterior sampling in the most convenient way, one needs to perform posterior predictive sampling based on the conditional Gaussian distribution for smoothed data and to conduct a two-dimensional interpolation for the covariance surface while restricting the interpolated surface to be symmetric and positive definite. When dealing with functional data on uncommon grids, our approach involves evaluating mean and covariance functions on the pooled grid, which can be extremely dense and cause numerical problems. For example, the covariance matrices on dense grids are likely to be singular, thus inverting these matrices in (3.2), (3.3), (3.5) would be problematic. In such a case, sampling from a multivariate normal distribution with a nearly singular covariance matrix and sampling from an Inverse-Wishart distribution (in the MCMC steps) with a nearly singular scale matrix might also be challenging. These are common numerical issues encountered in GP models due to dense grids.

Many solutions have been proposed to deal with the intensive computational burden and numerical issues caused by dense grids. The essential idea is to approximate a large matrix or its inverse using a low rank matrix. Comprehensive reviews can be found in [Rasmussen and Williams \(2006, Chapter 8.\)](#), [Quiñonero Candela et al. \(2007\)](#), and [Shi and Choi \(2011\)](#) for example. An efficient approximation method by adopting a linear random projection has also been proposed by [Banerjee et al. \(2013\)](#). In our simulations and real case studies, we employed numerical schemes such as using the technique of generalized inverse and converting a non-positive definite matrix to a positive definite matrix by replacing its non-positive eigenvalues with fairly small positive numbers, that is, the so-called *jittering*. Similar strategies have been adopted in the PACE ([Yao et al., 2005b](#)) as well. For problems involving large-scale matrix computations, more advanced low-rank approximation techniques can be integrated into our MCMC steps to further improve scalability. In terms of the computational scalability, our current MCMC algorithm is able to handle a moderately large dataset with $O(10^4)$ observational points. Alternatively, instead of representing the functional data on grids, we can also adopt basis representations for functional data with an appropriately selected basis system.

Acknowledgments

The authors would like to thank all colleagues in the PO1 project (supported by NIH grant PO1-CA-082710) for collecting the spectroscopy data and the Children’s Nutrition Research Center at the Baylor College of Medicine for providing the metabolic data (funded by National Institute of Diabetes and Digestive and Kidney Diseases Grant DK-74387 and the USDA/ARS under Cooperative Agreement 6250-51000-037). The research of Taeryon Choi was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education, Science and Technology (No. 2012R1A1A2041370). The authors are grateful for all of the comments from the editor, the associate editor, and two referees. Last but not least, the authors greatly appreciate Kirsten Herold of the Writing Lab (School of Public Health, University of Michigan) and Xiaowei Wu (Department of Statistics, Virginia Tech) for their help with editing this article.

References

- Banerjee, A., Dunson, D. B., and Tokdar, S. T. (2013). “Efficient Gaussian process regression for large datasets.” *Biometrika*, 100(1): 75–89. [2](#), [19](#)
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical modeling and analysis for spatial data*. Crc Press. [2](#)
- Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). “Gaussian predictive process models for large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4): 825–848. [2](#)
- Buys, T. P., Cantor, S. B., Guillaud, M., Adler-Storthz, K., Cox, D. D., Okolo, C., Arulogon, O., Oladepo, O., Basen-Engquist, K., Shinn, E., et al. (2012). “Optical technologies and molecular imaging for cervical neoplasia: a program project update.” *Gender medicine*, 9(1): S7–S24. [15](#)
- Cardot, H., Ferraty, F., and Sarda, P. (2003). “Spline estimators for the functional linear model.” *Statistica Sinica*, 13(3): 571–592. [2](#)
- Dawid, A. P. (1981). “Some matrix-variate distribution theory: notational considerations and a Bayesian application.” *Biometrika*, 68(1): 265–274. [5](#)
- Gelman, A. and Rubin, D. B. (1992). “Inference from iterative simulation using multiple sequences.” *Statistical science*, 457–472. [9](#)
- Hall, P., Poskitt, D. S., and Presnell, B. (2001). “A Functional DataAnalytic Approach to Signal Discrimination.” *Technometrics*, 43(1): 1–9. [2](#)
- Hitchcock, D. B., Casella, G., and Booth, J. G. (2006). “Improved estimation of dissimilarities by presmoothing functional data.” *Journal of the American Statistical Association*, 101(473): 211–222. [2](#)
- Kaufman, C. G., Sain, S. R., et al. (2010). “Bayesian functional ANOVA modeling using Gaussian process prior distributions.” *Bayesian Analysis*, 5(1): 123–149. [2](#)
- Leng, X. and Müller, H.-G. (2006). “Classification using functional data analysis for temporal gene expression data.” *Bioinformatics*, 22: 68–76. [3](#)
- Müller, H.-G. (2005). “Functional modeling and classification of longitudinal data.” *Scandinavian J. Statistics*, 32: 223–240. [3](#)
- Nguyen, X. and Gelfand, A. E. (2014). “Bayesian nonparametric modeling for functional analysis of variance.” *Annals of the Institute of Statistical Mathematics*, 66(3): 495–526. [2](#)
- Quiñonero Candela, J., E., R. C., and Williams, C. K. I. (2007). “Approximation Methods for Gaussian Process Regression.” Technical report, Applied Games, Microsoft Research Ltd. [19](#)
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
URL <http://www.R-project.org/> [10](#)

- Ramsay, J. O. and Dalzell, C. (1991). “Some tools for functional data analysis.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 539–572. 2
- Ramsay, J. O. and Silverman, B. W. (2002). *Applied functional data analysis: methods and case studies*, volume 77. Springer New York. 2
- (2005). *Functional data analysis*. Springer Series in Statistics. Springer, New York, second edition. 2, 3
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA. 2, 19
- 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*, 53: 233–243. 3
- Särkkä, S. and Aki, V. (2014). “MCMC Diagnostics for Matlab.” <http://becs.aalto.fi/en/research/bayes/mcmcdiag/>. 9
- Shi, J. Q. and Choi, T. (2011). *Gaussian process regression analysis for functional data*. CRC Press, Boca Raton, FL. 2, 19
- Shi, J. Q., Wang, B., Will, E. J., and West, R. M. (2012). “Mixed-effects Gaussian process functional regression models with application to dose-response curve prediction.” *Stat. Med.*, 31(26): 3165–3177. 2
- Stein, M. L. (1999). *Interpolation of spatial data: some theory for kriging*. Springer. 5
- Von Neumann, J. (1941). “Distribution of the ratio of the mean square successive difference to the variance.” *The Annals of Mathematical Statistics*, 12(4): 367–395. 7
- Wang, B. and Shi, J. Q. (2014). “Generalized Gaussian process regression model for non-Gaussian functional data.” *J. Amer. Statist. Assoc.*, 109(507): 1123–1133. 2
- Yamal, J.-M., Zewdie, G. A., Cox, D. D., Atkinson, E. N., Cantor, S. B., MacAulay, C., Davies, K., Adewole, I., Buys, T. P., and Follen, M. (2012). “Accuracy of optical spectroscopy for the detection of cervical intraepithelial neoplasia without colposcopic tissue information; a step toward automation for low resource settings.” *Journal of biomedical optics*, 17(4): 047002–047002. 15
- Yao, F., Müller, H.-G., and Wang, J.-L. (2005a). “Functional Data Analysis for Sparse Longitudinal Data.” *Journal of the American Statistical Association*, 100(470): 577–590. 3
- (2005b). “Functional linear regression analysis for longitudinal data.” *The Annals of Statistics*, 33(6): 2873–2903. 3, 7, 11, 19
- Zhang, H. (2004). “Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics.” *Journal of the American Statistical Association*, 99(465): 250–261. 5
- Zhu, H. and Cox, D. D. (2009). “A functional generalized linear model with curve selection in cervical pre-cancer diagnosis using fluorescence spectroscopy.” In

Optimality, volume 57 of *IMS Lecture Notes Monogr. Ser.*, 173–189. Inst. Math. Statist., Beachwood, OH. 2

Zhu, H., Strawn, N., and Dunson, B. D. (2014). “Bayesian graphical models for multivariate functional data.” *arXiv:1411.4158 [stat.ME]*.
URL <http://arxiv.org/abs/1411.4158> 5

Zhu, H., Vannucci, M., and Cox, D. D. (2010). “A Bayesian Hierarchical Model for Classification with Selection of Functional Predictors.” *Biometrics*, 66: 463–473. 2