

## Ultra-Efficient MCMC for Bayesian Longitudinal Functional Data Analysis

Thomas Y. Sun & Daniel R. Kowal

**To cite this article:** Thomas Y. Sun & Daniel R. Kowal (22 Jul 2024): Ultra-Efficient MCMC for Bayesian Longitudinal Functional Data Analysis, *Journal of Computational and Graphical Statistics*, DOI: [10.1080/10618600.2024.2362227](https://doi.org/10.1080/10618600.2024.2362227)

**To link to this article:** <https://doi.org/10.1080/10618600.2024.2362227>



[View supplementary material](#)



Published online: 22 Jul 2024.



[Submit your article to this journal](#)



Article views: 81



[View related articles](#)



CrossMark

[View Crossmark data](#)

# Ultra-Efficient MCMC for Bayesian Longitudinal Functional Data Analysis

Thomas Y. Sun<sup>a</sup>  and Daniel R. Kowal<sup>a,b</sup> 

<sup>a</sup>Department of Statistics, Rice University, Houston, TX; <sup>b</sup>Department of Statistics and Data Science, Cornell University, Ithaca, NY

## ABSTRACT

Functional mixed models are widely useful for regression analysis with dependent functional data, including longitudinal functional data with scalar predictors. However, existing algorithms for Bayesian inference with these models only provide either scalable computing or accurate approximations to the posterior distribution, but not both. We introduce a new MCMC sampling strategy for highly efficient and fully Bayesian regression with longitudinal functional data. Using a novel blocking structure paired with an orthogonalized basis reparameterization, our algorithm jointly samples the fixed effects regression functions together with all subject- and replicate-specific random effects functions. Crucially, the joint sampler optimizes sampling efficiency for these key parameters while preserving computational scalability. Perhaps surprisingly, our new MCMC sampling algorithm even surpasses state-of-the-art algorithms for frequentist estimation and variational Bayes approximations for functional mixed models—while also providing accurate posterior uncertainty quantification—and is orders of magnitude faster than existing Gibbs samplers. Simulation studies show improved point estimation and interval coverage in nearly all simulation settings over competing approaches. We apply our method to a large physical activity dataset to study how various demographic and health factors associate with intraday activity. Supplementary materials for this article are available online.

## ARTICLE HISTORY

Received June 2023  
Accepted May 2024

## KEYWORDS

Actigraphy data;  
Function-on-scalar  
regression; Gibbs sampler;  
Mixed models

## 1. Introduction

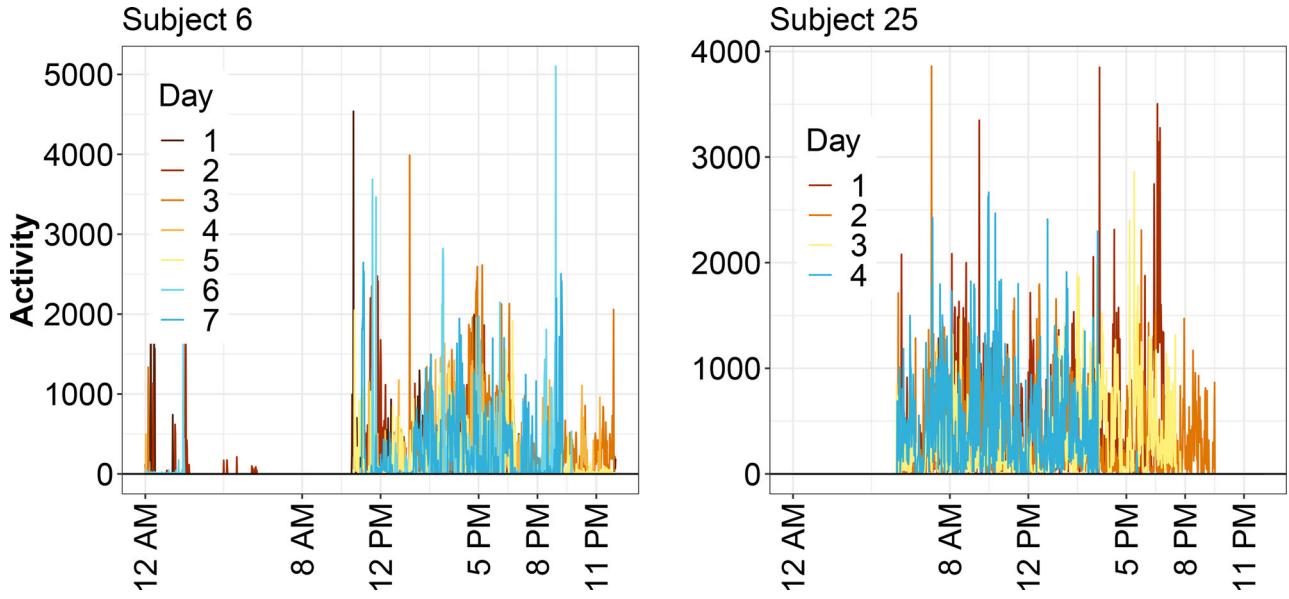
Functional data analysis (FDA) refers to the statistical analysis of data objects observed over a continuum, such as time or space, typically at high resolutions. FDA has been applied in a variety of important areas, including climate data (Besse, Cardot, and Stephenson 2000), electricity prices (Liebl 2013), COVID-19 dynamics (Boschi et al. 2021), and many others. FDA is particularly challenging when the functional data are dependent, which requires more sophisticated statistical models and more intensive computations. We focus on regression analysis with *longitudinal functional data*, which presents the simultaneous challenges of (i) within-curve dependencies, (ii) groupings among repeated (functional) measurements, and (iii) associations with (possibly many) scalar covariates.

To meet these challenges, *functional mixed models* (FMMs) have emerged as a powerful modeling tool. FMMs combine FDA and traditional mixed effects models to provide regression analysis of functional data in the presence of structured dependence. However, longitudinal functional datasets are often massive, with millions or billions of high-resolution measurements (Doherty et al. 2017; Lee et al. 2019). Thus, recent work has increasingly prioritized the *computational scalability* of FMMs. Naturally, the attendant computational burdens are exacerbated by the complexity of the statistical model, including various covariance structures and (possibly many) scalar covariates. As a result, there is high demand for inferential algorithms for FMMs

that both (i) scale to massive datasets and (ii) maintain these essential modeling capabilities.

These challenges are exemplified in our motivating application, the 2005–2006 National Health and Nutrition Examination Survey (NHANES) physical activity (PA) dataset. High resolution (minute-by-minute) PA data (Figure 1) are recorded across multiple days for each participant using hip-worn accelerometry devices and linked to a questionnaire that contains subject-level demographic and health information (Table 2). Figure 1 shows the PA levels of two randomly chosen subjects for every day they wore the device. Key aspects of the data become evident from this plot: the observations are high-resolution, noisy, autocorrelated, and exhibit considerable variation both between subjects and between days for a given subject. Notably, there are more than 72 million PA measurements on more than 10,000 participants in the study (Leroux et al. 2019). The goal is to link time-of-day PA with important health and demographic variables, while appropriately accounting for the prominent and complex dependencies among these high-dimensional data.

To enable statistical modeling and inference, we represent the PA as longitudinal functional data: the PA measurements are *functions* of time-of-day and the days are *repeated measurements* of these functional data for each subject. More generally, let  $Y_{i,j}(\tau)$  denote functional observations on a compact domain  $\tau \in \mathcal{T}$  (i.e., time-of-day) for within-subject replications  $j = 1, \dots, m_i$  (i.e., days) and subjects  $i = 1, \dots, n$ ,



**Figure 1.** Minute-by-minute physical activity (PA) measurements for two subjects in the NHANES study. The PA observations are noisy, autocorrelated, and exhibit considerable variation both between subjects and between days for a given subject.

with  $M = \sum_{i=1}^n m_i$ . The functional data are linked to  $L$  scalar covariates  $\mathbf{x}_{i,j} = (x_{i,j,1}, \dots, x_{i,j,L})'$  for subject  $i$  and replicate  $j$ . We study the following FMM for longitudinal functional regression:

$$Y_{i,j}(\tau) = \tilde{\alpha}_0(\tau) + \sum_{\ell=1}^L x_{i,j,\ell} \tilde{\alpha}_\ell(\tau) + \tilde{\gamma}_i(\tau) + \tilde{\omega}_{i,j}(\tau) + \epsilon_{i,j}(\tau), \\ \epsilon_{i,j}(\tau) \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2). \quad (1)$$

Model (1) extends traditional mixed effects models for longitudinal data to the functional data setting by allowing the *fixed effects*  $\{\tilde{\alpha}_\ell(\cdot)\}$  and the *random effects*  $\{\tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$  to vary as functions of  $\tau$ . First, the fixed effects  $\{\tilde{\alpha}_\ell(\cdot)\}$  are regression coefficient functions and describe the linear associations between  $\mathbf{x}_{i,j}$  and  $Y_{i,j}(\tau)$  at each  $\tau$ . Next, the random effects  $\{\tilde{\gamma}_i(\cdot)\}$  are shared among all functional observations for subject  $i$ , and thus account for subject-specific effects and dependencies among repeated (functional) observations. The replicate-specific random effects  $\{\tilde{\omega}_{i,j}(\cdot)\}$  account for within-curve dependencies (or “smooth” errors) that are unexplained by the fixed effects or subject-specific random effects. These effects are especially important for functional regression, even without repeated measurements (Reiss, Huang, and Mennes 2010; Kowal and Bourgeois 2020). Finally, the errors  $\{\epsilon_{i,j}(\cdot)\}$  describe any (non-smooth) measurement errors that remain. The random effects functions and errors  $\{\tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot), \epsilon_{i,j}(\cdot)\}$  are mean zero and mutually independent, with additional modeling assumptions discussed subsequently.

The FMM (1) is widely applicable and has appeared in many previous studies (Zipunnikov et al. 2014; Cederbaum et al. 2016; Zhu et al. 2019; Li et al. 2022), including both special cases without replications ( $m_i = 1$ ; Guo 2002) or covariates (Park and Staicu 2015) as well as generalizations for additive (Scheipl, Staicu, and Greven 2015) or non-Gaussian responses (Scheipl, Gertheiss, and Greven 2016). However, computational considerations are paramount, and often preclude the use of these generalized models even for moderately-sized datasets (Sergazinov et al. 2023). In response, Cui et al. (2022) proposed Fast

Univariate Inference (FUI) for (1), which seeks to dramatically simplify estimation by (i) fitting separate linear mixed models pointwise for each  $\tau$ , (ii) applying a smoother to the estimated pointwise fixed effects, and (iii) using asymptotic arguments or bootstrapping to obtain confidence bands for  $\{\tilde{\alpha}_\ell(\cdot)\}$ . Such a deconstruction sacrifices estimation efficiency and is difficult to apply for sparsely or irregularly sampled functional data. Like other frequentist approaches, FUI requires selection of tuning (smoothing) parameters and provides limited uncertainty quantification for certain parameters and predictions. Yet most important, we show subsequently that such a decomposition is *not* necessary to achieve scalable computing—and in fact that FUI is slower than our fully Bayesian approach (see Section 4).

In general, Bayesian approaches for FMMs are highly appealing due to the consolidated interpretation of fixed and random effects, as well as convenient uncertainty quantification of all model parameters and predictions. Morris and Carroll (2006) introduced Bayesian FMMs using a wavelet basis expansion for fixed and random effects. This approach has been extended and applied broadly (Morris et al. 2006; Zhu, Brown, and Morris 2011; Morris 2017; Lee et al. 2019), but relies on unique features of the wavelet basis—which may not be suitable for smoother functional data, cannot easily handle missing values, and is difficult to apply when the number of functional observations is non- dyadic. Bayesian FMMs commonly rely on Markov chain Monte Carlo (MCMC) sampling for posterior inference, typically using Gibbs sampling (Morris and Carroll 2006; Goldsmith and Kitago 2016; Lee et al. 2019) or Hamiltonian Monte Carlo (Goldsmith, Zipunnikov, and Schrack 2015). Advantageously, these algorithms provide exact (up to Monte Carlo error) inference, but are prohibitively slow, with running times in the hours or days even for moderately sized datasets (see Section 4). Further, certain MCMC sampling strategies are vulnerable to slow mixing and convergence: for example, Goldsmith and Kitago (2016) applied a Gibbs sampler that alternates sampling blocks for the fixed effects and the random effects, which is sensitive to the model

parameterization (i.e., centered vs. noncentered; Yu and Meng 2011) and can lead to poor exploration of the joint target distribution. Our analyses confirm the MCMC inefficiencies of this blocking strategy, which compounds the impact of the extremely lengthy computing times.

To reduce this computational burden, there has been recent development for variational Bayes (VB) approximations for Bayesian FMMs. VB substantially reduces computation times compared to existing MCMC algorithms (Goldsmith and Kitago 2016; Huo, Morris, and Zhu 2022). However, VB often provides poor uncertainty quantification, and thus Goldsmith and Kitago (2016) recommend it primarily as a tool to obtain quick initial estimates for model-building. Our simulations (Section 4) show that VB is adequate for point estimation, but falls considerably short in interval estimation. Thus, there is urgent demand for algorithms that can provide both *accurate* and *scalable* Bayesian inference for the FMM (1).

We address this significant gap in the literature. Specifically, we design an MCMC sampling algorithm for the FMM (1) that offers several unique features. First, we jointly sample *all* fixed and random effects functions  $\{\tilde{\alpha}_\ell(\cdot), \tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$ , which delivers superior Monte Carlo efficiency for these critical quantities. Second, we design this joint sampler using a careful blocking structure paired with an orthogonalized basis reparameterization, which leads to exceptionally fast sampling steps. The accompanying variance components are sampled in a separate block, yielding a convenient Gibbs sampler that is both *computationally* and *Monte Carlo* efficient. Perhaps surprisingly, our new MCMC sampling algorithm even surpasses state-of-the-art algorithms for frequentist estimation (Cui et al. 2022; Li et al. 2022) and VB approximations (Goldsmith and Kitago 2016) and is orders of magnitude faster than existing Gibbs samplers for FMMs. This superior scaling persists across number of subjects  $n$ , number of replicates  $m_i$  per subject, and number of covariates  $L$ , and is accompanied by more accurate point estimation and uncertainty quantification across nearly all tested scenarios. Applying our methods to the NHANES PA dataset, we demonstrate the significant practical impacts of our improved MCMC algorithm and provide posterior uncertainty quantification for key parameters. Compared to existing MCMC algorithms, we reduce the computation time from two weeks to only a few minutes.

The rest of this article is organized as follows. Section 2 presents the Bayesian FMM for longitudinal data. Section 3 describes the MCMC algorithm. Section 4 provides simulation analyses. We apply our model on the NHANES dataset in Section 5. We conclude with a discussion in Section 6. R code to implement our approach is available on GitHub at <https://github.com/thomasyun/FLFOSR>.

## 2. Basis Expansions and Prior Distributions

FMMs (1) are most commonly implemented using basis expansions, including splines (Guo 2002; Goldsmith and Kitago 2016), functional principal components (Park and Staicu 2015; Li et al. 2022), and wavelets (Morris and Carroll 2006; Huo, Morris, and Zhu 2022). The general basis representation of (1) is

$$Y_{i,j}(\tau) = \sum_{k=1}^K b_k(\tau) \beta_{k,i,j} + \epsilon_{i,j}(\tau), \quad \epsilon_{i,j}(\tau) \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2) \quad (2)$$

$$\beta_{k,i,j} = \alpha_{k,0} + \sum_{\ell=1}^L x_{i,j,\ell} \alpha_{k,\ell} + \gamma_{k,i} + \omega_{k,i,j} \quad (3)$$

where  $\{b_k(\cdot)\}_{k=1}^K$  are known basis functions and  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$  are unknown basis coefficients. Model (2)–(3) induces (1) under the identification  $\tilde{\alpha}_0(\tau) = \sum_{k=1}^K b_k(\tau) \alpha_{k,0}$  and similarly for the remaining fixed and random effects functions. The coefficients  $\{\beta_{k,i,j}\}$  are completely determined by  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$ , and merely serve as a placeholder for notational convenience. Because the basis functions  $\{b_k\}$  are known, estimation and inference for the coefficients  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$  is sufficient for estimation and inference of the functions  $\{\tilde{\alpha}_\ell(\cdot), \tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$ .

The choice of basis functions  $\{b_k\}$  must be paired carefully with the choice of prior on the coefficients  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$ , which together induce a prior for functions  $\{\tilde{\alpha}_\ell(\cdot), \tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$ . Typically, we select the prior on the coefficients to encourage certain properties for the functions, such as smoothness (Goldsmith, Zipunnikov, and Schrack 2015; Goldsmith and Kitago 2016; Kowal and Bourgeois 2020) or sparsity (Morris and Carroll 2006). Here, we also prioritize the resulting computational implications.

To motivate the general approach, suppose that we are interested in specifying a prior for a generic function under a basis expansion  $\tilde{\zeta}(\tau) = \sum_{k=1}^K b_k(\tau) \zeta_k$ . A common strategy is to assume a prior of the form  $\zeta = (\zeta_1, \dots, \zeta_K)' \sim N_K(\mathbf{0}, \sigma_\zeta^2 \mathbf{P}^-)$ , where  $\mathbf{P}$  is a known (roughness) penalty matrix and  $\sigma_\zeta^2$  is a (smoothness) parameter. For instance, suppose  $\mathbf{P}$  is the matrix of integrated squared second derivatives,  $[\mathbf{P}]_{k,k} = \int b_k''(\tau) b_k''(\tau) d\tau$ , where  $b_k''(\cdot)$  denotes the second derivative of  $b_k(\cdot)$ . Then the prior for  $\zeta$  may be expressed as

$$-2 \log p(\zeta \mid \sigma_\zeta^2) \stackrel{c}{=} \sigma_\zeta^{-2} \zeta' \mathbf{P} \zeta = \sigma_\zeta^{-2} \int \{\tilde{\zeta}''(\tau)\}^2 d\tau$$

and  $\stackrel{c}{=}$  denotes equality up to an additive constant. Thus, the log-prior on  $\zeta$  reproduces the classical roughness penalty on the function  $\tilde{\zeta}(\cdot)$ . In a Bayesian framework, we further may place a prior on  $\sigma_\zeta^2$  to learn the smoothness parameter; these details are discussed below in the context of model (2)–(3).

Although this approach is advantageous due to its generality and smoothness-inducing properties, it is unfavorable for computing, especially within the basis-expanded FMM (2)–(3). Suppose we observe data at  $\tau_1, \dots, \tau_T \in \mathcal{T}$  and let  $\mathbf{B}_0 = (\mathbf{b}_1, \dots, \mathbf{b}_K)$  be the  $T \times K$  basis matrix with  $\mathbf{b}_k = (b_k(\tau_1), \dots, b_k(\tau_T))'$ . We proceed generally with *any* basis matrix  $\mathbf{B}_0$  and penalty matrix  $\mathbf{P}$ , with the minimal requirement that  $\mathbf{P}^-$  is positive semidefinite. We seek to reparameterize these terms  $(\mathbf{B}_0, \mathbf{P})$ —at a one-time cost for all subsequent MCMC sampling—for more amenable computing. Specifically, we apply the orthogonalization strategy from Scheipl, Fahrmeir, and Kneib (2012), which uses the spectral decomposition of  $\mathbf{B}_0 \mathbf{P}^- \mathbf{B}_0'$  to form a basis matrix  $\mathbf{B}$  such that (i)  $\mathbf{B}' \mathbf{B} = \text{diag}(\{d_k\}_{k=1}^K)$  is a diagonal matrix and (ii)  $(\tilde{\zeta}(\tau_1), \dots, \tilde{\zeta}(\tau_T))' = \mathbf{B}_0 \zeta$  has the same distribution as  $\mathbf{B} \zeta^*$ , where  $\zeta^* \sim N(\mathbf{0}, \sigma_\zeta^2 \mathbf{I})$ . Here  $Y_{i,j}(\tau)$  and the unknown coefficient functions are assumed to lie in the same function space and are thus representable under the same set of basis functions. Notably, this strategy applies for any basis  $\{b_k\}$  and penalty matrix  $\mathbf{P}^-$ ; in our empirical

examples, we use B-splines with a penalty on the second-order differences of the coefficients.

Revisiting the FMM basis expansion (2)–(3), we henceforth assume, without loss of generality, that  $\mathcal{B}'\mathcal{B}$  is diagonal and that arbitrary (roughness) penalties may be incorporated via independent Gaussian priors. Thus, we specify the priors

$$\begin{aligned} [\alpha_{k,\ell} \mid \sigma_{\alpha_\ell}^2] &\stackrel{\text{indep}}{\sim} N(0, \sigma_{\alpha_\ell}^2), \quad [\gamma_{k,i} \mid \sigma_\gamma^2] \stackrel{\text{indep}}{\sim} N(0, \sigma_\gamma^2), \\ [\omega_{k,i,j} \mid \sigma_{\omega_i}^2] &\stackrel{\text{indep}}{\sim} N(0, \sigma_{\omega_i}^2) \end{aligned} \quad (4)$$

and assume that the reparameterization has already been completed. The variance parameters are shared across  $k$ , but this can be relaxed to allow adaptive shrinkage for the basis coefficients (Gao and Kowal 2024). In conjunction with the basis expansions, the priors (4) induce a Gaussian process prior for each fixed and random effect function, for example,

$$\tilde{\alpha}_0 \sim \mathcal{GP}(0, C_{\alpha_0}), \quad C_{\alpha_0}(\tau, u) = \sigma_{\alpha_0}^2 \sum_{k=1}^K b_k(\tau) b_k(u)$$

and similarly for the remaining terms.

Finally, we assume conditionally conjugate Gamma priors for the precision parameters,

$$\{\sigma_{\alpha_\ell}^{-2}, \sigma_\gamma^{-2}, \sigma_{\omega_i}^{-2}\} \stackrel{\text{iid}}{\sim} \text{Gamma}(a, b) \quad (5)$$

along with a Jeffreys' prior for the observation error variance,  $[\sigma_\epsilon^2] \propto 1/\sigma_\epsilon^2$ . These parameters determine both the smoothness of the corresponding functions as well as the various sources of variability within a function, between replicates for a subject, and between subjects. Other priors are available for variance parameters (Gelman 2006), and may be substituted into the proposed framework with minimal impact on the core MCMC sampling approach. However, we find that (5) offers excellent modeling performance and MCMC efficiency, and is not highly sensitive to the choice of  $(a, b)$  (see the supplementary material for a sensitivity analysis).

In practice, we observe the functional data  $Y_{i,j}(\tau)$  at discrete points  $\tau_1, \dots, \tau_T \in \mathcal{T}$ . For simplicity, we assume that these observation points are common for all  $(i, j)$ , but it is straightforward to accommodate sparsely or irregularly sampled functional data within a Bayesian framework (Kowal, Matteson, and Ruppert 2019). The likelihood implied under (2) with observed data  $\mathbf{Y}_{i,j} = (Y_{i,j}(\tau_1), \dots, Y_{i,j}(\tau_T))'$  is

$$Y_{i,j} = \mathcal{B}\beta_{i,j} + \epsilon_{i,j}, \quad \epsilon_{i,j} \stackrel{\text{iid}}{\sim} N_T(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I}_T) \quad (6)$$

where  $\beta_{i,j} = (\beta_{i,j,1}, \dots, \beta_{i,j,K})'$  is modeled via (3)–(4).

We emphasize that the modeling choices in (2)–(5) are meant to be broadly applicable—including generic basis expansions, penalty matrices, and variance components (or smoothness parameters) for each fixed and random effects function. The main contributions of this paper are not found in the uniqueness of this modeling strategy, but rather in our MCMC sampling algorithm for the general Bayesian FMM (1), which is presented in the next section. The essential features of our FMM specification are (i) the basis expansions (2) and (ii) the conditionally Gaussian priors (4), in conjunction with the aforementioned basis orthogonalization strategy (Scheipl, Fahrmeir, and Kneib 2012).

### 3. MCMC Algorithm

The primary challenge is to construct a sampling algorithm for the joint posterior distribution of the fixed and random effects functions  $\{\tilde{\alpha}_\ell(\cdot), \tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$  under the FMM (1). This task requires simultaneous consideration of (i) MCMC efficiency, including convergence and autocorrelation of the Markov chain, and (ii) computational scalability across all dimensions of the data: the number of subjects  $n$ , the number of replicates  $m_i$  per subject, number of covariates  $L$ , and the number of observation points  $T$  along each curve. We focus on the sampling steps for the fixed and random effects basis coefficients  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$ , which are sufficient for posterior inference for the fixed and random effects functions  $\{\tilde{\alpha}_\ell(\cdot), \tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$  under (2)–(3). The variance components  $\{\sigma_{\alpha_\ell}^2, \sigma_\gamma^2, \sigma_{\omega_i}^2\}$  are sampled in a separate block, resulting in a two-block Gibbs sampling algorithm (Algorithm 1).

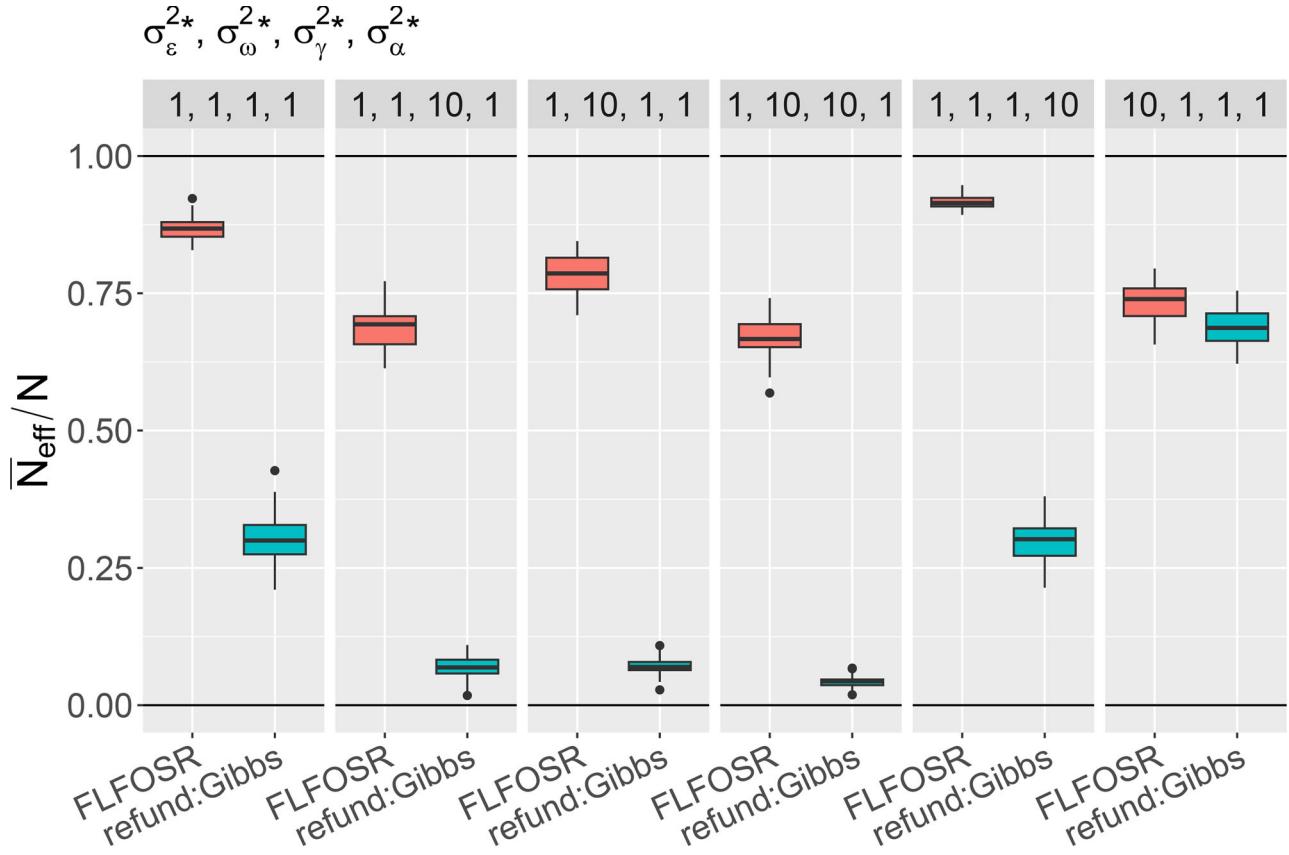
We propose a *joint* sampler for *all* fixed and random effects functions  $\{\tilde{\alpha}_\ell(\cdot), \tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$  by carefully decomposing the joint posterior distribution of the corresponding basis coefficients. First, let  $\boldsymbol{\alpha}_k = (\alpha_{k,0}, \alpha_{k,1}, \dots, \alpha_{k,L})'$ ,  $\boldsymbol{\gamma}_k = (\gamma_{k,1}, \dots, \gamma_{k,n})'$ , and  $\boldsymbol{\omega}_k = (\omega_{k,1}', \dots, \omega_{k,n}')'$ , where  $\boldsymbol{\omega}_{k,i} = (\omega_{k,i,1}, \dots, \omega_{k,i,m_i})'$ , and let  $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)'$  and similarly for  $\boldsymbol{\gamma}$  and  $\boldsymbol{\omega}$ . Similarly, define the diagonal variance matrices  $\boldsymbol{\Sigma}_\epsilon = \sigma_\epsilon^2 \mathbf{I}_M$ ,  $\boldsymbol{\Sigma}_\alpha = \text{diag}(\{\sigma_{\alpha_\ell}^2\}_{\ell=1}^L)$ ,  $\boldsymbol{\Sigma}_\gamma = \text{diag}(\{\sigma_\gamma^2\}_{i=1}^n) = \sigma_\gamma^2 \mathbf{I}_n$ , and  $\boldsymbol{\Sigma}_\omega = \text{diag}(\{\sigma_{\omega_i}^2\}_{i,j})$ , where the set of variances are first iterated through  $j = 1, \dots, m_i$  for  $i = 1$  and so on to yield  $M$  terms along the diagonal. Let  $\boldsymbol{\Sigma} = \{\boldsymbol{\Sigma}_\alpha, \boldsymbol{\Sigma}_\gamma, \boldsymbol{\Sigma}_\omega, \boldsymbol{\Sigma}_\epsilon\}$  contain all variance components.

The joint fixed and random effects posterior (conditional on  $\boldsymbol{\Sigma}$ ) is decomposable as

$$p(\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\Sigma}) = p(\boldsymbol{\alpha} \mid \mathbf{Y}, \boldsymbol{\Sigma}) p(\boldsymbol{\gamma} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}) p(\boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}) \quad (7)$$

where  $\mathbf{Y} = \{Y_{i,j}\}$  denotes all observed data. Our Gibbs sampler iterates between joint sampling blocks from (7) and  $[\boldsymbol{\Sigma} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}]$ ; the latter sampling step is straightforward (see Algorithm 1), so we focus on (7).

We sample from the joint posterior (7) by iteratively drawing from the three constituent terms: (i) the marginal posterior of the fixed effects  $\boldsymbol{\alpha}$ , (ii) the partial conditional posterior of the subject-specific random effects  $\boldsymbol{\gamma}$  given the fixed effects  $\boldsymbol{\alpha}$ , and (iii) the full conditional posterior of the replicate-specific random effects  $\boldsymbol{\omega}$ . In contrast to a Gibbs sampler that draws from the three *full conditional* distributions  $[\boldsymbol{\alpha} \mid \mathbf{Y}, \boldsymbol{\gamma}, \boldsymbol{\omega}, \boldsymbol{\Sigma}]$ ,  $[\boldsymbol{\gamma} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\Sigma}]$ , and  $[\boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}]$  (e.g., Goldsmith and Kitago 2016), or Gibbs samplers that sequentially draw each fixed effect  $\alpha_k$  (e.g. Morris and Carroll 2006), our approach delivers a direct Monte Carlo (not MCMC) sample from the joint posterior (conditional on  $\boldsymbol{\Sigma}$ ), and thus offers the potential for large gains in sampling efficiency (see Figure 2). Further, the joint sampler requires no consideration of centered versus non-centered parameterizations of the mixed effects model, which eliminates a recurring nuisance for Bayesian implementations of mixed and hierarchical models (Yu and Meng 2011). However, we must also carefully consider the resulting computational burden of these sampling steps. In particular, a sampling algorithm that offers Monte Carlo efficiency may nonetheless be infeasible



**Figure 2.** Average relative efficiency  $\bar{N}_{\text{eff}}/N$  for FLFOSR and `refund:Gibbs` across various simulation designs (columns) and 30 simulated datasets (boxplots). The proposed FLFOSR sampler shows consistently excellent and superior MCMC efficiency, while `refund:Gibbs` is extremely inefficient when  $\sigma_{\omega}^{2*}$  or  $\sigma_{\gamma}^{2*}$  is large.

in practice, if the raw computing times scale poorly in  $n$ ,  $m_i$ ,  $L$ , or  $T$ .

Our key innovation is that we provide convenient, closed-form, and highly scalable sampling steps for the constituent distributions in (7). First, observe that  $\{\beta_{k,i,j}\}$  in (2)–(3) contains all fixed and random effects coefficients, and thus may be viewed as a placeholder for each of the random variables in (7). Since the reparameterized basis matrix satisfies  $\mathbf{B}'\mathbf{B} = \text{diag}(\{d_k\}_{k=1}^K)$ , we compute  $(\mathbf{B}'\mathbf{B})^{-1}\mathbf{B}'\mathbf{Y}_{i,j} = (y_{1,i,j}, \dots, y_{K,i,j})'$ , where  $\{y_{k,i,j}\}$  are the coefficients from projecting the functional observations onto the basis matrix. Thus  $\left\| \mathbf{Y}_{i,j} - \mathbf{B}\beta_{i,j} \right\|_2^2 = \sum_{k=1}^K d_k (y_{k,i,j} - \beta_{k,i,j})^2$  and the joint likelihood (6) for  $\{\beta_{k,i,j}\}$  may be written

$$p(\mathbf{Y} \mid \{\beta_{k,i,j}\}, \Sigma) \propto \prod_{i=1}^n \prod_{j=1}^{m_i} \exp \left\{ -\frac{1}{2\sigma_{\epsilon}^2} \left\| \mathbf{Y}_{i,j} - \mathbf{B}\beta_{i,j} \right\|_2^2 \right\} \quad (8)$$

$$\propto \exp \left\{ -\frac{1}{2\sigma_{\epsilon}^2} \sum_{i=1}^n \sum_{j=1}^{m_i} \sum_{k=1}^K d_k (y_{k,i,j} - \beta_{k,i,j})^2 \right\} \quad (9)$$

up to constants that do not depend on  $\{\beta_{k,i,j}\}$ . Thus, the only dependence on  $T$  is via this projection step, which is a one-time cost for all subsequent MCMC sampling. In addition, the joint posterior factorizes across the basis coefficients,

$$p(\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega} \mid \mathbf{Y}, \Sigma) = \prod_{k=1}^K p(\boldsymbol{\alpha}_k, \boldsymbol{\gamma}_k, \boldsymbol{\omega}_k \mid \mathbf{Y}, \Sigma),$$

and we may sample the fixed and random effects coefficients separately for each  $k = 1, \dots, K$ . This strategy is parallelizable across  $k$  yet still maintains a joint sampler for all fixed and random effects coefficients. This offers a substantial simplification from the functional data likelihood in (8), which features  $T$ -dimensional terms  $\mathbf{Y}_{i,j}$ , especially since  $K \ll T$  in general. Finally, the likelihood (9) for  $\{\beta_{k,i,j}\}$  is proportional to the likelihood for  $\{\beta_{k,i,j}\}$  implied by the simple *working model*

$$y_{k,i,j} = \beta_{k,i,j} + \epsilon_{k,i,j}, \quad \epsilon_{k,i,j} \stackrel{\text{indep}}{\sim} N(0, \sigma_{\epsilon}^2/d_k) \quad (10)$$

$$= \alpha_{k,0} + \sum_{\ell=1}^L x_{i,j,\ell} \alpha_{k,\ell} + \gamma_{k,i} + \omega_{k,i,j} + \epsilon_{k,i,j}. \quad (11)$$

It is typically easier to derive the posterior distributions for  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$  in (7) using (10)–(11) instead of (6) or (8), yet the results are equivalent.

We describe the three sampling steps in ascending complexity, which reverses the actual sampling order from (7). Let  $\mathbf{y}_k = (y_{k,1,1}, \dots, y_{k,1,m_1}, \dots, y_{k,n,m_n})'$  denote the projected functional data ordered by replicates for each subject,  $\mathbf{X}$  be the  $M \times (L+1)$  design matrix with a column of ones for the intercept, and  $\mathbf{Z} = \text{bdiag} \{ \mathbf{1}_{m_i} \}_{i=1}^n$ , a block diagonal matrix with  $n$  columns and  $m_i$ -dimensional vectors of ones.

The simplest sampling step is the full conditional distribution for the replicate-specific coefficients:

$$\begin{aligned} [\boldsymbol{\omega}_k \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}] &\stackrel{\text{indep}}{\sim} N_M(\mathbf{Q}_{\omega_k}^{-1} \boldsymbol{\ell}_{\omega_k}, \mathbf{Q}_{\omega_k}^{-1}) \\ \mathbf{Q}_{\omega_k} &= \text{diag}(\{d_k \sigma_{\epsilon}^{-2} + \sigma_{\omega_i}^{-2}\}_{ij}) \\ \boldsymbol{\ell}_{\omega_k} &= d_k \sigma_{\epsilon}^{-2} \{ \mathbf{y}_k - (\mathbf{X} \boldsymbol{\alpha}_k + \mathbf{Z} \boldsymbol{\gamma}_k) \} \end{aligned} \quad (12)$$

for  $k = 1, \dots, K$ . Most important, the posterior precision matrix  $\mathbf{Q}_{\omega_k}$  is diagonal, and thus all the replicate-specific coefficients  $\boldsymbol{\omega}$  may be sampled independently (and in parallel) in  $\mathcal{O}(MK)$  computing time.

Next, we provide the sampling step for the partial conditional of the subject-specific random effects,  $[\boldsymbol{\gamma}_k \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}]$ , which marginalizes over  $\boldsymbol{\omega}$ . Under the prior (4), it is easy to see that the marginalized version of (11) is simply  $y_{k,i,j} = \alpha_{k,0} +$

$\sum_{\ell=1}^L x_{i,j,\ell} \alpha_{k,\ell} + \gamma_{k,i} + v_{k,i,j}$ , where  $v_{k,i,j} = \omega_{k,i,j} + \epsilon_{k,i,j} \stackrel{\text{indep}}{\sim} N(0, \sigma_{\omega_i}^2 + \sigma_{\epsilon}^2/d_k)$ . The requisite distribution is then

$$\begin{aligned} [\boldsymbol{\gamma}_k \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}] &\stackrel{\text{indep}}{\sim} N_n(\mathbf{Q}_{\gamma_k}^{-1} \boldsymbol{\ell}_{\gamma_k}, \mathbf{Q}_{\gamma_k}^{-1}) \\ \mathbf{Q}_{\gamma_k} &= \boldsymbol{\Sigma}_{\gamma}^{-1} + \mathbf{Z}' \mathbf{V}_k \mathbf{Z} = \text{diag}(\{(\sigma_{\gamma}^{-2} + d_k m_i (\sigma_{\epsilon}^2 + d_k \sigma_{\omega_i}^2)^{-1})\}_{i=1}^n) \\ \boldsymbol{\ell}_{\gamma_k} &= d_k \mathbf{Z}' \mathbf{V}_k (\mathbf{y}_k - \mathbf{X} \boldsymbol{\alpha}_k) \end{aligned} \quad (13)$$

where  $\mathbf{V}_k = \boldsymbol{\Sigma}_{\epsilon}^{-1} - d_k \boldsymbol{\Sigma}_{\epsilon}^{-1} (d_k \boldsymbol{\Sigma}_{\omega}^{-1} + \boldsymbol{\Sigma}_{\epsilon}^{-1})^{-1} \boldsymbol{\Sigma}_{\epsilon}^{-1} = \text{diag}(\{(\sigma_{\epsilon}^2 + d_k \sigma_{\omega_i}^2)^{-1}\}_{i,j})$  for  $k = 1, \dots, K$ . Crucially, the posterior precision matrix  $\mathbf{Q}_{\gamma_k}$  is diagonal, and thus all subject-specific coefficients  $\boldsymbol{\gamma}$  may be sampled independently (and in parallel) in  $\mathcal{O}(nK)$  computing time.

While the posterior precision matrices for  $[\boldsymbol{\omega}_k \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}]$  and  $[\boldsymbol{\gamma}_k \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}]$  are diagonal and easy to construct, we must also consider computation of the vectors  $\boldsymbol{\ell}_{\omega_k}$  and  $\boldsymbol{\ell}_{\gamma_k}$ . An important observation is that, for arbitrary matrix  $\Theta_{n \times n}$ , the operations  $\mathbf{Z} \Theta_{n \times n}$  (or  $\Theta_{n \times n} \mathbf{Z}'$ ) repeat row  $i$  (or column  $i$ ) of  $\Theta_{n \times n} m_i$  times, while  $\mathbf{Z}' \Theta_{M \times M}$  (or  $\Theta_{M \times M} \mathbf{Z}$ ) sums over the columns (or rows) for all rows (or columns) corresponding to the same subject  $i$ . The matrix notation obscures the simplicity of these computations, which are vectorizable and thus offer additional efficiency gains in R.

The remaining sampling step is the marginal posterior of the fixed effects,  $[\boldsymbol{\alpha}_k \mid \mathbf{Y}, \boldsymbol{\Sigma}]$ , which requires marginalization over *all* random effects coefficients  $\{\boldsymbol{\gamma}, \boldsymbol{\omega}\}$ . Once more, the working model (10)–(11) is particularly convenient for these derivations, and yields the distribution

$$\begin{aligned} [\boldsymbol{\alpha}_k \mid \mathbf{Y}, \boldsymbol{\Sigma}] &\stackrel{\text{indep}}{\sim} N_L(\mathbf{Q}_{\alpha_k}^{-1} \boldsymbol{\ell}_{\alpha_k}, \mathbf{Q}_{\alpha_k}^{-1}) \\ \mathbf{Q}_{\alpha_k} &= \boldsymbol{\Sigma}_{\alpha}^{-1} + d_k \mathbf{X}' (\mathbf{V}_k - \mathbf{W}_k) \mathbf{X} \\ \boldsymbol{\ell}_{\alpha_k} &= d_k \mathbf{X}' (\mathbf{V}_k - \mathbf{W}_k) \mathbf{y}_k \end{aligned} \quad (14)$$

where  $\mathbf{W}_k = (\mathbf{Z}' \mathbf{V}_k)' (\boldsymbol{\Sigma}_{\gamma}^{-1} + d_k \mathbf{Z}' \mathbf{V}_k \mathbf{Z})^{-1} \mathbf{Z}' \mathbf{V}_k$  and  $\mathbf{V}_k - \mathbf{W}_k = \text{diag}(\{(\sigma_{\epsilon}^2 + d_k \sigma_{\omega_i}^2 + d_k m_i \sigma_{\gamma}^2)^{-1}\}_{i,j})$  for  $k = 1, \dots, K$ . Unlike for the random effects coefficients, the posterior precision for  $\boldsymbol{\alpha}_k$  is not diagonal; indeed,  $\mathbf{Q}_{\alpha_k}$  is the *only* non-diagonal posterior precision matrix in our joint sampler for all  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$ . To alleviate this potential bottleneck when the number of covariates  $L$  is large, we apply the  $\mathcal{O}(L^3)$  sampling algorithm from Rue (2001) when  $L \leq M$  and the  $\mathcal{O}(M^2 L)$  sampling algorithm from Bhattacharya, Chakraborty, and Mallick (2016) when  $L > M$ . This sampling strategy has been successful in (non-longitudinal) function-on-scalars regression (Kowal and Bourgeois 2020). More generally, (14)

has the same structure as in (nonfunctional) Bayesian linear regression models, and thus we may adapt and apply new sampling strategies for that setting as the state-of-the-art advances (Nishimura and Suchard 2022).

In aggregate, we present our full MCMC sampling algorithm in [Algorithm 1](#). The joint sampling step for  $[\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\Sigma}]$  is highly scalable, and in fact delivers the same computational cost as a naive Gibbs sampler that instead use the three full conditional draws  $[\boldsymbol{\alpha} \mid \mathbf{Y}, \boldsymbol{\gamma}, \boldsymbol{\omega}, \boldsymbol{\Sigma}]$ ,  $[\boldsymbol{\gamma} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\Sigma}]$ , and  $[\boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}]$ —while offering the potential for substantial increases in MCMC efficiency.

---

**Algorithm 1:** MCMC sampling algorithm for the Bayesian FMM (1)–(6): *fast longitudinal function-on-scalar regression* (FLFOSR).

---

1. Sample the fixed and random effects coefficients  $[\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\Sigma}]$ :
    - (a) Sample  $[\boldsymbol{\alpha} \mid \mathbf{Y}, \boldsymbol{\Sigma}]$  from (14)
    - (b) Sample  $[\boldsymbol{\gamma} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}]$  from (13)
    - (c) Sample  $[\boldsymbol{\omega} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}]$  from (12)

and compute functions  $\tilde{\alpha}_\ell(\cdot) = \sum_{k=1}^K b_k(\cdot) \alpha_{k,\ell}$  and similarly for all  $\{\tilde{\gamma}_i(\cdot), \tilde{\omega}_{i,j}(\cdot)\}$ .
  2. Sample the variance components  $[\boldsymbol{\Sigma} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}]$ :
    - (a) Sample  $[\sigma_{\epsilon}^{-2} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}] \sim \text{Gamma}(MT/2, \sum_{i=1}^n \sum_{j=1}^{m_i} \|\mathbf{Y}_{i,j} - \mathbf{B} \boldsymbol{\beta}_{i,j}\|_2^2/2)$
    - (b) Sample  $[\sigma_{\alpha_\ell}^{-2} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}] \sim \text{Gamma}(a + K/2, b + \sum_{k=1}^K \alpha_{k,\ell}^2/2)$ ,  $\ell = 1, \dots, L$
    - (c) Sample  $[\sigma_{\gamma}^{-2} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}] \sim \text{Gamma}(a + nK/2, b + \sum_{k=1}^K \sum_{i=1}^n \gamma_{k,i}^2/2)$
    - (d) Sample  $[\sigma_{\omega_i}^{-2} \mid \mathbf{Y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}] \sim \text{Gamma}(a + m_i K/2, b + \sum_{k=1}^K \sum_{j=1}^{m_i} \omega_{k,i,j}^2/2)$ ,  $i = 1, \dots, n$ .
- 

From a practical perspective, we emphasize that many computations are parallelizable, including the sampling steps for all  $\{\boldsymbol{\gamma}, \boldsymbol{\omega}\}$ , while many constituent terms are highly vectorizable. For instance,  $\mathbf{V}_k$  and  $\mathbf{W}_k$  are diagonal, and computing  $(\mathbf{V}_k - \mathbf{W}_k) \mathbf{X}$  simply consists of multiplying each row of the matrix  $\mathbf{X}$  by the corresponding scalar value  $(\sigma_{\epsilon}^2 + d_k \sigma_{\omega_i}^2 + d_k m_i \sigma_{\gamma}^2)^{-1}$ . Within R, we maximize computational performance by avoiding matrix calculations wherever possible and opting for vectorized operations instead. This is in contrast to common FMM formulations that utilize matrix operations with Kronecker products or large block matrices, and thus face significant computational bottlenecks.

Lastly, we reiterate that [Algorithm 1](#) is applicable for *any* Bayesian FMM that satisfies a basis expansion (2)–(3) and assumes (conditionally) Gaussian priors (4). Thus, our approach is broadly applicable for many choices of basis functions, and remains valid for generic smoothness priors using the basis orthogonalization strategy of Scheipl, Fahrmeir, and Kneib (2012). We henceforth refer to our approach as *fast longitudinal function-on-scalar regression* (FLFOSR).

## 4. Simulations

We conduct a series of simulation studies to evaluate the performance of our approach against state-of-the-art Bayesian and frequentist competitors. Specifically, we assess (i) the MCMC efficiency and computational costs and (ii) the estimation accuracy and uncertainty quantification under various simulation designs.

### 4.1. Simulation Design

Functional observations are generated from the FMM (1) using the basis representation (2)–(3) using  $K^* = 5$  basis functions. We fix  $\alpha_{k,0}^* = 1$  and simulate the coefficients  $\alpha_{k,\ell}^* \sim N(0, \sigma_\alpha^{2*})$ ,  $\gamma_{k,i}^* \sim N(0, \sigma_\gamma^{2*})$ , and  $\omega_{k,i,j}^* \sim N(0, \sigma_\omega^{2*})$  independently, and then compute  $\beta_{k,i,j}^* = \alpha_{k,0}^* + \sum_{\ell=1}^L x_{i,\ell} \alpha_{k,\ell}^* + \gamma_{k,i}^* + \omega_{k,i,j}^*$ . The covariates  $\{x_{i,\ell}\}$  were generated from a standard normal distribution. The variance components  $\{\sigma_\alpha^{2*}, \sigma_\gamma^{2*}, \sigma_\omega^{2*}, \sigma_\epsilon^{2*}\}$  are varied in Sections 4.3 and 4.4 to evaluate performance under different sources of variability. Finally, functional data  $Y_{i,j}(\tau_t)$  are generated from (2), where the basis functions are orthogonalized B-spline basis functions and  $\{\tau_t\}_{t=1}^T$  is a grid of  $T = 144$  equally-spaced points in  $[0, 1]$ , which emulates the PA data in Section 5. We emphasize that although this data-generating process resembles the proposed Bayesian FMM, the same model structure is shared among all competing methods, and thus does not unfairly favor our approach.

### 4.2. Competing Methods

We compare the proposed FLFOSR approach against several Bayesian and frequentist methods for regression analysis with longitudinal functional data. The Bayesian competitors come from the widely-used `refund` package in R (Goldsmith et al. 2021), and feature both the Gibbs sampler (`refund:Gibbs`) and the VB approximation (`refund:VB`) from Goldsmith and Kitago (2016). Unlike FLFOSR, `refund:Gibbs` (i) uses the full conditional draws for all fixed and random effects instead of the joint sampler (Algorithm 1), (ii) does not orthogonalize the basis functions, and (iii) sets the hyperparameters of the variance components based on an initial estimate of the residual covariance matrix. We expect that (i) will inhibit MCMC efficiency and (ii) will decrease computational scalability.

Among frequentist approaches, there is a rich collection of recent strategies and software. We focus on FUI (Cui et al. 2022), which specifically advertises computational scalability for longitudinal function-on-scalar regression, and Li et al. (2022), which provides fixed effects inference for a variety of longitudinal correlation structures. We choose to work with the exchangeable correlation model since this mirrors our model assumptions. More general methods for FMMs or functional additive mixed models could be considered for this problem. However, the widely-used implementations (i.e., `pffr` in the `refund` package) are known to have severe bottlenecks for both computing and memory (Cui et al. 2022; Sergazinov et al. 2023). Thus, such methods were not included in the study.

For FLFOSR, we set the hyperparameters to be  $a = b = 0.1$  and assess sensitivity in the supplementary material. All

methods used  $K = 15$  basis functions. Any remaining settings were fixed at the default choices in the provided R functions. Simulations were performed on a Windows desktop with a 3.60 GHz Intel Xeon CPU with 32 GB of RAM.

### 4.3. Evaluating MCMC Efficiency and Computational Scalability

Among Bayesian methods, we measure MCMC efficiency using effective sample sizes (ESS), which equivalently summarize the autocorrelation in the Markov chain (Gelman et al. 2013). For each MCMC sampler, we generate  $N = 1000$  draws from a single chain after discarding the initial  $N_{\text{burn}} = 1000$  draws as burn-in, and then compute the pointwise ESS  $N_{\text{eff}}\{\tilde{\alpha}_\ell(\tau_t)\}$  of each fixed effect regression coefficient function  $\tilde{\alpha}_\ell(\cdot)$  for  $\ell = 1, \dots, L$  and  $t = 1, \dots, T$ . We also record  $s_N$  and  $s_{N_{\text{burn}}}$ , the total running times (in seconds) to obtain  $N$  and  $N_{\text{burn}}$  draws, respectively.

First, we compute the *average relative efficiency* across all covariates and observation points,

$$\bar{N}_{\text{eff}}/N = (LT)^{-1} \sum_{\ell=1}^L \sum_{t=1}^T N_{\text{eff}}\{\tilde{\alpha}_\ell(\tau_t)\}/N,$$

which exclusively measures MCMC efficiency (not computing time). We compare this metric between the competing MCMC samplers, FLFOSR and `refund:Gibbs`, using the simulation design from Section 4.1. Figure 2 presents results across various settings for the variance components  $\{\sigma_\alpha^{2*}, \sigma_\gamma^{2*}, \sigma_\omega^{2*}, \sigma_\epsilon^{2*}\}$  and  $n = 20$ ,  $m = 5$ , and  $L = 5$  (a smaller dataset is necessary for `refund:Gibbs`; see Figure 3). FLFOSR is dramatically more efficient than `refund:Gibbs`, with consistently excellent performance across all setting. By comparison, the MCMC efficiency of `refund:Gibbs` deteriorates significantly whenever  $\sigma_\gamma^{2*}$  or  $\sigma_\omega^{2*}$  is large. This result is unsurprising: for Gibbs samplers that alternate between drawing the fixed and the random effects from their respective full conditional distributions (e.g., `refund:Gibbs`), it is well-known that the model parameterization (centered vs. noncentered) is critical for MCMC efficiency (Yu and Meng 2011). In particular, the performance depends on the relative magnitudes of the variance components, which in practice are unknown. By comparison, FLFOSR samples the fixed and random effects functions *jointly*, and thus requires no consideration of the parameterization—and achieves excellent MCMC efficiency across a variety of settings.

Next, we summarize aggregate computational performance by measuring the time needed to generate 1000 *effective* samples, averaged over all covariates and observation points,

$$\bar{s}_{1000} = (LT)^{-1} \sum_{\ell=1}^L \sum_{t=1}^T \left[ s_{N_{\text{burn}}} + s_N \times \frac{1000}{N_{\text{eff}}\{\tilde{\alpha}_\ell(\tau_t)\}} \right]$$

which rewards both MCMC efficiency and computational scalability. This quantity allows some comparison between Bayesian and frequentist algorithms: although an MCMC algorithm may be run arbitrarily long to secure greater accuracy, 1000 effective samples is a reasonable and conservative target for general sampling-based inference.

Using the simulation design from [Section 4.1](#), we scale the size of simulated datasets across three separate dimensions: (i) the number of subjects  $n \in \{10, 20, 50, 100, 200\}$ , fixing  $m = 5$  and  $L = 5$ ; (ii) the number of replicates or repeated observations per subjects,  $m_i = m \in \{5, 10, 25, 50, 100, 150\}$ , fixing  $n = 10$  and  $L = 5$ ; and (iii) the number of predictors  $L \in \{5, 10, 25, 33, 50, 100, 200\}$ , fixing  $n = 30$  and  $m = 5$ . Data were generated using  $\sigma_\alpha^{2*} = \sigma_\gamma^{2*} = \sigma_\omega^{2*} = 1$ , and  $\sigma_\epsilon^{2*} = 10$ . The computing times for each algorithm are averaged across 30 datasets in each simulation setting.

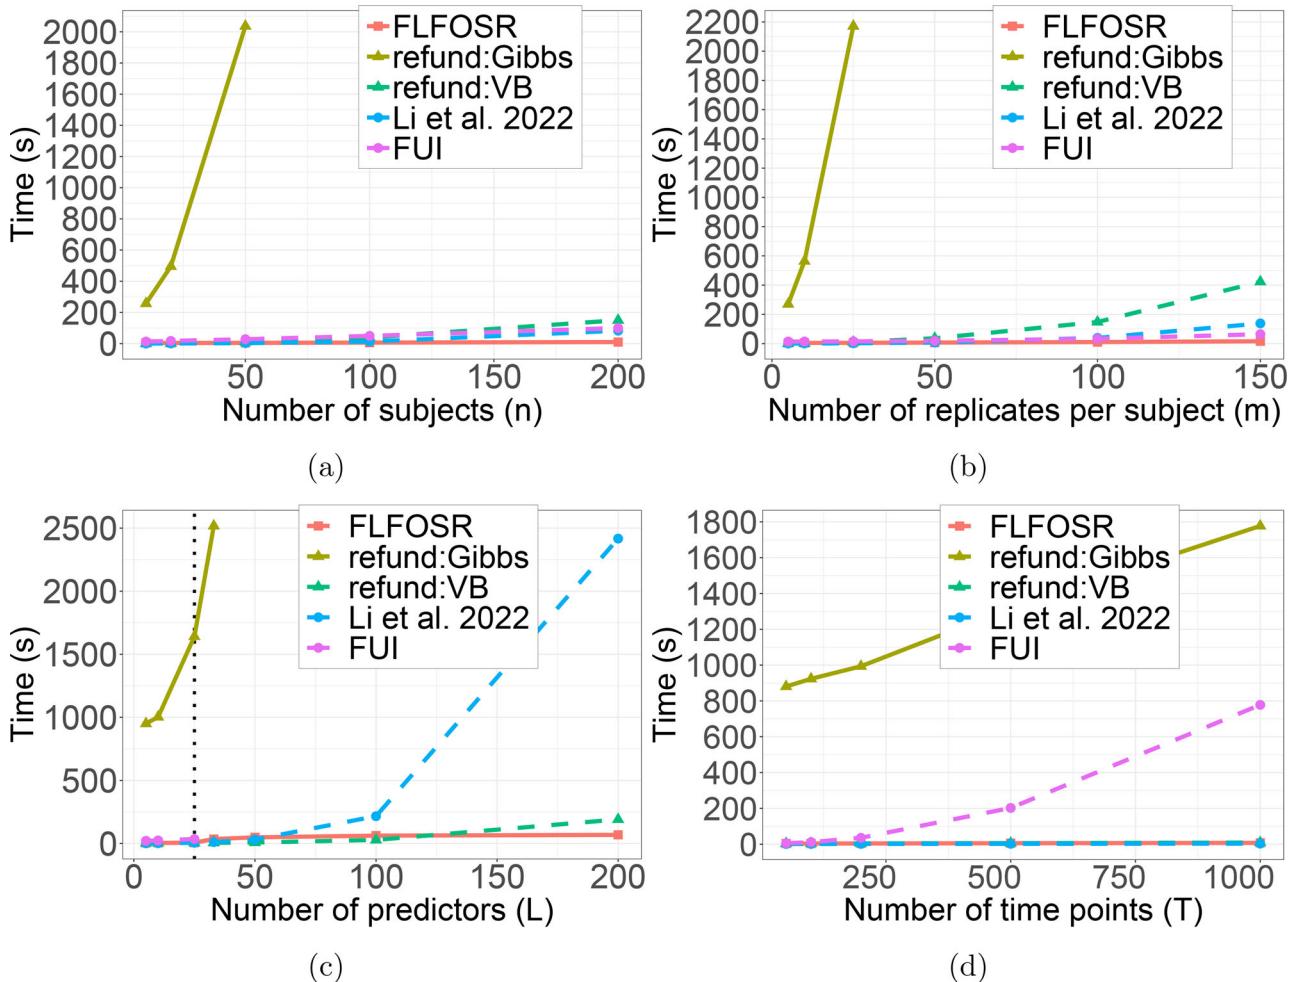
The computational performance of each method is summarized in [Figure 3](#) and detailed further in [Table 1](#). Once again, FLFOSR decisively outperforms the MCMC competitor (`refund:Gibbs`): FLFOSR delivers more efficient exploration of the posterior while reducing computational costs by several orders of magnitude. These results highlight the mutual importance of (i) the *joint* sampling step for  $\{\alpha_{k,\ell}, \gamma_{k,i}, \omega_{k,i,j}\}$  and (ii) the careful construction of fast constituent sampling steps ([Algorithm 1](#)). For instance, FLFOSR generates 1000 average effective samples in under 10 sec for a dataset with 144,000 data points ( $T = 144$ ,  $n = 200$ ,  $m = 5$ ). By comparison, `refund:Gibbs` requires several minutes for only  $n = 10$  subjects, and is not feasible even for moderately-sized datasets.

Yet most surprisingly, FLFOSR attains 1000 effective samples in *less time* than it takes for either the VB approximation or state-of-the-art frequentist algorithms ([Li et al. 2022](#); [Cui et al. 2022](#)) to terminate. These computing gains accelerate as each dimension  $n$ ,  $m$ , or  $L$  increases. Further, FUI scales poorly in  $L$  and requires  $L < n$ , while FLFOSR has no such restrictions. These results were obtained from a single machine without using any parallel processing. Thus, further scalability is attainable for FLFOSR.

These results are extraordinarily favorable: Bayesian inference for [\(1\)](#) is especially appealing due to the availability of full posterior uncertainty quantification for all fixed and random effects functions and predictions, yet is often eschewed in favor of frequentist methods due to computational limitations. FLFOSR obviates this tradeoff, and converts the computational performance from a disadvantage into a clear advantage for Bayesian FMMs.

#### 4.4. Evaluating Model Accuracy and Uncertainty Quantification

We assess the accuracy of each method by evaluating point and interval estimation for the fixed effect functions  $\{\tilde{\alpha}_\ell(\cdot)\}$ . We



**Figure 3.** Run time comparison of algorithms for longitudinal function-on-scalars regression, scaling by (a) the number of subjects  $n$ , (b) number of repeated observations per subject  $m_i$ , and (c) number of scalar predictors  $L$ , reported in seconds. Empirical computing time for MCMC methods (solid lines) refers to average time to 1000 effective samples,  $\bar{s}_{1000}$ , and other methods (dashed lines) report raw run time. Computing times were averaged across 30 replicate simulations. The proposed FLFOSR outperforms all competitors as each dimension grows. `refund:Gibbs` was omitted for some designs due to lack of competitiveness and FUI requires  $L < n$  (dotted black line in (c)).

**Table 1.** Evaluating algorithm performance via time to 1000 effective samples ( $\bar{s}_{1000}$ ), average relative efficiency ( $\bar{N}_{\text{eff}}/N$ ), and raw computing time (s) (in seconds).

$n$	$m$	$L$	$T$	FLFOSR		refund:Gibbs		refund:VB		Li et al.	FUI
				$\bar{s}_{1000}$	$\bar{N}_{\text{eff}}/N$	$\bar{s}_{1000}$	$\bar{N}_{\text{eff}}/N$	$s$	$s$	$s$	$s$
10	5	5	144	4.5	0.59	257.1	0.57	0.7	0.4	12.9	
50	5	5	144	4.5	0.86	2036.4	0.79	9.6	2.9	28.1	
200	5	5	144	9.9	0.9	—	—	149.9	82.6	99.6	
10	10	5	144	4.5	0.65	564	0.65	2	0.7	13.9	
10	50	5	144	6.9	0.75	—	—	36.3	7.2	19	
10	150	5	144	16	0.79	—	—	422.7	137.9	64.2	
30	5	5	144	4	0.81	950.2	0.75	3.9	1.2	20.3	
30	5	50	144	48.5	0.42	—	—	7.2	30.7	—	
30	5	200	144	68.2	0.46	—	—	190.6	2417	—	
30	5	5	50	3.7	0.80	880.6	0.73	3.6	1.1	4.9	
30	5	5	200	4.2	0.80	993.6	0.76	4.3	1.5	35.6	
30	5	5	1000	7.4	0.80	1777.2	0.76	9	3.2	778.1	

NOTE: The proposed MCMC algorithm (FLFOSR) is orders of magnitude faster than the MCMC competitor (refund:Gibbs) and even surpasses state-of-the-art VB and some frequentist competitors (Li et al. 2022; Cui et al. 2022) as each dimension grows. refund:Gibbs was omitted for some designs due to lack of competitiveness and FUI requires  $L < n$ .

measure point estimation accuracy using root mean squared error for the fixed effect functions,

$$\text{RMSE} = \sqrt{\frac{1}{LT} \sum_{\ell=1}^L \sum_{t=1}^T \{\hat{\alpha}_\ell(\tau_t) - \tilde{\alpha}_\ell^*(\tau_t)\}^2}$$

where  $\hat{\alpha}_\ell(\tau_t)$  and  $\tilde{\alpha}_\ell^*(\tau_t)$  are the estimated and true fixed effects functions, respectively, for the  $\ell$ th predictor at time  $\tau_t$ . For Bayesian methods, the estimator  $\hat{\alpha}_\ell(\tau_t)$  is the posterior mean. We evaluate uncertainty quantification using mean credible/confidence interval widths,

$$\text{MCIW} = \frac{1}{LT} \sum_{\ell=1}^L \sum_{t=1}^T \{\tilde{\alpha}_\ell^{(.975)}(\tau_t) - \tilde{\alpha}_\ell^{(.025)}(\tau_t)\}$$

paired with empirical coverage probability,

$$\text{ECP} = \frac{1}{LT} \sum_{\ell=1}^L \sum_{t=1}^T \mathbb{I}\{\tilde{\alpha}_\ell^{(.025)}(\tau_\ell) \leq \tilde{\alpha}_\ell^*(\tau_t) \leq \tilde{\alpha}_\ell^{(.975)}(\tau_t)\},$$

where  $(\tilde{\alpha}_\ell^{(.025)}(\tau_\ell), \tilde{\alpha}_\ell^{(.975)}(\tau_\ell))$  are 95% credible/confidence intervals. Ideal performance is achieved by nominal coverage  $\text{ECP} \geq 0.95$  (calibration) and small MCIW (sharpness).

We study the impact of differing sources of variability by varying the variance components  $\{\sigma_\alpha^{2*}, \sigma_\gamma^{2*}, \sigma_\omega^{2*}, \sigma_\epsilon^{2*}\}$  between 1 and 10 and set  $n = 20$ ,  $m = 5$ , and  $L = 5$ . The results are summarized in Figure 4. Broadly, the proposed FLFOSR is highly competitive across all scenarios for both point and interval estimation. Point estimation accuracy is comparable across all methods when between-subject variability is low ( $\sigma_\gamma^{2*} = 1$ ), but FLFOSR offers substantial gains in point estimation accuracy as between-subject variability increases ( $\sigma_\gamma^{2*} = 10$ ). Notably, FLFOSR achieves close to nominal coverage across all simulation designs, while refund:Gibbs, refund:VB, and Li et al. (2022) suffer from significant undercoverage. Clearly, the VB approximation is reasonably accurate for point estimation but unreliable for uncertainty quantification. FUI is the only competitor that achieves close to nominal coverage, yet the FUI intervals are much wider than the FLFOSR intervals and thus sacrifices some power to detect important effects. Further, FUI is the least accurate point estimator when  $\sigma_\gamma^{2*} = 10$ . When

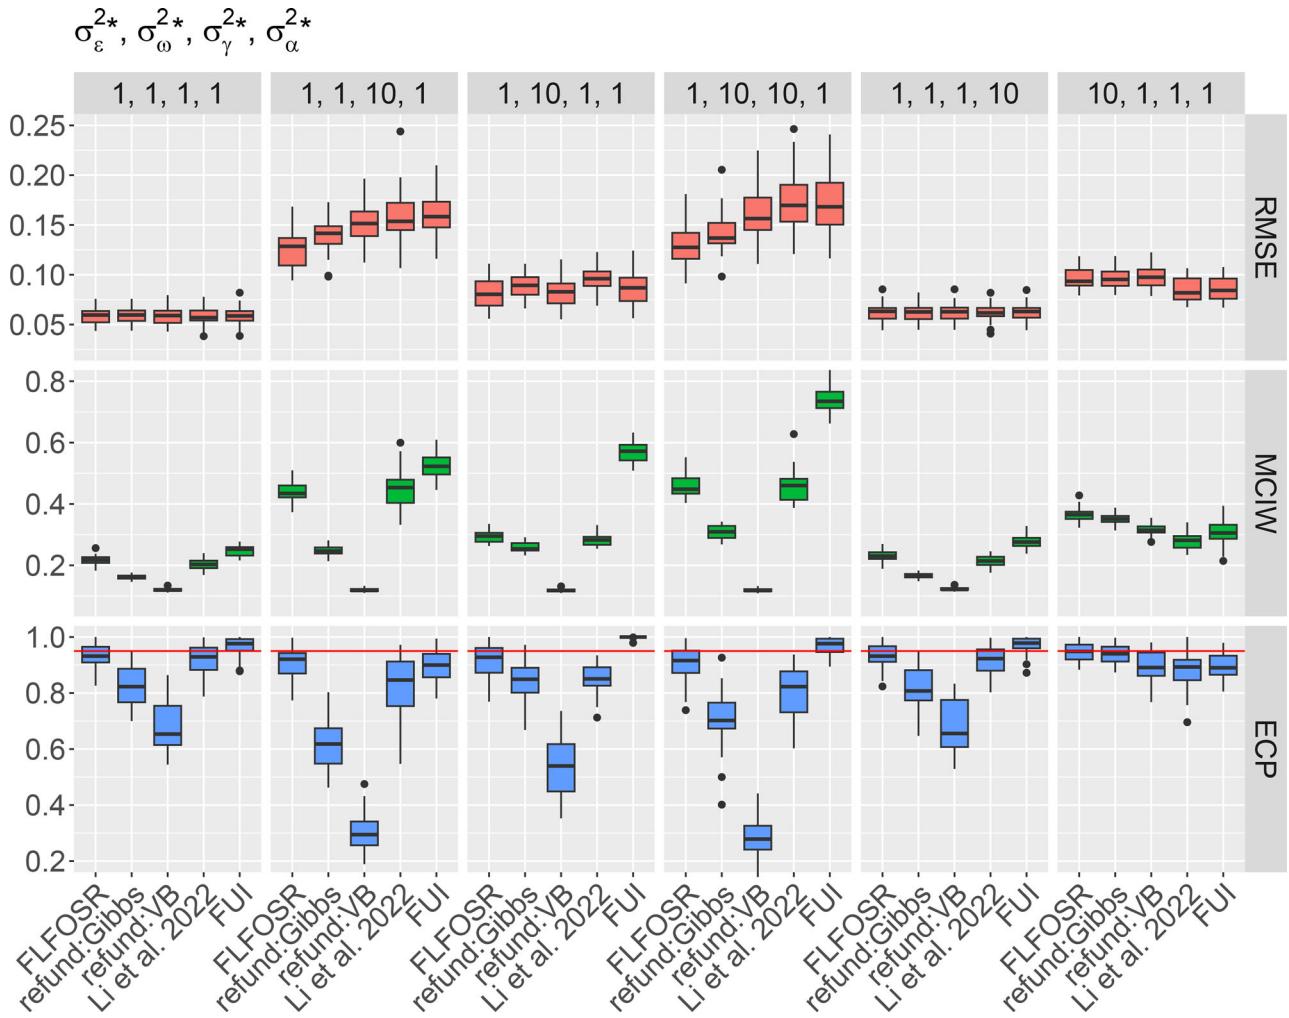
measurement error dominates ( $\sigma_\epsilon^{2*} = 10$ ), the Bayesian point estimates—including for FLFOSR—are less accurate. In such high-noise settings, this performance might be improved by replacing the priors (4) with more aggressive shrinkage priors (Gao and Kowal 2024).

In aggregate, the computational (Figures 2–3) and inferential (Figure 4) evaluations decisively favor FLFOSR. The proposed MCMC sampling algorithm (Algorithm 1) is more scalable than existing frequentist methods and more efficient than existing MCMC samplers, and delivers point and interval estimates that outperform state-of-the-art Bayesian and frequentist competitors. These significant gains amplify the additional benefits of Bayesian inference for FMMs, including uncertainty quantification for all fixed and random effects functions and predictions—which was previously not accessible even for moderately-sized datasets. By comparison, the competing frequentist approaches provide inference only for the fixed effects functions, and not for the random effects functions or predictions (Li et al. 2022; Cui et al. 2022).

## 5. Application

We apply FLFOSR to analyze physical activity (PA) data measured from wearable accelerometry devices. These devices provide high resolution PA measurements for individual subjects across multiple days. PA plays a major role in overall human health, with lower PA levels linked to higher all-cause mortality (Schmid, Ricci, and Leitzmann 2015; Smirnova et al. 2020). However, statistical analysis of PA data often includes substantial aggregation or averaging both within a day and across days for a given subject (Kowal 2022b; Hilden et al. 2023). Such pre-processing needlessly reduces the sample size and can obscure important sources of variability. In response, functional data analysis has become increasingly useful to analyze and model PA data (Morris et al. 2006; Sera et al. 2017; Leroux et al. 2019; Kowal and Bourgeois 2020; Kowal 2022a). Naturally, such approaches must confront the significant computational challenges associated with high resolution data for many individuals.

Using data from the 2005–2006 NHANES cohort, we model PA as a function of time-of-day  $\tau$  on days  $j = 1, \dots, m_i$  for subjects  $i = 1, \dots, n$ . Importantly, this representation allows



**Figure 4.** Root mean squared errors (top row), mean credible/confidence interval widths (middle row), and empirical coverage probabilities (bottom row; horizontal line denotes nominal coverage) for the fixed effects functions  $\tilde{\alpha}_\ell(\cdot)$  across various designs (columns) and 30 simulated datasets (boxplots). The proposed FLFOSR consistently delivers accurate point estimation and precise and well-calibrated uncertainty quantification.

for analysis of time-of-day PA patterns via a FMM (1), while also accounting for within-day autocorrelations, within-subject dependencies, and measurement errors. The raw activity data contains the acceleration in the x, y, and z-axis measured every 1/80th of a second by the accelerometer, aggregated at the minute level. We follow the pre-processing procedures outlined in Leroux et al. (2019) and the accompanying `nhanesdata` R package, which removes subjects with poor data quality or too few days of measured activity, and restrict our analysis to subjects aged 35–85. The functional data  $Y_{i,j}(\tau)$  are defined by computing the square-root of 10-min PA averages over the window  $\mathcal{T} = [4:00, 23:59]$ , which focuses on the waking hours and helps satisfy the FMM assumptions (1). These functional data are paired with important health and demographic variables (Table 2). The continuous variables are standardized to have mean 0 and standard deviation 1. The final dataset has  $n = 1723$  subjects with  $\text{median}(\{m_i\}_{i=1}^n) = 6$  days of observations per subject, for a total of  $M = 10,372$  days of measurements,  $T = 144$  measurements per day, and  $L = 20$  predictors per subject.

First, we compare the proposed FLFOSR approach with competing Bayesian methods (`refund:Gibbs` and `refund:VB`). We adopt the same hyperparameter and MCMC

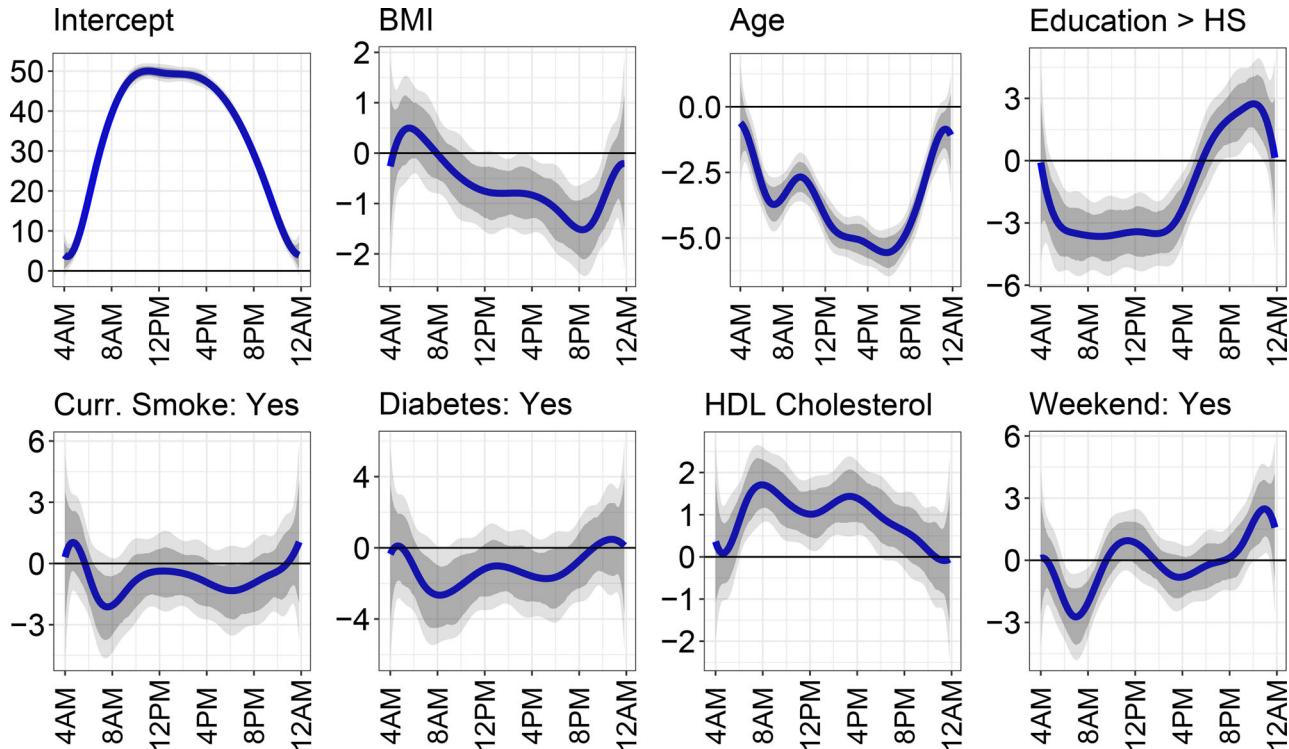
settings as in Section 4 and conduct a sensitive analysis in the supplementary material. However, we are unable to run `refund:Gibbs` or `refund:VB` due to memory limitations. Instead, we fit these competing methods on a random subsample of  $n = 278$  subjects, the largest subsample our memory could accommodate, solely for the purpose of computational comparisons. Table 3 shows the computational costs and MCMC efficiency for FLFOSR—fit to the whole dataset—compared to `refund:Gibbs` and `refund:VB`. FLFOSR required only about a minute to run, or about 5 min to generate 1000 effective samples from the posterior of the fixed effects functions. By comparison, `refund:Gibbs` required more than 21 hr to run, and due to MCMC inefficiencies, would require almost two weeks to generate 1000 effective samples—even on this much smaller dataset. Similarly, the VB approximation is dramatically slower than the FLFOSR MCMC sampling algorithm, again on the smaller dataset. Thus, the Goldsmith and Kitago (2016) strategy of (i) using VB to obtain initial estimates and (ii) using Gibbs sampling for the final results is neither necessary nor feasible: we instead should proceed exclusively with FLFOSR, which provides fully Bayesian inference on the complete dataset in a fraction of the time.

We summarize the FMM inference for select fixed effects functions  $\tilde{\alpha}_\ell(\cdot)$  in Figure 5 (see supplementary material for the remaining fixed effects functions). The results show interesting, but not necessarily unsurprising relationships between certain health factors and time-of-day PA. Age is a strong negative predictor of activity, especially in the afternoon hours. Current

**Table 2.** NHANES variables used in the regression analysis of PA data.

Variable	Values
<i>Response variable:</i>	
Activity level	[0, 440.0]
<i>Sociodemographic variables:</i>	
Gender	<b>Male</b> (53%), Female (47%)
Age (years)	[35, 84.8]
Race	<b>White</b> (54%), Black (21%), Hispanic (22%), Other (4%)
Education level	< HS (23%), = HS (53%), > HS (24%)
<i>Alcohol and nicotine use variables:</i>	
Drinks per week	[0, 105]
Smoke cigarettes	Current (20%), Former, (30%), <b>Never</b> (50%)
<i>Health-related variables:</i>	
Body Mass Index (kg/m <sup>2</sup> )	[15.9, 57.4]
HDL cholesterol (mg/dL)	[23, 188]
Total cholesterol (mg/dL)	[92, 458]
Systolic blood pressure (mmHg)	[80, 270]
Has congestive heart disease	Yes (4%), <b>No</b> (96%)
Has congestive heart failure	Yes (2%), <b>No</b> (98%)
Has cancer	Yes (9%), <b>No</b> (91%)
Has diabetes	Yes (10%), <b>No</b> (90%)
Has stroke	Yes (2%), <b>No</b> (98%)
<i>Other variables:</i>	
Is weekend	Yes (28%), <b>No</b> (72%)

NOTE: The range of observed values is reported for continuous variables and the percentage of each category is reported for categorical variables. Baseline categories used for categorical variables are bolded. In Education Level, "HS" refers to high school.



**Figure 5.** Posterior inference for the fixed effects functions  $\tilde{\alpha}_\ell(\cdot)$  for select covariates. The solid lines denote the posterior mean with 95% pointwise (dark band) and joint (light band) credible intervals.

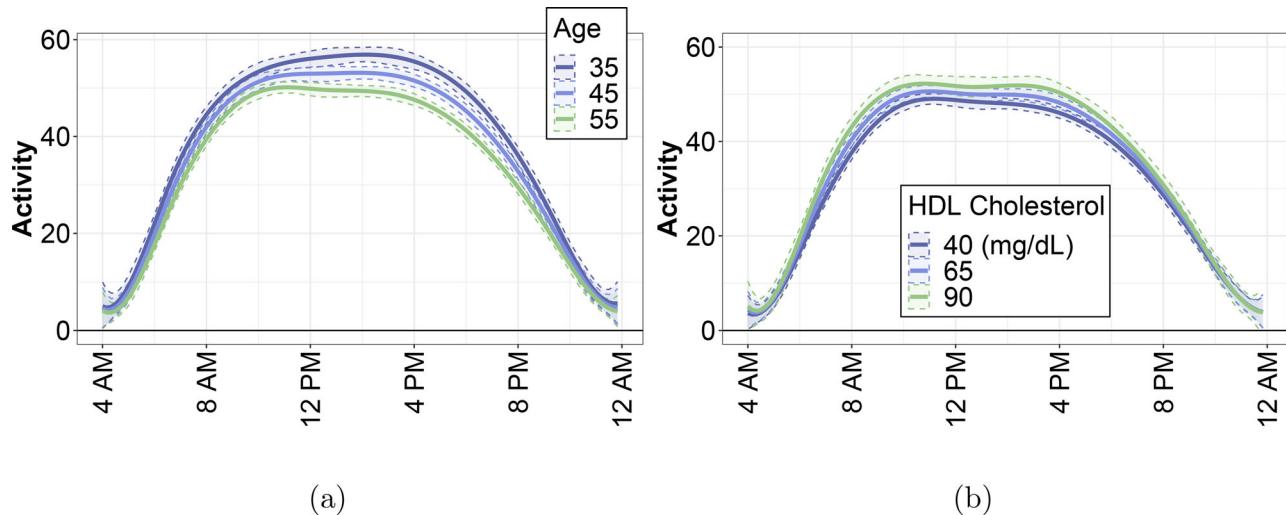
cigarette smokers are less active, predominantly in the morning, but this effect is not as strong later in the day. Individuals with higher educational attainment are less active throughout the daytime, but more active after in the early evening. Among health variables, certain comorbidities like diabetes are associated with less activity. HDL cholesterol, commonly referred to as "good cholesterol", is associated with higher activity throughout the day.

As an alternative visualization of these effects, Figure 6 shows the predicted population-level mean activity for varying levels of (a) age and (b) HDL cholesterol. The other predictors were set to the baseline and the mean values for categorical and continuous variables, respectively. As age increases from 35 to 55, there is a substantial overall decrease in predicted PA, which is particularly pronounced in the afternoon. HDL cholesterol levels in the data ranged from 23 mg/dL to 188 mg/dL, but levels below 40 mg/dL are broadly classified as "at risk" of heart disease and are generally recommended to be above 60 mg/dL (Expert Panel 2001). The right hand plot shows the predicted activity curves when raising HDL levels from the lower cutoff point up

**Table 3.** Computational comparisons among Bayesian algorithms for the NHANES dataset.

Method	Dataset	Time (N = 1000)	$\bar{s}_{1000}$	$\bar{N}_{\text{eff}}/N$
FLFOSR	Full (n = 1723)	1.3 min	4.8 min	0.27
refund: Gibbs	Reduced (n = 278)	21.3 hr	333.5 hr	0.06
refund: VB	Reduced (n = 278)	7.5 min	–	–

NOTE: Performance is evaluated using time to  $N = 1000$  MCMC samples, time to 1000 effective samples ( $\bar{s}_{1000}$ ), and average relative efficiency ( $\bar{N}_{\text{eff}}/N$ ). The proposed FLFOSR requires only a fraction of the computing time for a much larger dataset.



**Figure 6.** Mean activity level for varying levels of (a) age and (b) HDL cholesterol with 95% credible intervals for a subject with baseline categorical covariate levels and average values for continuous covariates.

into the recommended amount, where slightly heightened PA throughout waking hours is predicted as HDL levels increase.

## 6. Discussion

We introduced a new MCMC algorithm for Bayesian regression with longitudinal functional data. The algorithm applies for a broad and widely useful class of functional mixed models that include (i) basis expansions and (ii) fixed effects and subject- and replicate-specific random effects. Most notably, our MCMC sampler delivers unrivaled MCMC efficiency and unmatched computational scalability, and is empirically faster than state-of-the-art frequentist competitors—while also providing posterior uncertainty quantification for all model parameters and predictions. The scalability of our algorithm persists across all dimensions—the number subjects, replicates per subjects, covariates, and observation points—and thus establishes a new state-of-the-art for (Bayesian or frequentist) computing with longitudinal functional data. The proposed model showcases excellent point estimation accuracy and uncertainty quantification across a broad array of simulated data scenarios.

In practice, our MCMC algorithm enables complex Bayesian model-fitting with large functional datasets without the need for intensive computing resources. Indeed, we were able to perform fully Bayesian inference on a large physical activity dataset in about a minute using a personal desktop without any parallel processing. As the usage of wearable devices continues to grow, so does the need for highly scalable methods to study these high-resolution and high-dimensional data. For example, the UK Biobank study contains PA measurements of over 100,000 subjects across 7 days, measured in 5 sec intervals (Doherty et al. 2017), implying over 12 billion rows of PA measurements.

One computational aspect that we did not explore in this article is the memory usage of each algorithm. For the NHANES data, memory issues precluded use of the competing methods in the `refund` package on the whole dataset, while the proposed FLFOSR approach encountered no such problems. Thus, efficient memory usage may be another advantage of FLFOSR, but further analysis is needed.

## Supplementary Materials

**Additional results:** A document containing a sensitivity analysis and additional results for the application. (supplement.pdf, PDF)

**R code:** Code for model implementation and producing all simulation and application results in the paper. (code.zip, Zipped file)

## Disclosure Statement

The authors report there are no competing interests to declare.

## Funding

This material is based upon work supported by the National Science Foundation (SES-2214726).

## ORCID

Thomas Y. Sun  <http://orcid.org/0000-0002-5385-5083>

Daniel R. Kowal  <http://orcid.org/0000-0003-0917-3007>

## References

- Besse, P. C., Cardot, H., and Stephenson, D. B. (2000), “Autoregressive Forecasting of Some Functional Climatic Variations,” *Scandinavian Journal of Statistics*, 27, 673–687. [1]
- Bhattacharya, A., Chakraborty, A., and Mallick, B. K. (2016), “Fast Sampling with Gaussian Scale Mixture Priors in High-Dimensional Regression,” *Biometrika*, 103, 985–991. [6]
- Boschi, T., J. Di Iorio, L. Testa, M. A. Cremona, and F. Chiaromonte (2021). Functional Data Analysis Characterizes the Shapes of the First COVID-19 Epidemic Wave in Italy,” *Scientific Reports*, 11, 17054. [1]
- Cederbaum, J., Pouplier, M., Hoole, P., and Greven, S. (2016), “Functional Linear Mixed Models for Irregularly or Sparsely Sampled Data,” *Statistical Modelling*, 16, 67–88. [2]
- Cui, E., Leroux, A., Smirnova, E., and Crainiceanu, C. M. (2022), “Fast Univariate Inference for Longitudinal Functional Models,” *Journal of Computational and Graphical Statistics*, 31, 219–230. [2,3,7,8,9]
- Doherty, A., Jackson, D., Hammerla, N., Plötz, T., Olivier, P., Granat, M. H., White, T., van Hees, V. T., Trenell, M. I., Owen, C. G., Preece, S. J., Gillions, R., Sheard, S., Peakman, T., Brage, S., and Wareham, N. J. (2017), “Large Scale Population Assessment of Physical Activity Using Wrist



- Worn Accelerometers: The UK Biobank Study," *PloS One*, 12, e0169649. [1,12]
- Expert Panel. (2001), "Executive Summary of the Third Report of the National Cholesterol Education Program (NCEP) Expert Panel on Detection, Evaluation, and Treatment of High Blood Cholesterol in Adults (Adult Treatment Panel III)," *Journal of the American Medical Association*, 285, 2486–2497. [11]
- Gao, Y., and Kowal, D. R. (2024), "Bayesian Adaptive and Interpretable Functional Regression for Exposure Profiles," *The Annals of Applied Statistics*, 18, 642–663. [4,9]
- Gelman, A. (2006), "Prior Distributions for Variance Parameters in Hierarchical Models," *Bayesian Analysis*, 1, 515–534. [4]
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013), *Bayesian Data Analysis* (3rd ed.), Boca Raton, FL: CRC Press. [7]
- Goldsmith, J., and Kitago, T. (2016), "Assessing Systematic Effects of Stroke on Motor Control by Using Hierarchical Function-on-Scalar Regression," *Journal of the Royal Statistical Society, Series C*, 65, 215–236. [2,3,4,7,10]
- Goldsmith, J., Scheipl, F., Huang, L., Wrobel, J., Di, C., Gellar, J., Harezlak, J., McLean, M. W., Swihart, B., Xiao, L., Crainiceanu, C., and Reiss, P. T. (2021), *refund: Regression with Functional Data*, R package version 0.1-24. [7]
- Goldsmith, J., Zipunnikov, V., and Schrack, J. (2015), "Generalized Multi-level Function-on-Scalar Regression and Principal Component Analysis," *Biometrics*, 71, 344–353. [2,3]
- Guo, W. (2002), "Functional Mixed Effects Models," *Biometrics*, 58, 121–128. [2,3]
- Hilden, P., Schwartz, J. E., Pascual, C., Diaz, K. M., and Goldsmith, J. (2023), "How Many Days Are Needed? Measurement Reliability of Wearable Device Data to Assess Physical Activity," *PloS One*, 18, e0282162. [9]
- Huo, S., Morris, J. S., and Zhu, H. (2022), "Ultra-Fast Approximate Inference Using Variational Functional Mixed Models," *Journal of Computational and Graphical Statistics*, 1–32. [3]
- Kowal, D. R. (2022a), "Fast, Optimal, and Targeted Predictions Using Parameterized Decision Analysis," *Journal of the American Statistical Association*, 117, 1875–1886. [9]
- (2022b), "Subset Selection for Linear Mixed Models," *Biometrics*, 79, 1853–1867. [9]
- Kowal, D. R., and D. C. Bourgeois (2020), "Bayesian Function-on-Scalars Regression for High-Dimensional Data," *Journal of Computational and Graphical Statistics*, 29, 629–638. [2,3,6,9]
- Kowal, D. R., Matteson, D. S., and Ruppert, D. (2019), "Functional Autoregression for Sparsely Sampled Data," *Journal of Business & Economic Statistics*, 37, 97–109. [4]
- 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. [1,2]
- Leroux, A., Di, J., Smirnova, E., McGuffey, E. J., Cao, Q., Bayatmokhtari, E., Tabacu, L., Zipunnikov, V., Urbanek, J. K., and Crainiceanu, C. (2019), "Organizing and Analyzing the Activity Data in NHANES," *Statistics in Biosciences*, 11, 262–287. [1,9,10]
- Li, R., Xiao, L., Smirnova, E., Cui, E., Leroux, A., and Crainiceanu, C. M. (2022), "Fixed-Effects Inference and Tests of Correlation for Longitudinal Functional Data," *Statistics in Medicine*, 41, 3349–3364. [2,3,7,8,9]
- Liebl, D. (2013), "Modeling and Forecasting Electricity Spot Prices: A Functional Data Perspective," *The Annals of Applied Statistics*, 7, 1562–1592. [1]
- Morris, J. S. (2017), "Comparison and Contrast of Two General Functional Regression Modelling Frameworks," *Statistical Modelling*, 17, 59–85. [2]
- Morris, J. S., Arroyo, C., Coull, B. A., Ryan, L. M., Herrick, R., and Gortmaker, S. L. (2006), "Using Wavelet-Based Functional Mixed Models to Characterize Population Heterogeneity in Accelerometer Profiles: A Case Study," *Journal of the American Statistical Association*, 101, 1352–1364. [2,9]
- Morris, J. S., and Carroll, R. J. (2006), "Wavelet-based Functional Mixed Models," *Journal of the Royal Statistical Society, Series B*, 68, 179–199. [2,3,4]
- Nishimura, A., and Suchard, M. A. (2022), "Prior-Preconditioned Conjugate Gradient Method for Accelerated Gibbs Sampling in "Large  $n$ , Large  $p$ " Bayesian Sparse Regression," *Journal of the American Statistical Association*, 118, 2468–2481. [6]
- Park, S. Y., and Staicu, A.-M. (2015), "Longitudinal Functional Data Analysis," *Stat*, 4, 212–226. [2,3]
- Reiss, P. T., Huang, L., and Mennes, M. (2010), "Fast Function-on-Scalar Regression with Penalized Basis Expansions," *The International Journal of Biostatistics*, 6, 28. [2]
- Rue, H. (2001), "Fast Sampling of Gaussian Markov Random Fields," *Journal of the Royal Statistical Society, Series B*, 63, 325–338. [6]
- Scheipl, F., Fahrmeir, L., and Kneib, T. (2012), "Spike-and-Slab Priors for Function Selection in Structured Additive Regression Models," *Journal of the American Statistical Association*, 107, 1518–1532. [3,4,6]
- Scheipl, F., Gertheiss, J., and Greven, S. (2016), "Generalized Functional Additive Mixed Models," *Electronic Journal of Statistics*, 10, 1455–1492. [2]
- Scheipl, F., Staicu, A.-M., and Greven, S. (2015), "Functional Additive Mixed Models," *Journal of Computational and Graphical Statistics*, 24, 477–501. [2]
- Schmid, D., Ricci, C., and Leitzmann, M. F. (2015), "Associations of Objectively Assessed Physical Activity and Sedentary Time with All-Cause Mortality in US Adults: The NHANES Study," *PloS One*, 10, e0119591. [9]
- Sera, F., Griffiths, L. J., Dezateux, C., Geraci, M., and Cortina-Borja, M. (2017), "Using Functional Data Analysis to Understand Daily Activity Levels and Patterns in Primary School-Aged Children: Cross-Sectional Analysis of a UK-Wide Study," *PloS One*, 12, e0187677. [9]
- Sergazinov, R., Leroux, A., Cui, E., Crainiceanu, C., Aurora, R. N., Punjabi, N. M., and Gaynanova, I. (2023), "A Case Study of Glucose Levels During Sleep Using Multilevel Fast Function on Scalar Regression Inference," *Biometrics*, 79, 3873–3882. [2,7]
- Smirnova, E., Leroux, A., Cao, Q., Tabacu, L., Zipunnikov, V., Crainiceanu, C., and Urbanek, J. K. (2020), "The Predictive Performance of Objective Measures of Physical Activity Derived From Accelerometry Data for 5-Year All-Cause Mortality in Older Adults: National Health and Nutritional Examination Survey 2003–2006," *The Journals of Gerontology, Series A*, 75, 1779–1785. [9]
- Yu, Y., and Meng, X.-L. (2011), "To Center or Not to Center: That Is Not the Question—An Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Efficiency," *Journal of Computational and Graphical Statistics*, 20, 531–570. [3,4,7]
- Zhu, H., Brown, P. J., and Morris, J. S. (2011), "Robust, Adaptive Functional Regression in Functional Mixed Model Framework," *Journal of the American Statistical Association*, 106, 1167–1179. [2]
- Zhu, H., Chen, K., Luo, X., Yuan, Y., and Wang, J.-L. (2019), "FMEM: Functional Mixed Effects Models for Longitudinal Functional Responses," *Statistica Sinica*, 29, 2007–2033. [2]
- Zipunnikov, V., Greven, S., Shou, H., Caffo, B. S., Reich, D. S., and Crainiceanu, C. M. (2014), "Longitudinal High-Dimensional Principal Components Analysis with Application to Diffusion Tensor Imaging of Multiple Sclerosis," *The Annals of Applied Statistics*, 8, 2175–2202. [2]