

Inference for the Extended Functional Cox Model: A UK Biobank Case Study

ERJIA CUI^{1†}, ANGELA ZHAO^{2†*}, ANDREW LEROUX³, MARTIN A. LINDQUIST²,

CIPRIAN M. CRAINICEANU²

1. Division of Biostatistics and Health Data Science, University of Minnesota, USA

2. Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, USA

3. Department of Biostatistics and Informatics, University of Colorado Asnchutz Medical

Campus, USA

azhao29@jh.edu

SUMMARY

Multiple studies have shown that scalar summaries of objectively measured physical activity (PA) using accelerometers are the strongest predictors of mortality, outperforming all traditional risk factors, including age, sex, body mass index (BMI), and smoking. Here we show that diurnal patterns of PA and their day-to-day variability provide additional information about mortality. To do that, we introduce a class of extended functional Cox models and corresponding inferential tools designed to quantify the association between multiple functional and scalar predictors with time-to-event outcomes in large-scale (large n) high-dimensional (large p) datasets. Methods are applied to the UK Biobank study, which collected PA at every minute of the day for up to seven days, as well as time to mortality (93,370 participants with good quality accelerometry data and 931 events). Simulation studies show that methods perform well in realistic scenarios and scale up

[†]These two authors contributed equally to this work.

*To whom correspondence should be addressed.

to studies an order of magnitude larger than the UK Biobank accelerometry study. Establishing the feasibility and scalability of these methods for such complex and large data sets is a major milestone in applied Functional Data Analysis (FDA).

Key words: Functional data analysis; Survival analysis; Wearable devices; Biobank.

1. INTRODUCTION

We introduce a class of extended functional Cox models and corresponding inferential tools designed to quantify the association between multiple functional and scalar predictors and time-to-event outcomes in large-scale (large n) high-dimensional (large p) datasets. We show that these models can be viewed as a particular case of mixed effects models, which provides a principled and scalable inferential framework. Methods are applied to the UK Biobank study, which collected physical activity (PA) at every minute of the day for up to seven days, time to mortality, as well as many other behavioral and demographic covariates (93,370 participants with good quality accelerometry data and 931 events). Establishing the feasibility and scalability of these methods and providing ready to use, reproducible software that works on biobank-sized data sets is a major milestone in applied Functional Data Analysis (FDA).

The emergence of biobank data sets has raised many computational and methodological challenges that range from handling very large data sets, accounting for non-rectangular data structures (e.g., a “covariate” could be a scalar, such as age or BMI, or a high dimensional structure, such as minute-level PA data or images), to estimation and inference. Examples of such biobank studies are the UK Biobank (Allen *and others*, 2012), the United States All of Us (The All of Us Research Program Investigators, 2019), and the China Kadoorie Biobank Chen *and others* (2011). Here we focus on the UK Biobank, a large study conducted in the United Kingdom containing demographic, lifestyle, medical, and genetic information for around 500,000 participants.




Fig. 1. Top five panels: minute-level log-transformed PA averaged across days within individuals. Left column: (A1) the mean log-transformed PA over time for individuals alive (blue) and deceased (red) as of July 31, 2020. Middle column: lasagna plots for minute-level log-transformed PA for 9310 alive (B1) and 931 deceased (C1) participants. Right column: smoothed (orange) and unsmoothed (gray) data of one alive (D1) and one deceased (E1) participant overlayed with the population mean (cyan). Bottom panels display the same information as the top five panels but for the minute-level standard deviation across days within individuals.

In addition to this information, UK Biobank collected high-quality accelerometry data for around 100,000 participants, or about 20% of the original sample.

To illustrate the PA data structure, Figure 1 displays the minute-level PA data expressed as Euclidean Norm Minus One (ENMO, Bakrania *and others* (2016)) separated by individuals who were still alive or deceased as of July 31, 2020 in the UK Biobank. For reference, PA accelerometry data were collected between 2013 and 2015. For the purpose of this paper we used PA data summarized at the minute level for the duration of the accelerometry substudy, which could be up to seven days for every study participant. The top panels display the minute-level means of the log-transformed PA, while the bottom panels display the minute-level standard deviations of the log-transformed PA. Both of these measures are calculated at a given minute of the day across days within each individual. The panels in the middle columns in Figure 1 display the PA data for one individual per row where participants are ordered by age, and data are separated by individuals who were still alive (panels B1 and B2) and deceased (panels C1 and C2) by July 31, 2020. The x-axis represents time in minutes from midnight to midnight and the y-axis is the $\log(1 + \text{ENMO})$ value, where ENMO is calculated at the minute level. The left panels (A1 and A2) are the average across study participants within each group, where blue corresponds to study participants who were alive. More precisely, the blue line in panel A1 is obtained by taking the column averages in panel B1, while the red line in panel A1 is obtained by taking the column averages in panel C1. Together, these panels suggest that study participants who were alive had higher baseline mean activity and standard deviation between 7 am and 11 pm. The right columns provide examples of individual trajectories (gray), smooth estimators (orange), and corresponding population means (blue) for one participant who was alive (D1 and D2) and one who died (E1 and E2) before July 31, 2020. The panels D1 and E1 correspond to the mean, while D2 and E2 correspond to standard deviations across days at the same minute within the same person.

Using data from multiple studies (Agarwala *and others*, 2024; Ledbetter *and others*, 2022; Smirnova *and others*, 2020; Tabacu *and others*, 2020), including UK Biobank (Leroux *and others*, 2021), we have shown that summaries (e.g., mean, time in range) of the objectively measured minute level PA data displayed in Figure 1 are the strongest predictors of mortality, with better prediction performance than all traditional risk factors, including age, sex, body mass index (BMI), and smoking. The main goal of this paper is to investigate whether the complex, high-dimensional (1,440 observations per day), multivariate (minute-level mean and standard deviation at every minute over available days) PA data provides additional information about mortality, while accounting for other covariates. To do that we treat minute-level PA data as functional data (Crainiceanu *and others*, 2024; Cui *and others*, 2023; Kokoszka and Reimherr, 2017; Ramsay and Silverman, 2005). Some of earliest methods for studying the association between functional data and time-to-event outcomes have been introduced by Gellar *and others* (2015) followed by several different methods (Cui *and others*, 2021; Hao *and others*, 2021; Qu *and others*, 2016; Kong *and others*, 2018); see Chapter 7 in Crainiceanu *and others* (2024) for a recent review of methods. The current paper expands on the existing literature in several directions by: (1) allowing for the inclusion of multiple functional, nonparametric, and parametric terms in the log hazard; (2) expanding inference by constructing pointwise and Correlation and Multiplicity Adjusted (CMA, Crainiceanu *and others* (2024)) confidence intervals and p-values; (3) scaling up methods to biobank-size data sets with over 90,000 study participants and 1,440 observations per study participant; (4) extending the models to incorporate distributions as predictors (Ghosal *and others*, 2023a,b, 2022; Petersen *and others*, 2022) in a time-to-event model; and (5) simulating for the first time survival models with high dimensional predictors of the size and complexity of the UK Biobank. Finally, we provide reproducible software and discuss results of the UK Biobank data analysis. Such models, methods, and applications are crucial for increasing the relevance of functional data analysis methods in real world applications.

The paper is organized as follows. We first introduce key components of statistical models and inferential tools in Section 2. We then introduce the UK Biobank study, specific data processing steps, and the application results in Section 3. Section 4 presents a simulation study. We conclude with discussions in Section 5.

2. METHODS

2.1 The Extended Functional Cox Models

For study participant $i = 1, \dots, I$ we observe $Y_i = \min(T_i, C_i)$, where T_i is the event time, C_i is the censoring time, and let $\Delta_i = \mathbb{1}(T_i \leq C_i)$ be the binary event indicator. We assume that C_i is independent of T_i conditional on predictors, which is a realistic assumption in the UK Biobank where censoring is primarily administrative. Denote by $\mathbf{Z}_i = [Z_{i1}, \dots, Z_{ip}]^T \in \mathbb{R}^p$ the p scalar predictors and by $\mathbf{X}_{i1}, \mathbf{X}_{i2}, \dots, \mathbf{X}_{iR}$ the R functional predictors, where $\mathbf{X}_{ir} = \{X_{ir}(s_r), s_r \in \mathcal{S}_r\}$ and $\mathcal{S}_r$ could potentially be different domains of functional data. For the UK Biobank, $\mathbf{X}_{i1}$ and $\mathbf{X}_{i2}$ may represent the mean and standard deviation of the PA measure at the same minute of the day over multiple days, respectively. In this case $\mathcal{S}_1 = \mathcal{S}_2$ is the set of time points expressed in minutes from midnight to midnight. In addition, denote by $\mathbf{V}_i = [V_{i1}, \dots, V_{iq}]^T \in \mathbb{R}^q$ the q scalar predictors assumed to have additive (nonlinear or linear) effects on the log hazard. An extended functional Cox model is

$$\log \lambda_i(t) = \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + \sum_{l=1}^q f_l(V_{il}) + \sum_{r=1}^R \int_{\mathcal{S}_r} X_{ir}(s_r) \beta_r(s_r) ds_r, \quad (2.1)$$

where $\lambda_i(t) = \lambda_i(t | \mathbf{Z}_i, \mathbf{V}_i, \mathbf{X}_{i1}, \dots, \mathbf{X}_{iR})$ is the hazard function for the i th subject and $\lambda_0(\cdot)$ is the baseline hazard function. When $R = 1$ and $q = 0$, model (2.1) reduces to the linear functional Cox model proposed by Gellar *and others* (2015). The functional coefficient $\beta_r(\cdot)$ quantifies the effect of the r th functional predictor at different locations of the functional domain. For example, if $X_{ir}(s_r)$ is the mean value of the PA intensity at time s_r where $s_r = 1, \dots, 1440$ is the time

of day in minutes, $\beta_r(s_r)$ is interpreted as the effect of a one unit increase in PA intensity at time s_r on the log hazard while keeping the rest of PA constant. In addition, $f_l(\cdot)$ represents an unknown, smooth function that captures the effect of the l th scalar predictor on the log-hazard of mortality. We anticipate that such models will become increasingly common not because they are complex, but because they provide practical answers to a range of scientific questions. They also naturally extend models for functional survival, semiparametric, and Generalized Additive Models (GAM) regression. Moreover, these methods have to adapt to the reality of large data sets, such as UK Biobank, which contains thousands of scalar and high-dimensional variables for hundreds of thousands of study participants.

Model (2.1) is fit via basis expansion both for the additive and functional terms; for a general description of the approach see Chapters 3-7 in Crainiceanu *and others* (2024). For $l = 1, \dots, q$, denote by $\psi_l(v) = [\psi_{l1}(v), \dots, \psi_{K_{la}}(v)]^T$ the collection of K_{la} spline functions defined over the range of the l th covariate, V_l . For $r = 1, \dots, R$, denote by $\phi_r(s_r) = [\phi_{r1}(s_r), \dots, \phi_{K_{rb}}(s_r)]^T$ a collection of K_{rb} spline functions defined over $\mathcal{S}_r$. For $l = 1, \dots, q$, $f_l(V_{il}) = \sum_{k=1}^{K_{la}} u_{lk} \psi_{lk}(V_{il})$ and for $r = 1, \dots, R$, $\beta_r(s_r) = \sum_{k=1}^{K_{rb}} b_{rk} \phi_{rk}(s_r)$, where $\mathbf{u}_l = [u_{l1}, \dots, u_{lK_{la}}]^T$ and $\mathbf{b}_r = [b_{r1}, \dots, b_{rK_{rb}}]^T$ are unknown spline coefficients. With these bases expansions, Model (2.1) becomes

$$\begin{aligned} \log \lambda_i(t) &= \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + \sum_{l=1}^q f_l(V_{il}) + \sum_{r=1}^R \int_{\mathcal{S}_r} X_{ir}(s_r) \sum_{k=1}^{K_{rb}} b_{rk} \phi_{rk}(s_r) ds_r \\ &= \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + \sum_{l=1}^q \sum_{k=1}^{K_{la}} u_{lk} \psi_{lk}(V_{il}) + \sum_{r=1}^R \sum_{k=1}^{K_{rb}} b_{rk} \int_{\mathcal{S}_r} X_{ir}(s_r) \phi_{rk}(s_r) ds_r \quad (2.2) \\ &= \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + \sum_{l=1}^q \psi_l(V_{il})^T \mathbf{u}_l + \sum_{r=1}^R \mathbf{W}_{ir}^T \mathbf{b}_r, \end{aligned}$$

where $\mathbf{W}_{ir} = [W_{ir1}, \dots, W_{irK_{rb}}]^T$ and $W_{irk} = \int_{\mathcal{S}_r} X_{ir}(s_r) \phi_{rk}(s_r) ds_r$. This general structure is exactly the structure of the Cox regression model, though the covariates are designed specifically to handle functional data and nonparametric functions. Spline models are overparameterized and smoothing is induced by assumptions of the type $\mathbf{u}_l \sim N(\mathbf{0}, \sigma_{ul}^2 \boldsymbol{\Sigma}_{ul})$ and $\mathbf{b}_r \sim N(\mathbf{0}, \sigma_{br}^2 \boldsymbol{\Sigma}_{br})$, where $\boldsymbol{\Sigma}_{ul}$ and $\boldsymbol{\Sigma}_{br}$ are known and pre-specified for each specific spline basis, and σ_{ul}^2 and σ_{br}^2 are

unknown variance parameters acting as smoothing parameters on the spline coefficients. The idea can be traced back to the papers of Crainiceanu and Goldsmith (2010); Goldsmith *and others* (2011), which first described it for scalar-on-function regression (SoFR); for a historical perspective and complete details, see Crainiceanu *and others* (2024). These ideas were extended to time-to-event outcomes by maximizing the associated penalized partial log-likelihood (Crainiceanu *and others*, 2024; Cui, 2023; Leroux, 2020; Wood, 2017).

While equation (2.2) looks complicated, closer examination reveals that it can be viewed as a mixed effects model, where $\mathbf{u}_i$ and $\mathbf{b}_r$ are random effects parameters with a specific structure and $\boldsymbol{\gamma}$ are fixed effects parameters. Therefore, fitting model (2.2) is equivalent to fitting a survival model with random effects, which can be done for some specific cases, but not in general. Indeed, in a recent monograph Crainiceanu *and others* (2024) it is shown how to fit related models using the `refund` Goldsmith *and others* (2024) and `mgcv` Wood (2017) R packages, though they were applied to a smaller data set (NHANES) (Leroux *and others*, 2019). We take this approach and show how to adapt it to the more complex models considered here.

We now provide the R syntax to fit an extended functional Cox model when $p = 1$, $q = 1$, and $R = 2$. The formula can be extended to allow for more linear, additive, and functional terms. Suppose that the data are contained in the data frame `data` with the following fields: `time` is the observed time Y_i , `status` is the event indicator Δ_i , `s1Var` is a scalar predictor assumed to have linear effect on the log hazard, `saVar` is a scalar predictor assumed to have additive (nonlinear or linear) effect on the log hazard. In addition, `sFun1` is the matrix containing the first functional predictor where each row represents a subject and each column represents a location of the functional domain, and `sFun2` is the matrix of the second functional predictor stored in a similar format. For each functional predictor, we first create two matrices, `atmat` and `almat`, which are used to integrate functional observations over time.

```
data$almat1 <- I(matrix(1 / ncol(data$sFun1),
```

```

ncol = ncol(data$sFun1), nrow = nrow(data)))

data$atmat1 <- I(matrix(1:ncol(data$sFun1),

ncol = ncol(data$sFun1, nrow = nrow(data), byrow = TRUE))

data$almat2 <- I(matrix(1 / ncol(data$sFun2),

ncol = ncol(data$sFun2), nrow = nrow(data)))

data$atmat2 <- I(matrix(1:ncol(data$sFun2),

ncol = ncol(data$sFun2, nrow = nrow(data), byrow = TRUE))

```

After creating these matrices and storing them as four variables in `data`, the extended functional Cox model (2.1) can be fit using the `mgcv::gam` function as follows.

```

fit_efcm <- gam(time ~ slVar + s(saVar) +

s(atmat1, by = almat1 * sFun1) +

s(atmat2, by = almat2 * sFun2),

data = data, weights = status, family = cox.ph())

```

This approach to fitting complex functional Cox models is modular, as the additive term can be specified by applying the `s()` function to the variable, and including multiple functional predictors can be achieved by simply adding multiple `s(atmat, by = almat * sFun)` terms to the formula. Including additional nonparametric terms can similarly be done by adding terms of the type `s(sVar)`. Providing the model, describing the inferential approach, and introducing reproducible and extendable software implementation are important steps for promoting the use of functional data analysis methods in real data sets. The fact that these methods can be applied to data sets of the size and complexity of the UK Biobank data provides proof that the proposed methods are truly scalable and useful.

2.2 Correlation and Multiplicity Adjusted Inference

Consider the general case when $\beta(s) = \phi(s)^T \mathbf{b}$ is one of the functional parameters and we obtain an estimator $\hat{\mathbf{b}}$ of $\mathbf{b}$ and an estimator of its variance $\hat{\mathbf{V}}_{\mathbf{b}}$ (Gray, 1992; Verweij and Van Houwelingen, 1994). The pointwise 95% confidence interval for $\beta(s)$ can be constructed as $\hat{\beta}(s) \pm 1.96\sqrt{\phi(s)^T \hat{\mathbf{V}}_{\mathbf{b}} \phi(s)}$. Unfortunately, these intervals do not account for the multiple estimators or the complex correlations along the functional domain. To address this problem we focus on correlation and multiplicity adjusted (CMA) confidence intervals, a nomenclature introduced in Crainiceanu *and others* (2024).

The idea of constructing confidence intervals taking into account the correlation along the domain was first introduced as “simultaneous confidence bands” in the semiparametric regression setting by Ruppert *and others* (2003). It was later extended to the functional data setting by Crainiceanu *and others* (2012). Since then, the method was widely used in the literature (Cui *and others*, 2022a; Park *and others*, 2018; Sergazinov *and others*, 2023). Here we use the label “Correlation and Multiplicity Adjusted (CMA)” to explicitly refer to the inherent correlation of functional data and the many points where inference is conducted Crainiceanu *and others* (2024). Let $\boldsymbol{\beta} = \{\beta(s), s \in \mathcal{S}\}$. Under the assumption that the bias of $\hat{\boldsymbol{\beta}}$ is negligible and $\hat{\mathbf{b}}$ is multivariate normal, we have $\hat{\boldsymbol{\beta}} \sim \mathcal{N}(\boldsymbol{\beta}, \hat{\mathbf{V}}_{\boldsymbol{\beta}})$, where $\hat{\mathbf{V}}_{\boldsymbol{\beta}}(s, t) = \phi(s)^T \hat{\mathbf{V}}_{\mathbf{b}} \phi(t)$. Denote by $\hat{\mathbf{D}}_{\boldsymbol{\beta}} = \sqrt{\text{diag}(\hat{\mathbf{V}}_{\boldsymbol{\beta}})}$ the square root of the diagonal of $\hat{\mathbf{V}}_{\boldsymbol{\beta}}$. The goal of CMA inference is to find a critical value $q_{1-\alpha}$ such that $P(|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}| \leq q_{1-\alpha} \hat{\mathbf{D}}_{\boldsymbol{\beta}}) = 1 - \alpha$, where all inequalities are entrywise and simultaneous. In a recent monograph (Crainiceanu *and others*, 2024), several methods were proposed to obtain $q_{1-\alpha}$ in functional regression contexts. Here we use one of these approaches based on the assumption of multivariate normality of $\hat{\mathbf{b}}$, and hence $\hat{\boldsymbol{\beta}}$. Notice that $q_{1-\alpha}$ is essentially the equi-coordinate quantile of a multivariate normal distribution since $\hat{\boldsymbol{\beta}}$ can only be observed on a finite grid. Therefore, given $\hat{\mathbf{V}}_{\boldsymbol{\beta}}$, one can use existing software based on the multivariate normal distribution, such as the `mvtnorm::qmvnorm` function (Bornkamp, 2018), to calculate these quantiles. Suppose

that $\widehat{\mathbf{V}}_\beta$ is stored as a matrix `Vbeta`, the R syntax to obtain quantiles at level `alpha` is simply

```
q <- qmvnorm(1 - alpha, corr = cov2cor(Vbeta), tail = "both.tails")$quantile
```

This approach is very fast, accounts for the correlations of the estimators, and can be implemented in one line of code. It has the advantage that it avoids resampling, which is much slower. In some cases, especially when the dimension of β is very large or the normality assumption of $\widehat{b}$ is suspect, it is a good idea to have a plan B. Suppose that $\widehat{\beta}$ is evaluated on the grid $\{s_1, \dots, s_L\}$, the critical value $q_{1-\alpha}$ can be estimated as follows:

1. Simulate $\beta^{(b)}(s_1), \dots, \beta^{(b)}(s_L)$ from $\mathcal{N}(\widehat{\beta}, \widehat{\mathbf{V}}_\beta)$;
2. Calculate $r^{(b)} = \max_{l=1, \dots, L} \{|\beta^{(b)}(s_l) - \widehat{\beta}(s_l)| / \widehat{\mathbf{D}}_\beta(s_l)\}$;
3. Repeat Step 1-2 for $b = 1, \dots, B$ to obtain $r^{(1)}, \dots, r^{(B)}$.
4. $q_{1-\alpha}$ is estimated as the $(1 - \alpha)$ quantile of $r^{(1)}, \dots, r^{(B)}$.

After obtaining $q_{1-\alpha}$ using either approach, the CMA confidence interval for $\beta(s)$ at level α is constructed as $\widehat{\beta}(s) \pm q_{1-\alpha} \sqrt{\phi^T(s) \widehat{\mathbf{V}}_\beta \phi(s)}$. Since the critical value is estimated as the quantile of the maximum of the quantiles along the functional domain, it is guaranteed to be larger than or equal to the quantile of a standard normal distribution. As a result, the CMA confidence intervals are uniformly non-narrower than the pointwise confidence intervals at a given significance level.

These CMA confidence intervals can be calculated for any value of α . For every $s_l \in S$, we can then find the largest value of α for which the confidence interval does not include zero. We denote this probability by $p_{\text{pCMA}}(s)$ and refer to it as the pointwise correlation and multiplicity adjusted (pointwise CMA) p-value. As $q_{1-\alpha} \geq z_{1-\alpha/2}$, the pointwise CMA confidence intervals are non-narrower than the pointwise unadjusted confidence intervals. Therefore, the pointwise CMA p-values will be larger than or equal to the pointwise unadjusted p-values, leading to fewer or equal amounts of results being deemed “statistically significantly different from zero.” However, the tests will preserve the family-wise error rate (FWER) while accounting for test correlations.

Similarly, we define the global pointwise correlation and multiplicity adjusted (global CMA) p-value as the largest α level at which at least one confidence interval $\hat{\beta}(s_l) \pm q_{1-\alpha} \hat{D}_\beta(s_l)$, for $l = 1, \dots, L$ does not contain zero. If we denote by $p_{\text{gCMA}}(s_1, \dots, s_L)$ this p-value, it can be shown that

$$p_{\text{gCMA}}(s_1, \dots, s_L) = \min\{p_{\text{pCMA}}(s_1), \dots, p_{\text{pCMA}}(s_L)\}$$

because the null hypothesis is rejected if it is rejected at any point in the domain of $\beta(s)$.

The advantage of using these p-values over the unadjusted p-values is that the tests preserve their nominal level. The pointwise tests focus on testing whether a particular value of the function is zero, whereas the global tests focus on testing whether the entire function is zero simultaneously.

2.3 Functional Distributional Cox Models

So far, we have focused on modeling the effects of functional predictors as a function of time, such as physical activity intensity at a particular time of the day, on time to event, such as mortality. An alternative approach is to consider the distribution of the time series of minute-level activities Ghosal *and others* (2023b); Petersen *and others* (2022). Here we extend the Cox regression models to quantify the association between the distribution of high-dimensional predictors and time to event. We show that techniques used for scalar-on-function regression can be extended to Cox models with distributional predictors.

Assume that for the i th study participant, $i = 1, \dots, N$, we transform the time series data $X_i(s)$ into the estimated quantile function, where $\hat{Q}_i(p)$ is the p -quantile of the vector $\{X_i(s_1), \dots, X_i(s_L)\}$. For the UK Biobank study, this quantile function provides information about the time spent at every level of physical activity, is easy to understand, and may be useful when comparing results from different studies. The functional distributional Cox model has the following form

$$\log \lambda_i(t | \mathbf{Z}_i, \mathbf{X}_i) = \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + \int_0^1 \hat{Q}_i(p) \beta_Q(p) dp, \quad (2.3)$$

where $\beta_Q(p)$ quantifies the effect of the quantile corresponding to probability p on survival outcomes. For functional models, the functional effects are incorporated by taking the integral over the functional domain. In the functional distributional Cox model, however, the effects are incorporated by taking the integral along the probability interval $[0, 1]$. From here on, we propose to use exactly the same approaches described for functional Cox models. Moreover, these models can be extended to the additive distributional Cox model Cui *and others* (2021)

$$\log \lambda_i(t|\mathbf{Z}_i, \mathbf{X}_i) = \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + \int_0^1 F\{p, \hat{Q}_i(p)\} dp, \quad (2.4)$$

where $F(\cdot, \cdot)$ is an unknown bivariate twice differentiable function. This formulation allows a nonlinear association between quantile functions and hazard by introducing an additive structure. This model is beyond the scope of our current paper.

3. UK BIOBANK APPLICATIONS

3.1 Overview of UK Biobank

UK Biobank is a large prospective cohort study that includes 502,520 participants registered with the National Health Service in England, Wales, and Scotland between 2006 and 2010 (Allen *and others*, 2012). Comprehensive data on demographic, lifestyle, and health-related factors were collected. Extensive follow-up of participants' health and mortality status was conducted through linkage to national healthcare, death registries, and self-reports (Allen *and others*, 2012). The UK Biobank received ethical approval (REC reference for UK Biobank 11/NW/0382) and participants have provided written informed consent.

Among the 502,520 participants, 236,519 were invited to participate in an accelerometry substudy between 2013 and 2015. On average, individuals who were invited were younger and had better self-reported overall health than the rest of UK Biobank study participants (Leroux *and others*, 2021). Of these participants, 103,700 (43.9% response rate) agreed to join the study. In

this substudy, each individual was asked to wear an Axivity AX3 accelerometer on their dominant wrist for seven consecutive days immediately after receiving it. Accelerometry data was recorded continuously along three orthogonal axes at 100 Hz with a dynamic range of $\pm 8g$. After the monitoring period, participants were asked to mail-in the device back to the coordinating center using a pre-paid envelope (Doherty *and others*, 2017).

3.2 Data Processing

Several quality control steps were performed after the raw accelerometry data was collected by the coordinating center. Among the 103,700 participants, 7,164 were identified as having poor data quality due to insufficient wearing time or incorrect device calibration. After excluding these participants, there were 96,536 study participants with “good” quality accelerometry data, as determined by the UK Biobank team. The processed accelerometry data was released at multiple resolutions. In our analysis, we used the five-second average Euclidean Norm Minus One (ENMO, Van Hees *and others* (2013), expressed in milli- g ’s ($g = 9.81m/s^2$). The five-second data was then aggregated using 1-minute bins resulting in 1,440 observations per day per study participant. To control data quality, we further excluded study participants with fewer than 3 calendar days of accelerometry data and less than 1,368 minutes of estimated wear time (95% of the day.) The final data set consists of the minute-level accelerometry data from 93,370 study participants, and is stored as a matrix with 1,440 columns, the so called “1440 format” (Leroux *and others*, 2021). To the best of our knowledge, this is the largest, best designed, and most quality-controlled study incorporating objectively measured PA to date.

In addition to the accelerometry data, multiple demographic, clinical, and health variables are available in the UK Biobank. The covariates included in our analysis are described in detail in Section 3.3. The primary outcome in this paper is all-cause mortality, which was obtained from the UK Biobank data field 40000, “Date of death”. Individuals who died before the accelerom-

etry measurement or have an invalid date of death were excluded from analysis. Mortality data available for this study includes observations up to June 13, 2020.

To reduce skeweness of the data, minute-level ENMO were transformed using the function $f(\text{ENMO}) = \log(1 + \text{ENMO})$ (Di and others, 2017). In some analyses the study-participant time series data were transformed into a distribution, where quantiles were obtained for each probability on a fine grid of probabilities between $[0, 1]$.

3.3 Results

We start by investigating the association between physical activity at different times of the day and the risk of all-cause mortality after adjusting for age, sex, BMI, overall health rating, smoking, drinking, diabetes, cancer, stroke, coronary heart disease, and Townsend deprivation index. An extended functional Cox model (labeled M1) was fit using our estimation and inferential procedure described in Section 2. We define the time of the accelerometry measurement as the baseline time. Age at the time of the accelerometry was obtained by subtracting the year of the measurement from the year of birth. To better understand how age affects log hazard, we modeled the effect of age nonparametrically. Sex, BMI, smoking, drinking, overall health rating, and Townsend deprivation index were obtained from their respective data fields. Diabetes, stroke, cancer, and coronary heart disease were stored as binary indicators of whether or not the participant was diagnosed with the condition at baseline. Smoking and drinking indicators were assigned 1 if the participant ever smoked or ever drank alcohol at the time of their initial assessment and 0 if not. Participants with missing data for any of the above covariates were excluded ($n = 850$). These scalar covariates were added to our model as linear terms. PA data were further compressed by obtaining the mean and standard deviation of the minute-level $\log(1+\text{ENMO})$ values across days at each minute s of the day over all available days. For subject i , denote by $X_i(s)$ the mean of $\log(1+\text{ENMO})$ at minute s , and by $U_i(s)$ the standard deviation of $\log(1+\text{ENMO})$ at minute s .

of the day. With this notation our M1 model becomes

$$\log \lambda_i(t|\mathbf{Z}_i, \mathbf{X}_i) = \log \lambda_0(t) + \mathbf{Z}_i^T \boldsymbol{\gamma} + f_1(\text{age}_i) + \int_S X_i(s) \beta_X(s) ds + \int_S U_i(s) \beta_U(s) ds, \quad (3.5)$$

where $f_1(\cdot)$, $\beta_X(\cdot)$, and $\beta_U(\cdot)$ are unknown smooth function that are estimated using penalized splines (Eilers and Marx, 1996; Ruppert *and others*, 2003; Wood, 2017).

Suppose that `data` is a data frame organized as follows: `time` is the observed time Y_i , `status` is the event indicator Δ_i for mortality. The different scalar terms, $\mathbf{Z}_i$, are stored in their respective variable names. $X_i(s)$ and $U_i(s)$ are stored as two matrices of the same dimensions, `X` and `U`, where each row corresponds to $X_i(s)$ and $U_i(s)$, respectively. Similar to the general syntax introduced in Section 2.1, we first create two matrices, `atmat` and `almat`, which are used to integrate the rows of `X` and `U` over time.

```
data$almat <- I(matrix(1 / ncol(data$X),
                      ncol = ncol(data$X), nrow = nrow(data)))

data$atmat <- I(matrix(1:ncol(data$sFun),
                      ncol = ncol(data$sFun), nrow = nrow(data), byrow = TRUE))
```

The model was fit using the `mgcv::gam` function as follows.

```
M1 <- gam(time ~ drinking + Townsend + cancer +
           smoking + chd + overall_health + sex +
           s(age, k = 30, bs = "cr") +
           s(atmat, by = almat * X, bs = "cc", k = k) +
           s(atmat, by = almat * U, bs = "cc", k = k),
           data = data, weights = status, family = cox.ph())
```

Figure 2 displays the estimated nonlinear effects of age, $f_1(s)$, functional coefficients $\beta_X(s)$ and $\beta_U(s)$, and their corresponding unadjusted and CMA confidence intervals and p-values across time of day. The 95% pointwise unadjusted and CMA confidence intervals are shown as dark gray and light gray regions, respectively. The CMA confidence intervals are uniformly wider than the

pointwise unadjusted confidence intervals, which is expected, as discussed in Section 2.2. Similarly, the CMA p-values are larger than the pointwise unadjusted p-values.

The top panel in Figure 2 displays the associations between age and log hazard of all-cause mortality. With increased age, the risk of mortality also increases and is statistically significant (p-value < 0.0001). This relationship appears linear and every one-year increase in age is associated with approximately 7.1% increase in the hazard of mortality.

The second row of panels in Figure 2 displays the point estimators of $\beta_X(s)$ and $\beta_U(s)$, which are the effects of the mean and standard deviation of physical activity at the time s of the day. The dark gray areas correspond to the 95% pointwise unadjusted confidence intervals, while the light gray areas correspond to CMA confidence intervals. Based on the CMA confidence intervals for $\beta_X(s)$, we conclude that higher physical activity between 12pm and 7pm is significantly associated with a lower risk of all-cause mortality at $\alpha = 0.05$ even after accounting for other covariates and the standard deviation of physical activity. Note that this is slightly different from our previous results (Leroux *and others*, 2019), which indicate that physical activity between 8am and 8pm is significantly associated with lower risk of mortality. This is likely due to the fact that the current model accounts for the day-to-day variability, which is statistically significant exactly during the morning hours. These results are new and seem to indicate that higher day-to-day variability in the morning and higher afternoon physical activity are associated with lower risk of mortality. It is not immediately clear whether this provides new insights into “what to do”, but results may be due to the fact that physical activity is already relatively high in the morning for most individuals.

The CMA and unadjusted p-values for $X_i(s)$ and $U_i(s)$ were obtained by inverting the respective confidence intervals and shown in Figure 2. In addition to calculating a p-value for every time point, we also obtain the global p-value defined in Section 2.2. For $X_i(s)$, this value was smaller than 0.0001 and for $U_i(s)$, this value was 0.0016.

For scalar predictors, we observed that in individuals 50 years and older, men at birth have 51.8% increased risk of mortality (p-value < 0.0001) compared to women. The effect of Townsend deprivation index is also significant (p-value = 0.0282) with an increase of one unit being associated with a 2.6% increase in the risk of mortality. For reference, the Townsend index in our data varies between -6.26 and 10.59, has a standard deviation of 2.81 with higher values corresponding to higher poverty. Likewise, past cancer diagnosis increases the risk of mortality by 140% (p-value < 0.0001) and coronary heart disease increases it by 35.2% (p-value = 0.0034). Participants that were previous and current smokers have 15.9% (p-value = 0.0449) and 119% (p-value < 0.0001) increased risk of mortality compared to never smokers. Participants that have never drunk alcohol have 16.3% lower risk of mortality, though this was not statistically significant in our data set (p-value = 0.0668) compared to those who drink. Finally, participant who self-reported overall significant health problems had a higher risk of mortality. Those who rated themselves to have excellent, good, and fair health have 45.1% (p-value = 0.0128), 41.3% (p-value = 0.0149), 31.4% (p-value = 0.0264) lower risk of mortality compared to individuals who rated themselves as having poor health status, respectively. The remaining scalar coefficients were not statistically significant.

Next, we applied the distributional Cox models introduced in Section 2.3 to explore the association between quantiles of physical activity and mortality. This model, labeled M2, adjusted for the same set of scalar predictors as model M1. Although the formulas look formally similar, model M2 provides different insights into how different intensity levels of physical activity are associated with the hazard of mortality. For model M2 we used a non-cyclic cubic penalized with 119 knots to model the $\beta_Q(s)$ functional coefficient as the effect is not expected to be similar at $p = 0$ and $p = 1$. For each individual, a subject-specific quantile function was estimated and used to fit the model. The estimation results are shown on the right panel of Figure 3. The effect of quantile of physical activity is estimated to be linear in p , it is positive for $p \in [0, 0.38]$ and

negative for $p \in (0.38, 1]$. We have not assumed linearity in the effect, so the result seems to suggest that $\beta_Q(p)$ could be safely assumed to be linear. We do not conduct here formal tests of linearity as they exceed the scope of the current paper. Using CMA confidence intervals, we conclude that higher values of physical activity quantiles at $p = 0.45$ or higher and lower values of physical activity quantiles at $p = 0.3$ or lower are associated with lower risk of all-cause mortality (global p-value < 0.0001). A practical translation of these results is that being more active during the most active 10.8 hours of the day ($0.55 \times 24 = 13.2$) and being less active during the least active 7.2 (0.3×24) hours of the day is associated with lower risk of all-cause mortality. Interestingly, one expects individuals to sleep between 6 and 8 hours every day, which is likely to be the period of lowest physical activity. This finding complements the results based on the analysis of the time series of physical activity. In particular, it suggests that restful sleep/inactivity and high physical activity during periods of activity are associated with reduced risk of mortality. We believe that this is the first time such results are reported using CMA inference.

Similar to M1, several scalar predictors had statistically significant effects. Age (p-value < 0.0001) and Townsend deprivation index (p-value = 0.0158) are significantly and positively associated with the risk of mortality. Having coronary heart disease (p-value = 0.0036) or a biological male assignment at birth (p-value < 0.0001) also significantly increase the risk of mortality. Previous smokers (p-value = 0.0384) and current smokers (p-value < 0.0001) are at significantly higher risk than never smokers. Individuals who self-reported as having excellent (p-value = 0.0109), good (p-value = 0.0123), or fair (p-value = 0.0199) overall health have significantly lower risk of mortality compared to individuals who rated themselves as having poor health status, respectively.

Overall, these results highlight that the extended functional and distribution Cox models provide an efficient and interpretable solution to analyze the effects of physical activity on health outcomes in large-cohort studies. The applications of M1 and M2 reveal highly interpretable

and complementary findings. The models described in this paper would not be computationally feasible without the rapid development of statistical software. Moreover, the syntax presented in our paper is reproducible and can be easily modified to include different scalar, additive, and functional terms depending on the problem setting.

4. SIMULATIONS

Simulations were conducted to: (1) provide a general framework for simulating functional predictors and survival outcomes from biobank-scale datasets; (2) evaluate the inferential properties of the proposed methods. For simplicity, we focus on the basic form of the model, with no scalar predictors and one functional predictor.

4.1 *Synthetic Data Generation*

We start by creating synthetic data sets that contain physical activity data and time-to-event outcomes that are as close as possible to the structure of UK Biobank. This is very difficult, as simulations from parametric survival distributions did not work due to their extreme sensitivity to small changes in parameters. To address this problem, we conducted simulations based on the estimated survival function obtained from fitting the extended functional Cox models to the UK Biobank data Cui *and others* (2021). While these methods have been used in simpler models and smaller data sets, this was the first time that the extended functional Cox model was applied to a data set of the size and complexity of UK Biobank.

4.1.1 *Simulating Functional Predictors* The first step was to simulate functional data that is similar to the physical activity data in the UK Biobank. Denote by $X_i(s)$ the subject-specific minute level average physical MIMS and let $\hat{\mu}(s)$ be an estimator of the mean of $X_i(s)$. This could be obtained by taking a point-wise average, as the number of functions is very large, but one

could also further smooth this average. Using Functional Principal Component Analysis (FPCA) we estimate the eigenfunctions, $\hat{\phi}_k(s)$, and eigenvalues, $\hat{\lambda}_k$, of the $X_i(s) - \hat{\mu}(s)$ process. This was achieved using the Fast Covariance Estimation (FACE) (Xiao *and others*, 2016) implemented in the `refund::fpca.face` function Goldsmith *and others* (2024). Physical activity data was then generated from the model $\hat{X}_i(s) = \hat{\mu}(s) + \sum_{k=1}^M \xi_{ik} \hat{\phi}_k(s)$, where $\xi_{ik} \sim \mathcal{N}(0, \hat{\lambda}_k)$. Other distributions could be used, but here we use the normal distribution for simplicity. The non-negativity of the simulated data can be achieved by first applying the framework to log-transformed real data and then transforming the simulated data back to the original scale. For the purpose of this simulation we chose $M = 35$, which ensured that 99% of the variability was explained.

4.1.2 Simulating Survival Outcomes To simulate survival outcomes, we use the following framework described in Cui *and others* (2021). This method requires a functional Cox regression model fit on the real data, the real survival outcome data, and the simulated functional predictors. First, the estimated cumulative baseline hazard, $\hat{\Delta}_0(t) = \int_0^t \hat{\lambda}_0(u) du$, was extracted from the fitted model. Next, the linear predictor $\hat{\eta}_i = \int_{\mathcal{S}} \hat{X}_i(s) \hat{\beta}(s) ds$ was generated for each subject, where $\hat{X}_i(s)$ was simulated as described in Section 4.1.1 and $\hat{\beta}(s)$ was obtained from the fitted model. After obtaining $\hat{\Delta}_0(t)$ and $\hat{\eta}_i$, the estimated survival function is calculated as $\hat{S}_i(t) = \exp\{-e^{\hat{\eta}_i} \hat{\Delta}_0(t)\}$. The survival time, $\tilde{T}_i$, is then simulated from a random variable with cumulative distribution function $1 - \hat{S}_i(t)$. The censoring time $\tilde{C}_i$ was generated by sampling with replacement from the censoring times in the UK Biobank. The simulated survival outcomes are obtained as $\tilde{Y}_i = \min(\tilde{T}_i, \tilde{C}_i)$ and $\tilde{\Delta}_i = \mathbb{1}(\tilde{T}_i \leq \tilde{C}_i)$, respectively.

4.2 Estimation Accuracy

The aim in this section is to evaluate how close are the results based on the simulated data to the results based on the real UK Biobank data. To do this, we fit a linear functional Cox model to

the simulated data and compare the functional coefficient, $\beta(s)$, to the true coefficient. Here the true coefficient, denoted as $\hat{\beta}_{\text{UKB}}(s)$, is the coefficient obtained from the real UK Biobank data where the total sample size was $N = 93,370$ and the number of events was 931. We will simulate data of various total sample sizes, while maintaining the same proportion of events (about 1%).

We conducted 200 simulations for each of the following sample sizes: $N = 50,000$, 100,000, 500,000, and 1,000,000. All simulations were conducted conditional on $\hat{\beta}_{\text{UKB}}(s)$, which is considered to be the “true functional coefficient”. Running such simulations is non-trivial, especially for larger sample sizes. For example, the 200 simulations for the M1 model with $N = 1,000,000$ observations required up to 400GB of memory and 2.5 days (about 18 minutes per data set) on a cluster computing system. The code for our simulation study is available in our Github repository.

The average estimated coefficients from simulated data with different sample sizes are shown in different colors in Figure 4. The true functional coefficient, $\hat{\beta}_{\text{UKB}}(s)$, is shown as the black line. The estimated average coefficients from simulated data are close to the true coefficient and get closer as sample size increases. This suggests that estimators have small bias and the bias gets smaller with increased sample size. To further quantify the estimation accuracy of simulated data, the mean squared error (MSE) and the empirical coverage of the 95% pointwise and CMA confidence intervals were calculated for each sample size and the results are shown in Table 1. As the sample size increases, the MSE decreases. The MSE value is 0.004 for $N = 1,000,000$, compared to 0.291 for $N = 50,000$, or less than 2%. Furthermore, as the sample size increases, the empirical coverage of both of the pointwise and CMA confidence intervals increases. However, they are both below the nominal level for $N = 100,000$, which is the approximate size of the original sample. Figure 5 displays the empirical coverage of the pointwise 95% confidence intervals along the functional domain. For smaller sample sizes, under coverage was observed at several locations of the domain where there is a change in the curvature. Thus, it is likely that under

Table 1. The mean squared error and mean confidence interval (CI) coverage of the estimated functional coefficient based on 200 simulations with different sample sizes and proportion of cases.

Sample Size, N	Mean Squared Error	Pointwise CI Coverage	CMA CI Coverage
50,000	0.291	87.7%	74.5%
100,000	0.312	91.3%	81.5%
500,000	0.006	94.8%	93.5%
1,000,000	0.004	95.0%	92.0%

coverage may be partially attributed to biases in the point estimates in those regions, though non-normality of estimators could also be a factor. Note that CMA confidence intervals have lower coverage than the average coverage of pointwise confidence intervals. This is expected because average pointwise confidence intervals is the average over all time points of the proportion of simulations where the pointwise confidence intervals cover $\hat{\beta}_{\text{UKB}}(s)$. In contrast, the coverage of CMA confidence intervals is defined as the proportion of times out of 200 simulations when the confidence interval covers $\hat{\beta}_{\text{UKB}}$ at all time points. Thus, the CMA performance is more sensitive to bias anywhere along the time range. In fact this is more comparable to the minimum coverage of the pointwise estimators than the average; see Figure 5.

5. DISCUSSION

Inferential methods for survival analysis were extended to incorporate multiple functional predictors, nonparametric effects of scalar covariates, and parametric effects. For the first time, correlation and multiplicity adjusted (CMA) confidence bands and p-values for the functional coefficients were calculated. Distributional predictors (where the functional data are replaced by the quantile function of the observed values) were also incorporated and treated as a standard scalar-on-function regression model. Models were applied to the UK Biobank data set, which is one order of magnitude larger than previous data sets used for functional survival models Crainiceanu *and others* (2024). The associated software is easy to use and scalable as it takes around 20 minutes

to fit the model to the UK Biobank. For reproducibility purposes, the associated R code for our analysis is available at https://github.com/ang-zhao/biobank_functional_cox_mortality.

An ever larger body of evidence indicates that summaries of objectively measured physical activity are the strongest predictor of mortality and many diseases in people over 50 Cui *and others* (2022b); Ledbetter *and others* (2022); Leroux *and others* (2021); Meng *and others* (2023); Smirnova *and others* (2020); Zhao *and others* (2023). Thus, there is a need to investigate whether the daily patterns of PA contain more information than these summaries. This paper validates some of the published scientific results in a much larger data set and provides new scientific insights: (1) increased physical activity during the day and reduced physical activity during the night are associated with lower risk of death; (2) higher day to day variability in the morning conditional on the same level of activity in the morning is associated with lower mortality risk; (3) conditional on the functional effects of the mean and standard deviation of physical activity, the effect of age on the log hazard is indistinguishable from being linear; (4) when distributional predictors are used, the functional effect is estimated to be linear, positive below quantile $p = 0.30$ (consistent with positive effects of less physical activity during rest periods) and negative above quantile $p = 0.55$ (consistent with positive effects of more physical activity during active periods); and (5) even with very large studies, estimating and quantifying the inferential properties of estimators is important, especially when the number of outcomes is small (e.g., 1% of the total sample size).

Another important contribution of the paper is to show how to simulate survival data with functional predictors that closely mimic the UK Biobank data. Simulations indicate that inference results are close to their nominal level for sample sizes close to the UK Biobank and estimation and coverage probabilities continue to improve with increased sample size at the same percentage of observed outcomes (1%).

Functional regression models, and applied functional data analysis more generally, like other

statistical methods, will continue to be successful if it can adapt to larger and more complex data, provide insights that may not be available from data summaries, or suggest simpler models without substantial loss of information. Here we propose a set of analytic tools that satisfies these requirements, which provides a necessary step forward in the analysis of high-dimensional functional data.

REFERENCES

- AGARWALA, N., ZIPUNNIKOV, V., CRAINICEANU, C.M. AND LEROUX, A. (2024). Quantifying the time-varying association between objectively measured physical activity and mortality in us older adults over a 12-year follow-up period: the NHANES 2003–2006 study. *BMJ Evidence-Based Medicine*.
- ALLEN, N., SUDLOW, C., DOWNEY, P., PEAKMAN, T., DANESH, J., ELLIOTT, P., GALLACHER, J., GREEN, J., MATTHEWS, P., PELL, J., SPROSEN, T. *and others*. (2012). UK biobank: Current status and what it means for epidemiology. *Health Policy and Technology*, 123–126.
- BAKRANIA, K., YATES, T., ROWLANDS, A. V., ESLIGER, D. W., BUNNEWELL, S., SANDERS, J., DAVIES, M., KHUNTI, K. AND EDWARDSON, C. L. (2016). Intensity thresholds on raw acceleration data: Euclidean norm minus one (enmo) and mean amplitude deviation (mad) approaches. *PLOS One* **11**(10), e0164045.
- BORNKAMP, B. (2018). Calculating quantiles of noisy distribution functions using local linear regressions. *Computational Statistics* **33**, 487–501.
- CHEN, Z., CHEN, J., COLLINS, R., GUO, Y., PETO, R., WU, F., LI, L. AND CHINA KADOORIE BIOBANK (CKB) COLLABORATIVE GROUP. (2011). China kadoorie biobank of 0.5 million

- people: survey methods, baseline characteristics and long-term follow-up. *International Journal of Epidemiology* **40**(6), 1652–1666.
- CRAINICEANU, C. M. AND GOLDSMITH, J. (2010). Bayesian functional data analysis using WinBUGS. *Journal of Statistical Software* **32**(11), 1–33.
- CRAINICEANU, C. M., GOLDSMITH, J., LEROUX, A. AND CUI, E. (2024). *Functional Data Analysis with R*. Chapman and Hall/CRC.
- CRAINICEANU, C. M., STAICU, A., RAY, S. AND PUNJABI, N. (2012). Bootstrap-based inference on the difference in the means of two correlated functional processes. *Statistics in Medicine* **31**(26), 3223–3240.
- CUI, E. (2023). Functional data analysis methods for large scale physical activity studies [Ph.D. Thesis]. J.s Hopkins University.
- CUI, E., CRAINICEANU, C. M. AND LEROUX, A. (2021). Additive functional cox model. *Journal of Computational and Graphical Statistics* **30**(3), 780–793.
- CUI, E., LEROUX, A., SMIRNOVA, E. AND CRAINICEANU, C. M. (2022a). Fast univariate inference for longitudinal functional models. *Journal of Computational and Graphical Statistics* **31**(1), 219–230.
- CUI, E., LI, R., CRAINICEANU, C. M. AND XIAO, L. (2023). Fast multilevel functional principal component analysis. *Journal of Computational and Graphical Statistics* **32**(2), 366–377.
- CUI, E., THOMPSON, E. C., CARROLL, R. J. AND RUPPERT, D. (2022b). A semiparametric risk score for physical activity. *Statistics in Medicine* **41**(7), 1191–1204.
- DI, J., LEROUX, A., URBANEK, J., VARADHAN, R., SPIRA, A. P., SCHRACK, J. AND ZIPUNNIKOV, V. (2017). Patterns of sedentary and active time accumulation are associated with mortality in us adults: The NHANES study. *BioRxiv*, 182337.

- DOHERTY, A., JACKSON, D., HAMMERLA, N., PLOTZ, 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 others.* (2017). Large scale population assessment of physical activity using wrist worn accelerometers: The UK biobank study. *PLOS One* **12**(2), e0169649.
- EILERS, P. H. C. AND MARX, B. D. (1996). Flexible smoothing with b-splines and penalties. *Statistical Science* **11**(2), 89–121.
- GELLAR, J. E., COLANTUONI, E., NEEDHAM, D. M. AND CRAINICEANU, C. M. (2015). Cox regression models with functional covariates for survival data. *Statistical Modelling* **15**(3), 256–278.
- GHOSAL, R., GHOSH, S. K., SCHRACK, J. A AND ZIPUNNIKOV, V. (2023a). Distributional outcome regression and its application to modelling continuously monitored heart rate and physical activity. *arXiv preprint arXiv:2301.11399*.
- GHOSAL, R., VARMA, V. R., VOLFSO, D., HILLEL, I., URBANEK, J., HAUSDORFF, J. M., WATTS, A. AND ZIPUNNIKOV, V. (2023b). Distributional data analysis via quantile functions and its application to modeling digital biomarkers of gait in alzheimer’s disease. *Biostatistics* **24**(3), 539–561.
- GHOSAL, R., VARMA, V. R., VOLFSO, D., URBANEK, J., HAUSDORFF, J. M., WATTS, A. AND ZIPUNNIKOV, V. (2022). Scalar on time-by-distribution regression and its application for modelling associations between daily-living physical activity and cognitive functions in alzheimer’s disease. *Scientific Reports* **12**(1), 11558.
- GOLDSMITH, J., BOBB, J., CRAINICEANU, C. M., CAFFO, B. AND REICH, D. (2011). Penalized functional regression. *Journal of Computational and Graphical Statistics* **20**(4), 830–851.

- GOLDSMITH, J., SCHEIPL, F., HUANG, L., WROBEL, J., DI, C., GELLAR, J., HAREZLAK, J., MCLEAN, M. W., SWIHART, B., XIAO, L., CRAINICEANU, C. M., REISS, P. T. *and others*. (2024). *refund: Regression with Functional Data*. R package version 0.1-36.
- GRAY, R. J. (1992). Flexible methods for analyzing survival data using splines, with applications to breast cancer prognosis. *Journal of the American Statistical Association* **87**(420), 942–951.
- HAO, M., LIU, K., XU, W. AND ZHAO, X. (2021). Semiparametric inference for the functional cox model. *Journal of the American Statistical Association* **116**(535), 1319–1329.
- KOKOSZKA, P. AND REIMHERR, M. (2017). *Introduction to functional data analysis*. Chapman and Hall/CRC.
- KONG, D., IBRAHIM, J. G., LEE, E. AND ZHU, H. (2018). Flcrm: Functional linear cox regression model. *Biometrics* **74**(1), 109–117.
- LEDBETTER, M. K., TABACU, L., LEROUX, A., CRAINICEANU, C. M. AND SMIRNOVA, E. (2022). Cardiovascular mortality risk prediction using objectively measured physical activity phenotypes in nhanes 2003–2006. *Preventive medicine* **164**, 107303.
- LEROUX, A. (2020). Statistical methods for the analysis of functional data under models with complex association structures [Ph.D. Thesis]. J.s Hopkins University.
- LEROUX, A., DI, J., SMIRNOVA, E., MCGUFFEY, E. J., CAO, Q., BAYATMOKHTARI, E., TABACU, L., ZIPUNNIKOV, V., URBANEK, J. K AND CRAINICEANU, C. M. (2019). Organizing and analyzing the activity data in nhanes. *Statistics in biosciences* **11**, 262–287.
- LEROUX, A., XU, S., KUNDU, P., MUSCHELLI, J., SMIRNOVA, E., CHATTERJEE, N. AND CRAINICEANU, C. M. (2021). Quantifying the predictive performance of objectively measured physical activity on mortality in the UK Biobank. *The Journals of Gerontology: Series A* **76**(8), 1486–1494.

- MENG, Q., CUI, E., LEROUX, A., MOWRY, E. M., LINDQUIST, M. A. AND CRAINICEANU, C. M. (2023). Quantifying the association between objectively measured physical activity and multiple sclerosis in the UK Biobank. *Medicine and Science in Sports and Exercise* **55**(12), 2194–2202.
- PARK, S. Y., STAICU, A., XIAO, L. AND CRAINICEANU, C. M. (2018). Simple fixed-effects inference for complex functional models. *Biostatistics* **19**(2), 137–152.
- PETERSEN, A., ZHANG, C. AND KOKOSZKA, P. (2022). Modeling probability density functions as data objects. *Econometrics and Statistics* **21**, 159–178.
- QU, S., WANG, J. AND WANG, X. (2016). Optimal estimation for the functional cox model. *The Annals of Statistics* **44**(4), 1708–1738.
- RAMSAY, J. O. AND SILVERMAN, B. W. (2005). *Functional Data Analysis*. Springer.
- RUPPERT, D., WAND, M. P. AND CARROLL, R. J. (2003). *Semiparametric Regression*, Number 12. Cambridge University Press.
- SERGAZINOV, R., LEROUX, A., CUI, E., CRAINICEANU, C. M., 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**(4), 3873–3882.
- SMIRNOVA, E., LEROUX, A., CAO, Q., TABACU, L., ZIPUNNIKOV, V., CRAINICEANU, C. M. 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**(9), 1779–1785.
- TABACU, L., LEDBETTER, M., LEROUX, A., CRAINICEANU, C.M. AND SMIRNOVA, E. (2020). Quantifying the varying predictive value of physical activity measures obtained from wearable

- accelerometers on all-cause mortality over short to medium time horizons in nhanes 2003-2006. *Sensors* **21**(1), 1779–1785.
- THE ALL OF US RESEARCH PROGRAM INVESTIGATORS. (2019). The “all of us” research program. *New England Journal of Medicine* **381**, 668–676.
- VAN HEES, V. T., GORZELNIAK, L., DEAN LEÓN, E. C., EDER, M. A., PIAS, M., TAHERIAN, S., EKELEND, U., RENSTRÖM, F., FRANKS, P. W., HORSCH, A. *and others*. (2013). Separating movement and gravity components in an acceleration signal and implications for the assessment of human daily physical activity. *PLOS One* **8**(4), e61691.
- VERWEIJ, P. J. M. AND VAN HOUWELINGEN, H. C. (1994). Penalized likelihood in cox regression. *Statistics in Medicine* **13**(23-24), 2427–2436.
- WOOD, S. N. (2017). *Generalized Additive Models: an Introduction with R*. CRC Press.
- XIAO, L., ZIPUNNIKOV, V., RUPPERT, D. AND CRAINICEANU, C.M. (2016). Fast covariance estimation for high-dimensional functional data. *Statistics and Computing* **26**(1), 409–421.
- ZHAO, A., CUI, E., LEROUX, A., LINDQUIST, M. A. AND CRAINICEANU, C. M. (2023). Evaluating the prediction performance of objective physical activity measures for incident parkinson’s disease in the uk biobank. *Journal of Neurology* **270**(12), 5913–5923.




Fig. 2. Top panel: estimated effects of age, $f_1(\text{age})$, in the M1 model modeled using penalized splines. Second row: estimated functional coefficients, $\beta_X(s)$ (left panel) and $\beta_U(s)$ (right panel) together with 95% pointwise unadjusted (dark gray) and CMA (light gray) confidence intervals. The dashed red line corresponds to 0. Third row: pointwise unadjusted (dark solid lines) and CMA (light gray lines) p-values corresponding to $\beta_X(s)$ (left panel) and $\beta_U(s)$ (right panel).




Fig. 3. Left panel: estimated functional coefficient $\beta_Q(p)$ from the M2 model as a function of the probability p . The model accounts for the list of covariates described in section 3.3. The 95% pointwise and CMA confidence intervals are shown by dark gray and light gray areas, respectively. A dotted red line is shown at 0. Right panel: Pointwise unadjusted and CMA p-values.




Fig. 4. The black line represents the estimated functional coefficient from a model fitted on the real UK Biobank data. The functional coefficient from models fitted on simulated data of different sample sizes from 200 simulations are plotted in color: $N = 50,000$ (red line), $100,000$ (orange line), $500,000$ (blue line), and $1,000,000$ (purple line).




Fig. 5. The empirical coverage of the pointwise 95% confidence intervals at each location of the functional domain under different sample sizes from 200 simulations.