Computational Statistics
https://doi.org/10.1007/s00180-024-01591-1

ORIGINAL PAPER

Analysing kinematic data from recreational runners using functional data analysis

Edward Gunning[1] · Steven Golovkine[1] · Andrew J. Simpkin[2] · Aoife Burke[3] · Sarah Dillon[3,4] · Shane Gore[3,5] · Kieran Moran[3,5] · Siobhan O'Connor[3] · Enda White[3] · Norma Bargary[1]

Received: 24 November 2023 / Accepted: 13 December 2024

Abstract

We present a multivariate functional mixed effects model for kinematic data from a large number of recreational runners. The runners' sagittal plane hip and knee angles are modelled jointly as a bivariate function with random effects functions accounting for the dependence among bilateral measurements. The model is fitted by applying multivariate functional principal component analysis (mv-FPCA) and modelling the mv-FPCA scores using scalar linear mixed effects models. Simulation and bootstrap approaches are introduced to construct simultaneous confidence bands for the fixed effects functions, and covariance functions are reconstructed to summarise the variability structure in the data and thoroughly investigate the suitability of the proposed model. In our scientific application, we observe a statistically significant effect of running speed on both joints. We observe strong within-subject correlations, reflecting the highly idiosyncratic nature of running technique. Our approach is applicable to modelling multiple streams of smooth biomechanical data collected in complex experimental designs.

Keywords Biomechanics · Functional data analysis · Mixed-effects model · Multivariate functional data

✉ Edward Gunning
edward.gunning@pennmedicine.upenn.edu

[1] MACSI, Department of Mathematics and Statistics, University of Limerick, Limerick, Ireland

[2] School of Mathematical and Statistical Sciences, University of Galway, Galway, Ireland

[3] Centre for Injury Prevention and Performance, Athletic Therapy and Training; School of Health and Human Performance, Dublin City University, Dublin, Ireland

[4] School of Allied Health, Faculty of Education and Health Science, University of Limerick, Limerick, Ireland

[5] Insight SFI Research Centre for Data Analytics, Dublin City University, Dublin, Ireland

Published online: 06 February 2025

1 Introduction

Advances in data collection, processing and storage technologies have led to an increased volume of data produced for biomechanics and human movement research (Ferber et al. 2016). Forces (kinetics) or displacement (kinematics) are measured hundreds or thousands of times per second during a single movement, leading to datasets characterised by high-dimensional observations. Functional data analysis (FDA) (Ramsay and Silverman 2005) is particularly well-suited to modelling human movement data as it treats a time series of kinetic or kinematic data as a single function (or curve) rather than as a sequence of discrete measurements. This allows a more comprehensive analysis than reducing the time series to a single summary value (e.g., peak angle) or ignoring the time dependence in the high-dimensional sequences of measurements (Hébert-Losier et al. 2015; Pataky et al. 2015; Warmenhoven et al. 2021). Applications of FDA in biomechanics and human movement research include: describing the effects of orthoses on running or walking (Coffey et al. 2011; Zhang et al. 2017), clustering runners according to footfall pattern (Liebl et al. 2014), predicting fatigue in recreational athletes (Wu et al. 2019) and classifying different forms of activity (Aguilera-Morillo and Aguilera 2020).

Our motivating dataset comes from the Dublin City University (DCU) Running Injury Surveillance Centre (RISC) study, which aims to investigate the relationship between clinical and biomechanical variables and running-related injuries (RRIs) among novice and recreational runners. Although recreational running is one of the most popular recreational hobbies in the world and it provides substantial positive benefits for health and well-being, RRIs present a considerable barrier to participation and several other negative consequences, e.g., negative health aspects and financial costs (Hespanhol Junior et al. 2017). Despite this, our understanding of RRIs is limited, especially with respect to biomechanical factors, which has motivated studies to investigate the relationship between biomechanical variables and RRIs in populations of recreational runners. In particular, there has been a large focus on the population of recently-injured runners (see, e.g., Bramah et al. 2018; Becker et al. 2017; Mann et al. 2015), as history of a recent RRI is the strongest risk factor for suffering a new one. It is hypothesised that recently-injured runners might retain some of the movement characteristics that contributed to the previous injury, or adopt compensatory mechanisms that cause them to be re-injured (Saragiotto et al. 2014). Findings of these studies have largely been conflicting, in part perhaps because they have employed traditional statistical techniques using discrete kinematic variables (e.g., Ceyssens et al. 2019; Willwacher et al. 2022). The ability of FDA methods to preserve the salient structure in time-dependent biomechanical data could lead to more comprehensive analyses that improve our understanding of RRIs and biomechanical factors.

Male and female runners between 18 and 64 years of age participated in the RISC study. Whole-body kinematic data were recorded during a three-minute treadmill run, where the participant ran at a self-selected speed that reflected their

typical training pace. In addition, they completed a survey detailing their demographics, injury history (i.e., retrospective injury information) and training habits and were monitored for the occurrence of RRIs for a 12-month period (i.e., prospective injury information); see Table 1 for summary characteristics of the participants in the dataset. For this dataset, the relationship between injury history and scalar clinical (Dillon et al. 2021) and scalar biomechanical (Burke et al. 2022) variables has been examined, but approaches that preserve the full biomechanical time series data have not been employed. Focusing on the hip and knee angles in the sagittal plane (Fig. 1), we aim to to characterise the effect of retrospective injury status on the full biomechanical time series, while accounting for and understanding the effects of other factors, e.g., sex, running speed and age. *Function-on-scalar regression models* (Faraway 1997; Ramsay and Silverman 2005) are an appropriate tool for characterising these relationships, where the biomechanical time series' are treated as the functional response variable(s), modelling their dependence on scalar covariates, e.g., injury status, sex, running speed and age. Morris (2015, Sect. 5) provides a comprehensive review of conventional function-on-scalar regression models.

Conventional function-on-scalar regression models assume independent observations, and do not handle dependence induced by repeated observations from the same individual. However, these dependencies frequently arise in biomechanics for a number of reasons, e.g., multiple strides, trials or repetitions of a movement, or measurements from both sides of the body. In our case, although we have computed an average of all strides on the right and left side separately (Fig. 1), further averaging across the right and left sides to produce a single bilateral average curve could lead to a substantial loss of information and it could potentially bias subsequent analyses if large asymmetries exist. *Functional mixed effects* (or *multilevel*) models, which are the analogue of classical scalar mixed effects models (Laird and Ware 1982; Bates et al. 2015), extend conventional function-on-scalar regression models to handle repeated measures settings and more complex dependence structures. The literature on functional mixed effects models is rich—early pioneering work was by

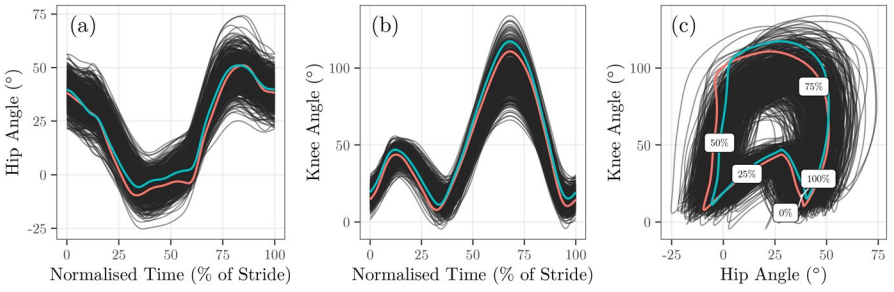

Fig. 1 The dataset used in this analysis. **a** The hip angle functions. **b** The knee angle functions. **c** The knee angle functions plotted against the hip angle functions in an angle-angle diagram. In each plot, the right and left side observations for a single participant are highlighted in turquoise and red, respectively. The data have been time normalised and registered in preparation for analysis as described in Sect. 3.1, and are evaluated on a grid of 101 points $t = 0, 1, \ldots, 100$ for plotting

Table 1 Summary characteristics of the participants in the RISC dataset included in this analysis

	Mean	Std. Dev.	
Speed (km h^{-1})	11.0	1.6	
Age (years)	43.3	9.0	
Weight (kg)	72.4	12.9	
Height (cm)	172.9	9.7	
		N	(%)
Retrospective Injury Status	Never Injured	50	17.4
	Injured > 2 yr. ago	67	23.3
	Injured 1 – 2 yr. ago	51	17.7
	Injured < 1 yr. ago	120	41.7
Sex	Male	176	61.1
	Female	112	38.9

Morris and Carroll (2006); Guo (2002), later developments by Scheipl et al. (2015); Cui et al. (2022), reviews are provided by Morris (2015); Liu and Guo (2012); Morris (2017) and a recent application in running biomechanics by Matabuena et al. (2023).

Rather than fitting separate (univariate) functional mixed effects models to the data from the knee and hip, it makes sense from a methodological and applied perspective to model them collectively. From a statistical perspective, sharing information among functional variables can lead to improved parameter estimates (Volkmann et al. 2021; Zhu et al. 2017), and from a biomechanical perspective it is preferable to model and interpret the knee and hip jointly (Fig. 1 (c)) because they work together as parts of a system and understanding their interaction (i.e., coordination) is crucial (Glazier 2021). *Multivariate* functional data analysis techniques (see, e.g., Gòrecki et al. 2018) concern the analysis of multiple functional variables (e.g., the knee and hip angles), and they have been shown to be useful for understanding co-ordination among multiple joints in sports biomechanics (Ryan et al. 2006; Trounson et al. 2020). *Multivariate* (or *multiple-response*) functional mixed effects models extend classical univariate functional mixed effects models to handle multiple functional variables as outcomes. However, the literature on these models is more scarce than in the univariate case[1], with just three main approaches proposed (Goldsmith and Kitago 2016; Volkmann 2021; Zhu et al. 2017) and they have yet to be applied in human-movement/running biomechanics.

Goldsmith and Kitago (2016) developed a bespoke bivariate functional mixed effects model for kinematic data from a motor control experiment, where linear fixed effects of scalar covariates and subject-specific random effects were modelled using penalised splines. Their model was fitted in a Bayesian framework (using both

[1] Methods for multivariate functional regression and inference *without* random effects/multilevel structures have been developed by Jiang et al. (2022); Diquigiovanni et al. (2022); Zhu et al. (2022); Liu et al. (2022); Li and Xiao (2023).

variational approximations and full Markov Chain Monte Carlo (MCMC) sampling). Volkmann et al. (2021) proposed an alternative approach, by extending the univariate Functional Additive Mixed Model (FAMM) to the multivariate setting. In this model, smooth non-linear effects of scalar covariates and multiple layers of random effects were modelled using penalised splines and multivariate Functional Principal Components (mv-FPCs), respectively. It is fitted in a Frequentist framework by recasting the functional model as a large scalar additive mixed model and using the **mgcv** software (Wood 2011), readily accommodating functions that are sparsely or irregularly measured with error. Finally, Zhu et al. (2017) extended the Bayesian Functional Mixed Model (BayesFMM) *basis modelling* approach of Morris and Carroll (2006) to handle multivariate functional data. Their approach involves projecting each multivariate functional observation onto a set of basis functions and then modelling each basis coefficient separately using Bayesian scalar linear mixed effects models. This “divide and conquer" strategy makes it scalable to large datasets and facilitates the specification of a variety of complex random effects structures. Therefore, we use the general approach of Zhu et al. (2017) to model the RISC dataset, with modifications that are motivated by the application at hand.

In particular, we present a Frequentist implementation of the Bayesian basis modelling approach, which was noted as a possible extension by Zhu et al. (2017) but not pursued. This allows the model to be fitted using existing open-source mixed effects modelling software. However, it does not produce posterior samples for pointwise and simultaneous inference of fixed effects, so for this we adapt existing Frequentist resampling and simulation techniques. The basis modelling approach makes the assumption that each basis coefficient can be modelled separately, though the suitability of this assumption is not always checked in practice. As such, we present an approach to graphically assess the suitability of this assumption for our application by comparing covariance reconstructions to unstructured estimates. Finally, we extend the intraclass correlation coefficient (ICC) for univariate functional data (Di et al. 2009) to the multivariate case, to summarise the degree of intra-subject correlation in our application.

The remainder of the article is structured as follows. In Sect. 2, we describe our proposed methodology and its implementation. Section 3 contains the data analysis and results of our scientific application. We close with a discussion in Sect. 4. A simulation study, additional methodological and application details, and a sensitivity analysis using alternative modelling approaches are contained in an online supplementary material. Supplementary R code is available at https://github.com/FAST-ULxNUIG/RISC1-fda-manuscript-01-code.

2 Methodology

2.1 Model

We denote the bivariate functional observation from the *i*th individual on side *j* as

$$\mathbf{y}_{ij}(t) = (y_{ij}^{(hip)}(t), y_{ij}^{(knee)}(t))^\top,$$

for $i = 1, \ldots, N$ where N is the total number of individuals, $j \in \{\text{left}, \text{right}\}$ and $t \in [0, T]$ which is a normalised time interval. We let $\mathbf{x}_{ij} = (x_{ij1}, \ldots, x_{ijA})^\top$ denote the vector of length A of scalar covariates for subject i on side j. The covariates could be subject specific (e.g., sex, height) or subject-and-side specific (e.g., an indicator for a subject's dominant side).

The bivariate functional mixed effects model is

$$\mathbf{y}_{ij}(t) = \boldsymbol{\beta}_0(t) + \sum_{a=1}^{A} x_{ija}\boldsymbol{\beta}_a(t) + \mathbf{u}_i(t) + \boldsymbol{\varepsilon}_{ij}(t),$$

where the bivariate function $\boldsymbol{\beta}_0(t) = (\beta_0^{(hip)}(t), \beta_0^{(knee)}(t))^\top$ is the intercept function, the bivariate function $\boldsymbol{\beta}_a(t) = (\beta_a^{(hip)}(t), \beta_a^{(knee)}(t))^\top$ is the fixed effect regression coefficient function corresponding to the ath covariate, the bivariate function $\mathbf{u}_i(t) = (u_i^{(hip)}(t), u_i^{(knee)}(t))^\top$ is the functional random intercept for the ith subject and $\boldsymbol{\varepsilon}_{ij}(t) = (\varepsilon_{ij}^{(hip)}(t), \varepsilon_{ij}^{(knee)}(t))^\top$ is the functional random error specific to the ith subject on side j.

The model is the bivariate functional analogue of a scalar linear mixed effects model with a single grouping variable (e.g., Laird and Ware 1982). For $a = 1, \ldots, A$, the fixed-effect function $\boldsymbol{\beta}_a(t)$ captures how the ath scalar covariate influences the "expected level and shape" of the bivariate functional response (Bauer et al. 2018). We assume that the bivariate functional random intercepts $\mathbf{u}_i(t)$, $i = 1, \ldots, N$ are independent realisations of a zero-mean multivariate Gaussian process with a matrix-valued covariance function $\mathbf{Q}$. The bivariate functional random intercepts take into account the grouping structure in the data, i.e., that the left and right side hip and knee angle functions from the same subject are likely to be similar and should share a subject-specific average function. A standard multivariate function-on-scalar regression model without these random effects would ignore this intra-subject correlation, effectively treating an individual's observations from the right and left side as independent. Analogous to random intercepts in scalar linear mixed models, they can be thought of as capturing the correlation between observations from the same subject, or accounting for average differences between subjects. We assume that the bivariate functional random errors $\boldsymbol{\varepsilon}_{ij}(t)$, $i = 1, \ldots, N$, $j \in \{\text{left}, \text{right}\}$ are independent realisations of a zero-mean multivariate Gaussian process with a matrix-valued covariance function $\mathbf{S}$. They are often referred to as "curve-level functional random effects" because they capture correlation within, rather than between, functional observations (Morris 2015).

We stack all functional terms in the model to give

$$\mathbf{Y}(t) = \mathbf{X}\mathbf{B}(t) + \mathbf{Z}\mathbf{U}(t) + \mathbf{E}(t), \tag{1}$$

where the matrix $\mathbf{Y}(t) = (\mathbf{y}_{1,\text{left}}(t) \mid \cdots \mid \mathbf{y}_{N,\text{right}}(t))^\top$ represents the functional observations, the matrix $\mathbf{B}(t) = (\boldsymbol{\beta}_0(t) \mid \cdots \mid \boldsymbol{\beta}_A(t))^\top$ represents the functional fixed effects, the matrix $\mathbf{U}(t) = (\mathbf{u}_1(t) \mid \cdots \mid \mathbf{u}_N(t))^\top$ represents the functional random

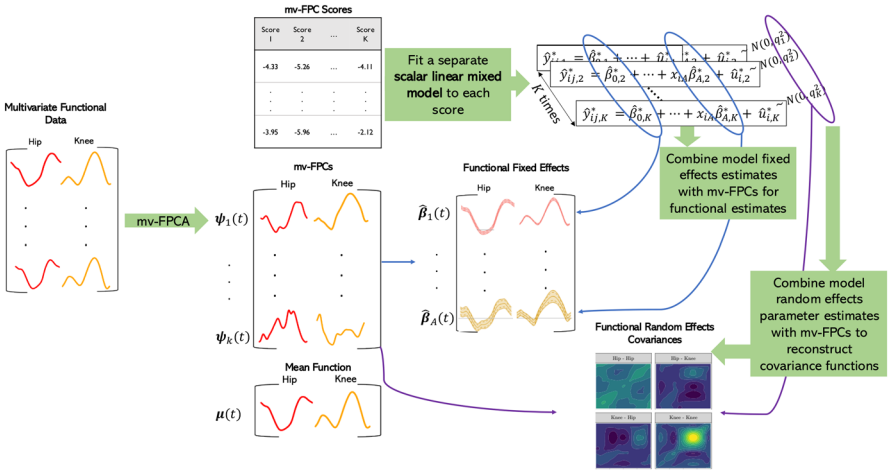

Fig. 2 A flowchart summarising the main sequence of steps in our modelling approach

effects, the matrix $\mathbf{E}(t) = (\boldsymbol{\varepsilon}_{1,\text{left}}(t) \mid \cdots \mid \boldsymbol{\varepsilon}_{N,\text{right}}(t))^\top$ represents the functional random errors and $\mathbf{X}$ and $\mathbf{Z}$ are $2N \times (A+1)$ and $2N \times N$ design matrices for the fixed and random effects, respectively. Written in this way, the model is a bivariate version of the general functional mixed model (Morris and Carroll 2006).

Our approach for fitting the model, described in the remainder of this section, can be summarised as follows. First, the multivariate functional data are expanded on a basis of multivariate functional principal components (mv-FPCs) (Sect. 2.2). Scalar mixed models are fitted to each of the resulting FPC scores independently (Sect. 2.3). Estimates of the model parameters are combined across the bivariate FPCs to give estimates of the functional model terms (Sects. 2.4 and 2.5). The main steps are also summarised graphically in Fig. 2. A short simulation study to assess this approach in realistic data-generating scenarios is contained in Sect. 4 in the Supplement.

2.2 Basis expansion

The first step in the basis modelling strategy is to represent each individual observation using a *basis expansion*, that is

$$\mathbf{y}_{ij}(t) = \sum_{k=1}^{K} y^*_{ijk} \boldsymbol{\psi}_k(t),$$

where $\{y^*_{ijk}\}_{k=1}^K$ are scalar basis coefficients and $\{\boldsymbol{\psi}_k(t) = (\psi_k^{hip}(t), \psi_k^{knee}(t))^\top\}_{k=1}^K$ are bivariate basis functions. As noted by Ramsay and Silverman (2005) and Morris (2015), the type of basis chosen should suit the characteristics of the data at hand. Popular choices for univariate functions are wavelets or FPCs (Morris and Carroll 2006; Aston et al. 2010; Lee et al. 2019).

For multivariate functional data, such as the kinematic data in our application, multivariate FPCs are a natural choice of basis, because they capture common variation across the dimensions of the multivariate functional data. That is, in our application, mv-FPCA produces a set of basis functions that are useful for capturing variation in the hip and knee jointly.

Implicitly, this is the same type of basis that Zhu et al. (2017) used to represent the multivariate functional data in their work. Though they did not explicitly mention mv-FPCA, they expanded the functional data within each dimension on an orthonormal univariate basis and then performed a standard PCA of the combined matrix of basis coefficients, using the resulting PCA scores as new basis coefficients in the modelling. As demonstrated by Happ and Greven (2018), this is one approach to estimating mv-FPCs. However, there are other approaches to estimating the mv-FPCs that are equally valid. In our application, we chose to construct the mv-FPC basis from univariate B-spline (non-orthonormal) basis expansions within each dimension (Happ and Greven 2018), as B-splines are known to be suited to representing smooth functions with local features (Ramsay and Silverman 2005; Morris 2015). The mv-FPCA basis could also be estimated directly from discrete observations of the multivariate functional data Li et al. (2020); Ramsay and Silverman (2005).

The only requirement is that the basis expansion is *near-lossless*, which according to Morris (2017, p. 72), means it is "sufficiently rich such that for all practical purposes it can recapitulate the observed functional data". In general, this property is controlled by K, the number of mv-FPCs retained. Because we calculate the mv-FPCA from univariate expansions, the richness of the basis additionally depends on the number of univariate basis functions used within each dimension. We use a large number of univariate basis functions within each dimension, as the estimated mv-FPCA has been shown to be sensitive to using too few univariate basis functions (Golovkine et al. 2023). When choosing K, the number of mv-FPCs to keep, retaining a larger number will give a closer fit to the observed data, while retaining fewer mv-FPCs ("truncation" or "compression") makes the representation smoother and reduces computation time taken to model the scores. As our data are smooth and we assume that they are measured without error, we choose a large number of mv-FPCs to explain a high variance-explained threshold of 99.99% (Zhu et al. 2017). However, if we wanted to induce regularisation and avoid over-fitting, we could more carefully choose K through cross-validation for near-lossless basis expansions (Zohner 2021, Chap. 2).

Here and in the analysis that follows, we assume that the overall mean has been subtracted from $\mathbf{Y}(t)$ and we add it back when interpreting results. We let $\mathbf{\Psi}(t)$ denote the $K \times 2$ matrix containing the mv-FPCs so that we can write $\mathbf{Y}(t) = \mathbf{Y}^*\mathbf{\Psi}(t)$, where $\mathbf{Y}^*$ is the matrix of basis coefficients (i.e., mv-FPCA scores) which is obtained by projecting the $\mathbf{Y}(t)$ onto the mv-FPCs. The central idea of the basis modelling paradigm is to use the same basis for all terms in the model (1), i.e., $\mathbf{B}(t) = \mathbf{B}^*\mathbf{\Psi}(t)$, $\mathbf{U}(t) = \mathbf{U}^*\mathbf{\Psi}(t)$ and $\mathbf{E}(t) = \mathbf{E}^*\mathbf{\Psi}(t)$ so that the following "basis-space" model can be fitted instead

$$\mathbf{Y}^* = \mathbf{XB}^* + \mathbf{ZU}^* + \mathbf{E}^*,$$

which is obtained by projecting both sides of (1) onto $\mathbf{\Psi}(t)$ (see, e.g., Morris 2017, Sect. 3.1.4). This simplifies the task from fitting a bivariate functional mixed model (the "data-space" model) to fitting a multivariate scalar linear mixed model (the "basis-space" model). Because mv-FPCA scores are (marginally) uncorrelated, we attain further simplification by assuming that the columns of $\mathbf{Y}^*$ are approximately independent and can be modelled separately. This reduces the problem to fitting a series of univariate scalar linear mixed models to the columns of $\mathbf{Y}^*$, reducing computation times and memory requirements.

2.3 Estimation

A Gaussian scalar linear mixed effects model is fitted separately to each FPC score, i.e., each column of $\mathbf{Y}^*$. The model for the kth basis coefficient, $k = 1, \ldots, K$, is

$$y^*_{ijk} = \beta^*_{0k} + \sum_{a=1}^{A} x_{ija}\beta^*_{ak} + u^*_{ik} + \varepsilon^*_{ijk},$$

where $u^*_{ik} \overset{i.i.d.}{\sim} \mathcal{N}(0, q_k)$ and $\varepsilon^*_{ijk} \overset{i.i.d.}{\sim} \mathcal{N}(0, s_k)$. Here, the Gaussian specification for the random effects and random errors follows from the assumption of a Gaussian process for $\mathbf{u}_i(t)$ and $\boldsymbol{\varepsilon}_{ij}(t)$, and hence that their projection onto an orthonormal basis yields Gaussian coefficients (Golovkine et al. 2022). The model can be estimated using either Bayesian or Frequentist methods; Zhu et al. (2017) took a Bayesian approach using a custom MCMC algorithm written in MATLAB and C. We opt for a frequentist approach because it allows a fast and straightforward implementation using standard open-source software; we use the `lmer()` function from the **lme4** (Bates et al. 2015) R (R Core Team 2022) package to fit the models using REML. An introduction to REML estimation of linear mixed models is given by Wood (2017, Sect. 2.4.5) and Pinheiro and Bates (2000, Sect. 2.2.5).

Implicitly, modelling each basis coefficient separately assumes that $\text{Cov}(u^*_{ik}, u^*_{ik'}) = 0$ and $\text{Cov}(\varepsilon^*_{ijk}, \varepsilon^*_{ijk'}) = 0$ for $k \neq k'$. Although the mv-FPCA step produces basis coefficients that are marginally uncorrelated, the projections of the individual processes $\mathbf{u}_i(t)$ and $\boldsymbol{\varepsilon}_{ij}(t)$ onto the mv-FPCA basis are not guaranteed to be uncorrelated across k. However, this assumption is commonly made in basis modelling approaches for functional mixed models because it brings about simplifications in modelling and computation while maintaining a high degree of flexibility (Aston et al. 2010; Lee et al. 2019; Zhu et al. 2017). In Sect. 2.5, we describe a way to graphically assess the extent to which this assumption is reasonable for each process, based on the reconstruction of their respective covariance functions.

2.4 Fixed effects

We do not try to interpret the individual models fitted to the basis coefficients. Instead we combine the estimated parameters across coefficients with the basis functions to

reconstruct the functional model terms. This step is referred to as "transforming the estimates back to the data space". For a given fixed-effect function $\boldsymbol{\beta}_a(t)$, we use the estimates $\widehat{\beta}^*_{ak}$ of β^*_{ak}, $k = 1, \ldots, K$ to construct an estimate

$$\widehat{\boldsymbol{\beta}}_a(t) = \sum_{k=1}^{K} \widehat{\beta}^*_{ak} \boldsymbol{\psi}_k(t) = \boldsymbol{\Psi}(t)^\top \widehat{\boldsymbol{\beta}}^*_a, \quad \text{where } \widehat{\boldsymbol{\beta}}^*_a = (\widehat{\beta}^*_{a1}, \ldots, \widehat{\beta}^*_{aK})^\top.$$

2.4.1 Pointwise confidence intervals

Pointwise confidence intervals for $\boldsymbol{\beta}_a(t)$ can be constructed based on a Gaussian approximation $\widehat{\boldsymbol{\beta}}^*_a \sim \mathcal{N}_K(\boldsymbol{\beta}_a, \widehat{\boldsymbol{\Sigma}}_a)$, where $\widehat{\boldsymbol{\Sigma}}_a = \text{diag}\left\{\widehat{\text{Var}}\left(\widehat{\beta}^*_{a1}\right), \ldots, \widehat{\text{Var}}\left(\widehat{\beta}^*_{aK}\right)\right\}$. This gives the pointwise variance function

$$\widehat{\text{Var}}(\widehat{\boldsymbol{\beta}}_a(t)) \approx \boldsymbol{\Psi}(t)^\top \widehat{\boldsymbol{\Sigma}}_a \boldsymbol{\Psi}(t),$$

so that an approximate pointwise confidence interval can be constructed as $\widehat{\boldsymbol{\beta}}_a(t) \pm q_{1-\alpha/2} \times \widehat{\text{SE}}(\widehat{\boldsymbol{\beta}}_a(t))$ where $\widehat{\text{SE}}(\widehat{\boldsymbol{\beta}}_a(t))$ is the square-root of $\widehat{\text{Var}}(\widehat{\boldsymbol{\beta}}_a(t))$ and $q_{1-\alpha/2}$ is the $(1-\alpha/2)$th quantile of the standard Gaussian distribution. These are Wald intervals because they are based on the Gaussian approximation for each $\widehat{\beta}^*_{ak}$ and are only asymptotically valid because the estimate $\widehat{\text{Var}}(\widehat{\beta}^*_{ak})$ is used in place of the true $\text{Var}(\widehat{\beta}^*_{ak})$ (Kenward and Roger 1997). Despite this, Wald intervals are quick and straightforward to compute and are returned by default by standard mixed model software. For our application with a large number of study participants the approximation should be reasonable so they are a convenient tool. In Sect. 2.4.2, we describe a more computationally intensive bootstrap technique for constructing simultaneous confidence bands which can also be used to construct pointwise intervals.

2.4.2 Simultaneous confidence bands

Pointwise confidence intervals for bivariate functional parameters only provide coverage within a given dimension $p \in \{\text{hip}, \text{knee}\}$ at a specific point $t \in [0, T]$. They will not, in general, provide nominal coverage for the entire function $\boldsymbol{\beta}_a(t)$ because of the multiple-testing problem (Degras 2017). We define a simultaneous confidence band as the band $[\boldsymbol{\beta}_{a,L}(t),\ \boldsymbol{\beta}_{a,U}(t)]$ providing simultaneous coverage

$$P(\beta^{(p)}_a(t) \in [\beta^{(p)}_{a,L}(t),\ \beta^{(p)}_{a,U}(t)], \forall\, p \in \{\text{hip}, \text{knee}\} \text{ and } t \in [0, T]) \approx 1 - \alpha.$$

The band can be thought of as providing an adjustment for multiple testing along the whole domain $[0, T]$ and across the hip and knee dimensions.

Resampling or simulation techniques are typically used to build simultaneous confidence bands. A sketch of the general procedure first introduced by Ruppert et al. (2003, Sec. 6.5) for scatterplot smoothing, which has been subsequently been shown to work well for univariate functional data (Crainiceanu et al. 2012; Park et al. 2018; Cui et al. 2022), is given in Algorithm 1. The algorithm admits a number of ways to construct the bands in our application, differing in how the samples $\widehat{\boldsymbol{\beta}}_{a(1)}(t), \ldots \widehat{\boldsymbol{\beta}}_{a(R)}(t)$ and the estimate $\widehat{\mathrm{SE}}(\widehat{\boldsymbol{\beta}}_a(t))$ are obtained (Step 1). The Wald approximation in Sect. 2.4.1 can be used to simulate samples from $\mathcal{N}_K(\widehat{\boldsymbol{\beta}}_a, \widehat{\boldsymbol{\Sigma}}_a)$. Alternatively, parametric or non-parametric bootstrap techniques can be used to obtain the samples and estimate $\widehat{\boldsymbol{\Sigma}}_a$. We opt for the non-parametric bootstrap, where bootstrap samples are created by resampling subject indices with replacement, hence called the "bootstrap of subjects" (Crainiceanu et al. 2012; Park et al. 2018; Cui et al. 2022). Each time a subject appears in a bootstrap sample, they are assigned a new pseudo-ID which is used in model estimation. We use the bootstrap to estimate $\boldsymbol{\Sigma}_a$ and then draw samples from $\mathcal{N}_K(\widehat{\boldsymbol{\beta}}_a, \widehat{\boldsymbol{\Sigma}}_a)$ for Step 1 of Algorithm 1, however the bootstrap samples could also be used directly (Crainiceanu et al. 2012).

Algorithm 1 Level α simultaneous confidence bands for $\boldsymbol{\beta}_a(t)$ (Crainiceanu et al. 2012).

Data: $\widehat{\boldsymbol{\beta}}_a(t)$, $\widehat{\mathrm{SE}}(\widehat{\boldsymbol{\beta}}_a(t))$.
Result: Simultaneous confidence bands of
$\{\boldsymbol{\beta}_a(t), p \in \{\text{hip, knee}\} \text{ and } t \in [0, T]\}$.
1. Obtain samples $\widehat{\boldsymbol{\beta}}_{a(1)}(t), \ldots \widehat{\boldsymbol{\beta}}_{a(R)}(t)$ by simulation or bootstrap;
for $r = 1, \ldots, R$ **do**
  2. Calculate $z_r = \max_{t,p}\{|\widehat{\boldsymbol{\beta}}_a(t) - \widehat{\boldsymbol{\beta}}_{a(r)}(t)|/\widehat{\mathrm{SE}}(\widehat{\boldsymbol{\beta}}_a(t))\}$;
end
3. Compute $z_{(1-\alpha)}$, the $(1-\alpha)$th empirical quantile of $\{z_1, \ldots, z_R\}$;
4. The simultaneous confidence band is calculated as

$$\widehat{\boldsymbol{\beta}}_a(t) \pm z_{(1-\alpha)}\widehat{\mathrm{SE}}(\widehat{\boldsymbol{\beta}}_a(t)).$$

2.5 Covariance reconstruction

In scalar linear mixed effects models, we are not concerned with estimating the random effects themselves; instead we try to estimate the parameters that describe the random effects' distributions, i.e., the variance and covariance parameters (Faraway 2016, p. 195). Analogously in bivariate functional mixed effects models, we are concerned with

estimation of the auto- and cross-covariance functions describing the bivariate functional random effects.

The bivariate functional random intercepts are given by

$$\mathbf{u}_i(t) = \sum_{k=1}^{K} u_{ik}^* \boldsymbol{\psi}_k(t) = \boldsymbol{\Psi}(t)^\top \mathbf{u}_i^*,$$

and due to the independence assumption for the basis coefficients, we have

$$\text{Cov}\left(\mathbf{u}_i^*\right) = \mathbf{Q}^* = \text{diag}\{q_1, \ldots, q_K\},$$

where $q_1, \ldots, q_K$ are random-intercept variances from the scalar mixed models (Sect. 2.3). Therefore, the matrix-valued covariance function for the bivariate functional random intercepts is given by

$$\mathbf{Q}(t, t') = \text{Cov}(\mathbf{u}_i(t), \mathbf{u}_i(t')) = \boldsymbol{\Psi}(t)^\top \mathbf{Q}^* \boldsymbol{\Psi}(t'), \quad t, t' \in [0, T].$$

Similarly the matrix-valued covariance function for the bivariate functional random error is

$$\mathbf{S}(t, t') = \text{Cov}(\boldsymbol{\varepsilon}_{ij}(t), \boldsymbol{\varepsilon}_{ij}(t')) = \boldsymbol{\Psi}(t)^\top \mathbf{S}^* \boldsymbol{\Psi}(t'), \quad t, t' \in [0, T],$$

where $\mathbf{S}^* = \text{diag}\{s_1, \ldots, s_K\}$. In practice, we replace q_k and s_k by their estimates $\widehat{q}_k$ and $\widehat{s}_k$ to obtain the reconstructions $\widehat{\mathbf{Q}}$ of $\mathbf{Q}$ and $\widehat{\mathbf{S}}$ of $\mathbf{S}$.

As mentioned in Sect. 2.3, the independence assumption for the basis coefficients restricts $\mathbf{Q}^*$ and $\mathbf{S}^*$ to be diagonal, limiting the types of covariance structures that can be estimated. Lee et al. (2019) recommend checking this assumption graphically by plotting the reconstructed covariance functions. For functions on large and possibly high-dimensional grids (e.g., images), it has typically not been feasible to compute unrestricted covariance estimates to compare the model reconstructions with. In this work, we obtain fully unstructured estimates of the covariance functions by extending the multilevel FPCA method of Di et al. (2009) to multivariate functional data. By comparing the model and unstructured estimates graphically, we can assess whether the diagonal assumptions for $\mathbf{Q}^*$ and $\mathbf{S}^*$ are reasonable. Full details on the calculation of the unstructured estimates are provided in Sect. 2 in the Supplement.

2.6 Functional intraclass correlation coefficient

Random-intercept scalar mixed models allow a partitioning of variability into between-subjects and within-subjects elements through the intraclass correlation coefficient (ICC) (Faraway 2016, Sect. 8.1). Di et al. (2009, Sect. 2.2) extended the ICC to univariate functional data by integrating each term over the functional domain. We further extend it to multivariate functional data by integrating over the functional domain and summing over the dimensions. The multivariate functional ICC for our model is

$$\text{ICC} = \frac{\sum_{k=1}^{K} q_k}{\sum_{k=1}^{K} q_k + \sum_{k=1}^{K} s_k}.$$

In our application, it can be interpreted as the proportion of variability in the hip and knee angles (after accounting for fixed effects) attributable to differences between subjects. The remainder (1 − ICC) represents the proportion attributable to differences within subjects between the left and right sides (asymmetry). Further details on the ICC are provided in Sect. 3 in the Supplement.

3 Data analysis and results

3.1 Data extraction, segmentation and alignment

This section summarises the data collection, extraction and preparation for analysis. As per the Vicon Plug in Gait model (Vicon Motion Systems, Oxford, UK), 28 reflective markers (14 mm in diameter) were placed at bony landmarks on the lower limbs, pelvis and trunk with an additional two markers placed on the anterior aspect of the mid tibia and mid thigh bilaterally. After a dynamic warm-up including treadmill running (FlowFitness, Runner-DTM2500i, Netherlands) for 6 min at a speed of 9 km h^{-1}, participants completed a three-minute run at a self-selected pace that best represented their typical training pace. During the first minute of this three-minute run, kinematic data were collected using a 17-camera, three-dimensional motion analysis system (Vantage, Vicon, Oxford, UK) recording at 200Hz. The marker trajectories were then filtered using a fourth-order zero-lag Butterworth filter at 15Hz, chosen by residual analysis (Winter 1979). Functional joints and minimisation of soft tissue were calculated using the "OSSCA" method in NEXUS 2 (Taylor et al. 2010). Sagittal plane hip and knee angles were then extracted bilaterally.

The extracted data were segmented into individual strides at the initial contact of the foot with the ground, which was identified as the first occurrence of two events: 1) the first negative vertical acceleration of the toe maker, and 2) the peak vertical acceleration of the heel marker. Both events were identified within a search window defined between the local maxima of the toe marker anterior position and the subsequent local minima of the ankle marker vertical position. For each stride, the time-argument values were then linearly re-scaled so that all curves shared the normalised domain [0, 100], where 0 represents the start of a stride and 100(%) represents the end (i.e., linear time/ length normalisation, Helwig et al. 2011). When discretisation of the functions was required, e.g., for plotting or computing the simultaneous bands, a grid of 101 points $t = 0, 1, \ldots, 100$ was used. Landmark registration (Kneip and Gasser 1992) was performed to further reduce timing variation in the functions. A single landmark was chosen to align the functional data from each stride—the peak of the knee flexion angle. This landmark was chosen because it is clear and well-defined for every stride and easy to identify using a simple grid search. The hip and knee angles were registered simultaneously to this landmark to preserve the temporal correlation between them.

3.2 mv-FPCA

As described in Sect. 3.1, the raw marker trajectories were filtered to remove observational error. Therefore, no additional smoothing was performed to avoid oversmoothing and dampening features in the data. Instead, the first-stage basis-function expansion interpolated, rather than smoothed, the data and reduced its dimension (i.e., reduced a large number of observation points, differing between curves, to a smaller number of common basis coefficients).

First, a B-spline basis was chosen to represent the univariate functional data in each dimension because it is a flexible basis and is well suited to smooth functions, such as the kinematic data at hand (Ramsay and Silverman 2005; Morris 2015). We found that $K_{hip} = K_{knee} = 80$ B-spline basis functions were sufficient to approximate the functional data from each stride almost perfectly. The basis coefficients were computed by ordinary least squares because no smoothing was required. Given the basis representation of the functional data for each individual stride, the reduced dataset of left and right side averages used in the analysis (Fig. 1) was obtained by averaging the basis coefficients of all strides for a given subject on a given side of the body. Computed from the univariate B-spline expansions, the bivariate FPCA yielded $\widetilde{K} = 38$ bivariate FPCs, satisfying the 99.99% variance explained threshold. As expected, the majority of the variance was explained by the leading FPCs, e.g., 95% of the variance was explained by the first seven FPCs, and 99% by the first 13. Additional information on the basis transformation is provided in Section 5.1 in the Supplement.

3.3 Fixed effects

The fitted model was

$$\mathbf{y}_{ij}(t) = \beta_0(t) + \underbrace{\sum_{a=1}^{3} x_{ia}\beta_a(t)}_{\text{Injury Status}} + \text{speed}_i \times \beta_4(t) + \text{sex}_i \times \beta_5(t) + \text{age}_i \times \beta_6(t) + \text{weight}_i \times \beta_7(t) + \text{height}_i \times \beta_8(t) + \mathbf{u}_i(t) + \varepsilon_{ij}(t),$$

where x_{i1}, x_{i2} and x_{i3} are dummy-coded variables representing the "Injured more than 2 years ago", "Injured 1-2 years ago" and "Injured less than 1 year ago" categories of the retrospective injury status variable and the reference category is "Never injured", speed_i is the self-selected running speed of subject i in km h^{-1}, sex_i is a dummy-coded variable for sex of subject i (0 = male, 1 = female), age_i is the age of subject i in years, weight_i is the weight of subject i in kilograms and height_i is the height of subject i in centimetres. All numeric variables were centred to make the intercept function more interpretable. The regression coefficient functions for numerical and dummy-coded variables can be interpreted analogously to multiple

linear regression; e.g., $\beta_4^{(hip)}(t)$ represents the expected change in the hip angle at t for a 1-km h^{-1} increase in speed with all other variables held constant, and $\beta_5^{(knee)}(t)$ represents the expected difference in the knee angle at t between females and males with all other variables held constant.

Figure 3 shows the estimated regression coefficient functions. The solid lines represent the point estimates, the shaded ribbons represent the 95% simultaneous bands and the dashed lines represent 95% pointwise confidence intervals. Results obtained via the Wald and bootstrap approaches were practically indistinguishable so only the bootstrap intervals are shown. The simultaneous bands are about 1.5 times as wide as the pointwise intervals. The confidence bands for the retrospective injury status regression coefficient functions contain zero (solid grey horizontal line) for all t, meaning that there is no evidence of a difference between any of the categories and the reference category of “Never injured". Similarly, there is limited evidence of an age, height, weight or sex effect; although the simultaneous bands do not contain zero at certain points, the magnitude of each effect is small. However, self-selected speed has a strong effect in both the hip and knee dimensions—the coefficient function has a distinct shape and the confidence band only contains zero when the function is changing from positive to negative. In the Supplement (Sec. 5.2), we provide a more intuitive visualisation to interpret this effect, where we predict the hip and knee angles from our model at different speeds, and present the predicted joint angle curves relative to time and relative to one another on an angle-angle diagram.

3.4 Random effects

Figure 4 displays filled-contour plots of the estimated multivariate covariance functions. The random-effects covariance function $\mathbf{Q}$ (top panel) is reconstructed by the model almost perfectly—the model estimates (left panel) appear very similar to the unstructured estimates (right panel). The random-error covariance function $\mathbf{S}$ (bottom panel) is also well estimated, the model captures the general shape of the function. However, there are certain parts of $\mathbf{S}$ which the model cannot reconstruct. For example, in the knee-knee component of $\mathbf{S}$ in the region $t, t' \in [75, 100]$ there is a discrepancy between the model and unstructured estimates. We show in Section 5.3 in the Supplement that this discrepancy is due to the diagonal assumption for $\mathbf{S}^*$ and can be resolved by allowing a small number of non-zero off-diagonal correlations in $\mathbf{S}^*$. Overall, however, we can conclude that the covariance functions are reconstructed well.

The functional ICC was estimated at 0.78 (bootstrap 95% CI = [0.76, 0.81]), indicating that approximately 78% of the variability in the average hip and knee angle functions (after accounting for the fixed effects) can be explained by subject-level differences, and 22% of variability is due to differences within subjects between sides. This high degree of “clustering" highlights the presence of idiosyncratic running patterns that are consistent across both sides of the body.

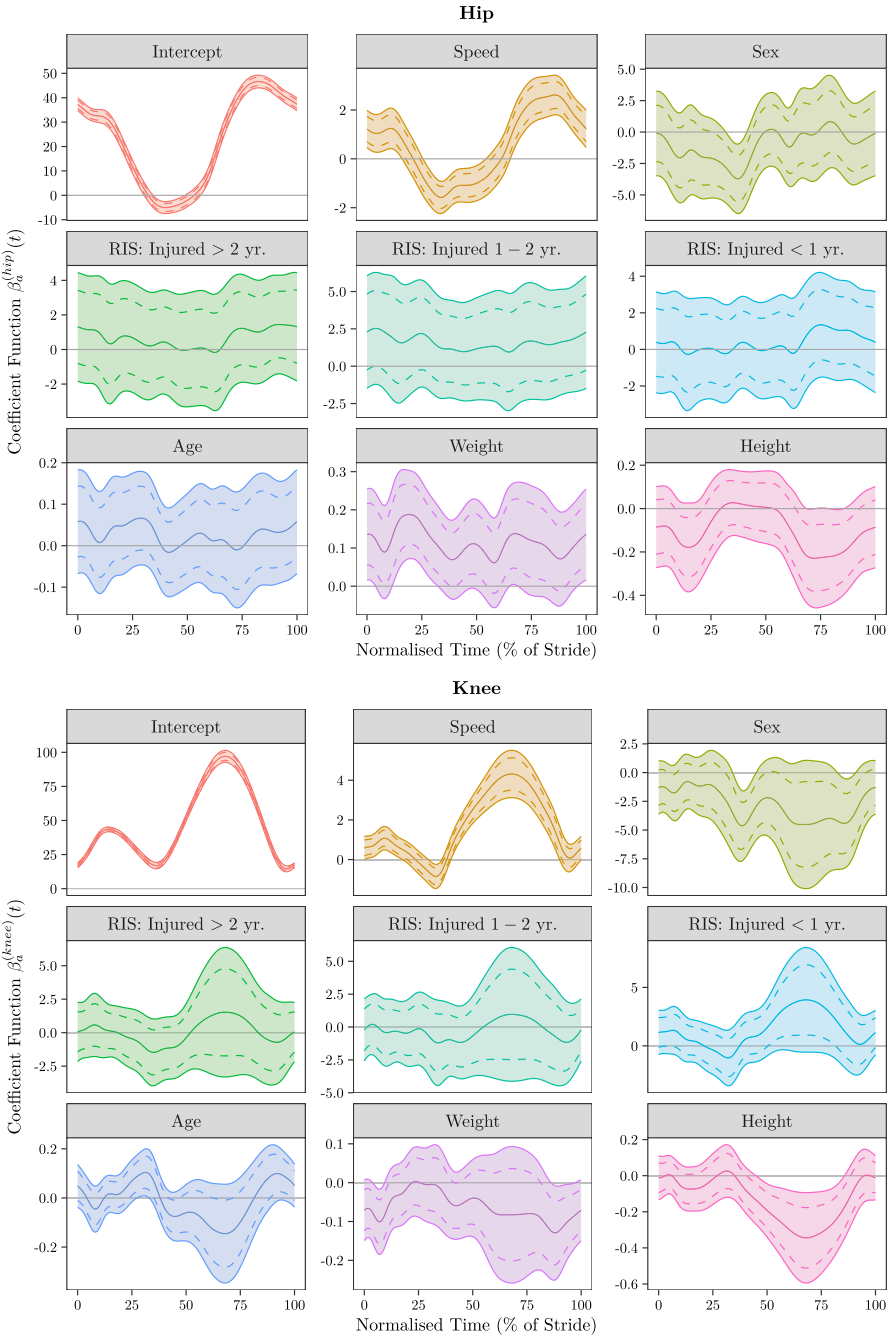

Fig. 3 The estimated fixed-effect hip (top panel) and knee (bottom panel) regression coefficient functions. The solid line represents the point estimate function. The shaded ribbons represent 95% simultaneous confidence bands obtained via bootstrap. The dashed lines represent 95% pointwise confidence intervals obtained via bootstrap. Note: Wald confidence intervals are omitted because they are almost identical to those obtained with the bootstrap approach

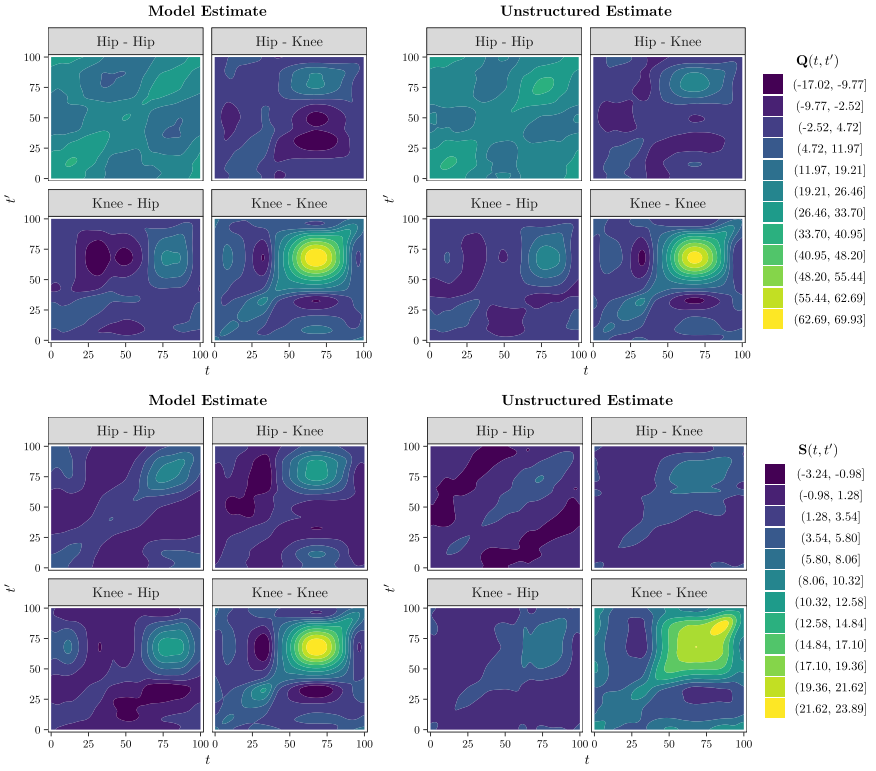

Fig. 4 Filled-contour plots of the multivariate covariance functions. Top panel: The model (left) and unstructured (right) estimates of the multivariate functional random effects covariance function $\mathbf{Q}(t, t')$. Bottom panel: The model (left) and unstructured (right) estimates of the multivariate functional random error covariance function $\mathbf{S}(t, t')$

4 Discussion

This article has presented a multivariate functional mixed model for kinematic data from recreational runners collected during a treadmill run. Using an existing basis modelling approach (Zhu et al. 2017), we project the multivariate functional data onto a mv-FPCA basis to reduce it to a to a set of uncorrelated scores and fit a series of scalar linear mixed models to the scores. We have provided a frequentist implementation of the model which means it can be fitted using existing open-source software and adapted bootstrap and simulation techniques for inference on the fixed-effect function estimates. We introduced reconstructions and comparisons of the multivariate covariance structures to graphically assess the model, which demonstrated that the assumptions being made on these structures were appropriate for our data application. We also extended the univariate functional ICC to summarise the degree of intra-subject correlation in our application, showing strong correlations

in runner's bilateral kinematics (or equivalently, high inter-subject variability/ idiosyncrasy).

From a scientific perspective, we did not detect evidence of a statistically significant effect of retrospective injury status on the kinematic data after accounting for the other covariates. It is difficult to directly compare this result with existing literature due to differences in study designs, and also because previous analyses have focused on discrete kinematic variables whereas we modelled kinematics over the course of the full running stride. For example, the findings are consistent with the work of Messier et al. (2018) who, in a large *prospective* study of runners who were all asymptomatic at baseline, found that discrete measures of knee flexion were not significantly different between those who did and not become injured. On the other hand, Bramah et al. (2018) found a significant difference in knee flexion at the start of the stride between injured and uninjured runners in a retrospective study, were the injured runners *were* symptomatic at baseline. Although our model did aim to comprehensively characterise effects of injury status on both hip and knee kinematics across the whole running stride, we cannot rule out the possibility that our injury groupings were too broad (both in terms of injury types/ location and the times since the runners were symptomatic), that our sample size was insufficient to detect a difference using these groupings, or that other joints or planes of motion are more important for RRIs.

In contrast, we did find a strong, statistically significant effect of running speed on hip and knee kinematics—participants who run faster tend to so by producing greater hip and knee flexion at various stages throughout the movement. The functional effects that we have characterised using FDA are understood qualitatively in the biomechanical literature—Grimshaw et al. (2007, p. 256) explain that "As speed increases, the flexion of hip and knee joints during the swing phase increases, this serves to reduce the moment of inertia of the limb, thus allowing for a faster swing. There may also be a slight increase in the degree of knee flexion at impact". Additionally, almost identical qualitative effects of running speed were found by Orendurff et al. (2018) in experiments where the individuals ran on a treadmill at multiple different speeds (i.e., speed was a within-subject variable). However, their statistical modelling approach was limited to simply plotting the group average curves at the different speeds without any inference, and then reducing the curve data to discrete variables (e.g., peak knee flexion) and performing a repeated measures ANOVA on them, treating speed as a categorical covariate. Although this might be sufficient in some applications, including speed as a continuous covariate in our multivariate functional mixed effects model allows us to obtain estimates that appropriately characterise the effect of speed across the whole running stride, obtain simultaneous inference on this effect that is valid across the whole running stride and across both the hip and knee joints, and to make predictions of full curves at different running speeds (Supp. Sect. 5.2), while still accounting for repeated measures. An added benefit of modelling the hip and knee jointly rather than fitting separate univariate models is that it leads to intuitive visualisations of combined effects on hip-knee kinematics using angle-angle diagrams (Supp. Sect. 5.2), which are a practically useful tool for biomechanics researchers and practitioners conducting coordination research (Lamb and Bartlett 2017).

Some limitations and extensions of this work are as follows. The kinematic data had already undergone filtering in the extraction step, as is typical for human movement data collected using motion capture systems, so further smoothing was not applied. However, in other scenarios where the data are less smooth, it may be desirable to regularise the estimated fixed effects functions. This could be achieved by pre-smoothing the individual functional observations in the first-stage basis transformation or retaining fewer FPCs in the second stage. However, in certain situations, heavily pre-smoothing individual observations may neglect uncertainty in their estimates in downstream analysis (Bauer et al. 2018). The fixed effects estimates could also be post-smoothed by evaluating them on a grid and employing any scatterplot smoother (Fan and Zhang 2000; Cui et al. 2022). Finally, variable selection could be used in the scalar linear mixed effects models, which would lead to a sparse representation of the fixed-effect functions, i.e., each fixed-effect function would be represented by a small number of FPCs (Morris and Carroll 2006; Aston et al. 2010).

In the second simulation scenario, and our scientific application, the random-effect and random-error covariance functions are reconstructed with error because the diagonal assumption for $\mathbf{Q}^*$ and $\mathbf{S}^*$ is too restrictive to fully capture the covariance structures. However, the approximation still works well to provide approximate fixed effects inference and summaries of the variance structure, i.e., the ICC.

If better estimates of the covariance functions were required, a modification could be made to the current approach to allow a small number of off-diagonal elements in $\mathbf{Q}^*$ and $\mathbf{S}^*$ to be non-zero. We show in Section 5.3 in the Supplement that unrestricted versions of $\mathbf{Q}^*$ and $\mathbf{S}^*$ can be estimated using the algorithm of Fieuws and Verbeke (2006), and variants of the graphical LASSO (Friedman et al. 2008) used to select which off-diagonal elements to retain. The final model could then be fitted with certain FPCs modelled in pairs or small groups, rather than completely independently. We worked with linear-time normalised and landmark-registered curves, but did not include the respective parameters of these transformations (i.e., curve lengths and landmark times) in subsequent analysis. It is likely that these parameters also depend on the covariates used in our model. A unified modelling of amplitude and phase (see, e.g., Hadjipantelis et al. 2015) could be achieved by modelling the phase variation parameters along with the mv-FPC scores, likely allowing for correlation among them. We leave further investigation of this approach to future work.

Two main extensions of the model and application will be pursued. Our first goal is to extend the model to include all strides rather than an average for each side. As the strides admit a time ordering, *longitudinal functional data analysis* methods will be required—simply adding another level to the current model and ignoring the ordering of the strides may not be sufficient. There are a number of papers on univariate longitudinal functional data (Greven et al. 2010; Park and Staicu 2015; Lee et al. 2019), however we are developing bespoke methodology to handle the multivariate three-level case. The second extension is to include kinematic data from other joints, such as the ankle or pelvis, or from the other two planes of motion (i.e., frontal and transverse) in the model. This extension is more straightforward methodologically, but will be more computationally demanding and may provide interesting scientific results.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s00180-024-01591-1.

Acknowledgements The authors thank the editors and reviewers for their detailed comments, which greatly helped to improve the manuscript.

Funding Science Foundation Ireland (SFI) grant numbers 18/CRT/6049 (EG), 19/FFP/7002 (SG, AJS and NB), and SFI/12/RC/2289_P2 (RISC running dataset), co-funded by the European Regional Development Fund.

Declarations

Conflict of interest No relevant financial or non-financial interests to disclose.

References

Aguilera-Morillo MC, Aguilera AM (2020) Multi-class classification of biomechanical data: a functional LDA approach based on multi-class penalized functional PLS. Stat Model 20(6):592–616. https://doi.org/10.1177/1471082X19871157

Aston JAD, Chiou J-M, Evans JP (2010) Linguistic pitch analysis using functional principal component mixed effect models. J R Stat Soc: Ser C: Appl Stat 59(2):297–317. https://doi.org/10.1111/j.1467-9876.2009.00689.x

Bates D, Mächler M, Bolker B, Walker S (2015) Fitting linear mixed-effects models using lme4. J Stat Softw 67(1):1–48. https://doi.org/10.18637/jss.v067.i01

Bauer A, Scheipl F, Küchenhoff H, Gabriel A-A (2018) An introduction to semiparametric function-on-scalar regression. Stat Model 18(3–4):346–364. https://doi.org/10.1177/1471082X17748034

Becker J, James S, Wayner R, Osternig L, Chou L-S (2017) Biomechanical factors associated with achilles tendinopathy and medial tibial stress syndrome in runners. Am J Sports Med 45(11):2614–2621. https://doi.org/10.1177/0363546517708193

Bramah C, Preece SJ, Gill N, Herrington L (2018) Is there a pathological gait associated with common soft tissue running injuries? Am J Sports Med 46(12):3023–3031. https://doi.org/10.1177/0363546518793657

Burke A, Dillon S, O'Connor S, Whyte EF, Gore S, Moran KA (2022) Comparison of impact accelerations between injury-resistant and recently injured recreational runners. PLoS ONE 17(9):0273716. https://doi.org/10.1371/journal.pone.0273716

Ceyssens L, Vanelderen R, Barton C, Malliaras P, Dingenen B (2019) Biomechanical risk factors associated with running-related injuries: a systematic review. Sports Med 49(7):1095–1115. https://doi.org/10.1007/s40279-019-01110-z

Coffey N, Harrison AJ, Donoghue OA, Hayes K (2011) Common functional principal components analysis: a new approach to analyzing human movement data. Hum Mov Sci 30(6):1144–1166. https://doi.org/10.1016/j.humov.2010.11.005

Crainiceanu CM, Staicu A-M, Ray S, Punjabi N (2012) Bootstrap-based inference on the difference in the means of two correlated functional processes. Stat Med 31(26):3223–3240. https://doi.org/10.1002/sim.5439

Cui E, Leroux A, Smirnova E, Crainiceanu CM (2022) Fast univariate inference for longitudinal functional models. J Comput Graph Stat 31(1):219–230. https://doi.org/10.1080/10618600.2021.1950006

Degras D (2017) Simultaneous confidence bands for the mean of functional data. Wiley Interdiscip Rev Comput Stat 9(3):1397. https://doi.org/10.1002/wics.1397

Di C-Z, Crainiceanu CM, Caffo BS, Punjabi NM (2009) Multilevel functional principal component analysis. Ann Appl Stat 3(1):458–488. https://doi.org/10.1214/08-AOAS206SUPP

Dillon S, Burke A, Whyte EF, O'Connor S, Gore S, Moran KA (2021) Do injury-resistant runners have distinct differences in clinical measures compared with recently injured runners? Med Sci Sports Exerc 53(9):1807–1817. https://doi.org/10.1249/MSS.0000000000002649

Diquigiovanni J, Fontana M, Vantini S (2022) Conformal prediction bands for multivariate functional data. J Multivar Anal 189:104879. https://doi.org/10.1016/j.jmva.2021.104879

Fan J, Zhang J-T (2000) Two-step estimation of functional linear models with applications to longitudinal data. J R Stat Soc Series B Stat Methodol 62(2):303–322. https://doi.org/10.1111/1467-9868.00233

Faraway JJ (1997) Regression analysis for a functional response. Technometrics 39(3):254–261. https://doi.org/10.2307/1271130

Faraway JJ (2016) Extending the Linear Model with R : Generalized Linear, Mixed Effects and Nonparametric Regression Models, Second Edition. Chapman & Hall/CRC, New York. https://doi.org/10.1201/9781315382722

Ferber R, Osis ST, Hicks JL, Delp SL (2016) Gait biomechanics in the era of data science. J Biomech 49(16):3759–3761. https://doi.org/10.1016/j.jbiomech.2016.10.033

Fieuws S, Verbeke G (2006) Pairwise fitting of mixed models for the joint modeling of multivariate longitudinal profiles. Biometrics 62(2):424–431. https://doi.org/10.1111/j.1541-0420.2006.00507.x

Friedman J, Hastie T, Tibshirani R (2008) Sparse inverse covariance estimation with the graphical lasso. Biostatistics 9(3):432–441. https://doi.org/10.1093/biostatistics/kxm045

Glazier PS (2021) Beyond animated skeletons: how can biomechanical feedback be used to enhance sports performance? J Biomech 129:110686. https://doi.org/10.1016/j.jbiomech.2021.110686

Goldsmith J, Kitago T (2016) Assessing systematic effects of stroke on motorcontrol by using hierarchical function-on-scalar regression. J R Stat Soc: Ser C: Appl Stat 65(2):215–236. https://doi.org/10.1111/rssc.12115

Golovkine S, Gunning E, Simpkin AJ, Bargary N (2023) On the use of the Gram matrix for multivariate functional principal components analysis. arXiv:2306.12949 [stat]. https://doi.org/10.48550/arXiv.2306.12949

Golovkine S, Klutchnikoff N, Patilea V (2022) Clustering multivariate functional data using unsupervised binary trees. Comput Stat Data Anal 168:107376. https://doi.org/10.1016/j.csda.2021.107376

Gòrecki T, Krzyśko M, Waszak L, Wolyński W (2018) Selected statistical methods of data analysis for multivariate functional data. Stat Pap 59(1):153–182. https://doi.org/10.1007/s00362-016-0757-8

Greven S, Crainiceanu CM, Caffo B, Reich D (2010) Longitudinal functional principal component analysis. Electron J Stat 4:1022–1054. https://doi.org/10.1214/10-EJS575

Grimshaw P, Fowler N, Lees A, Burden A (2007) BIOS instant notes in sport and exercise biomechanics. Routledge, Oxon. https://doi.org/10.4324/9780203488300

Guo W (2002) Functional mixed effects models. Biometrics 58(1):121–128. https://doi.org/10.1111/j.0006-341x.2002.00121.x

Hadjipantelis PZ, Aston JAD, Müller HG, Evans JP (2015) Unifying amplitude and phase analysis: a compositional data approach to functional multivariate mixed-effects modeling of mandarin Chinese. J Am Stat Assoc 110(510):545–559. https://doi.org/10.1080/01621459.2015.1006729

Happ C, Greven S (2018) Multivariate functional principal component analysis for data observed on different (Dimensional) domains. J Am Stat Assoc 113(522):649–659. https://doi.org/10.1080/01621459.2016.1273115

Helwig NE, Hong S, Hsiao-Wecksler ET, Polk JD (2011) Methods to temporally align gait cycle data. J Biomech 44(3):561–566. https://doi.org/10.1016/j.jbiomech.2010.09.015

Hespanhol Junior LC, Mechelen W, Verhagen E (2017) Health and economic burden of running-related injuries in dutch trailrunners: a prospective cohort study. Sports Med 47(2):367–377. https://doi.org/10.1007/s40279-016-0551-8

Hébert-Losier K, Pini A, Vantini S, Strandberg J, Abramowicz K, Schelin L, Häger CK (2015) One-leg hop kinematics 20 years following anterior cruciate ligament rupture: data revisited using functional data analysis. Clin Biomech (Bristol) 30(10):1153–1161. https://doi.org/10.1016/j.clinbiomech.2015.08.010

Jiang J, Lin H, Zhong Q, Li Y (2022) Analysis of multivariate non-gaussian functional data: a semiparametric latent process approach. J Multivar Anal 189:104888. https://doi.org/10.1016/j.jmva.2021.104888

Kenward MG, Roger JH (1997) Small sample inference for fixed effects from restricted maximum likelihood. Biometrics 53(3):983–997. https://doi.org/10.2307/2533558

Kneip A, Gasser T (1992) Statistical tools to analyze data representing a sample of curves. Ann Stat 20(3):1266–1305. https://doi.org/10.1214/aos/1176348769

Laird NM, Ware JH (1982) Random-effects models for longitudinal data. Biometrics 38(4):963–974. https://doi.org/10.2307/2529876

Lamb PF, Bartlett RM (2017) Assessing movement coordination. In: Biomechanical Evaluation of Movement in Sport and Exercise, 2nd edn. Routledge, London. https://doi.org/10.4324/9780203095546

Lee W, Miranda MF, Rausch P, Baladandayuthapani V, Fazio M, Downs JC, Morris JS (2019) Bayesian semiparametric functional mixed models for serially correlated functional data, with application to glaucoma data. J Am Stat Assoc 114(526):495–513. https://doi.org/10.1080/01621459.2018.1476242

Li C, Xiao L, Luo S (2020) Fast covariance estimation for multivariate sparse functional data. Stat 9(1):245. https://doi.org/10.1002/sta4.245

Li R, Xiao L (2023) Latent factor model for multivariate functional data. Biometrics 79(4):3307–3318. https://doi.org/10.1111/biom.13924

Liebl D, Willwacher S, Hamill J, Brüggemann G-P (2014) Ankle plantarflexion strength in rearfoot and forefoot runners: a novel clusteranalytic approach. Hum Mov Sci 35:104–120. https://doi.org/10.1016/j.humov.2014.03.008

Liu X, Ma S, Chen K (2022) Multivariate functional regression via nested reduced-rank regularization. J Comput Graph Stat 31(1):231–240. https://doi.org/10.1080/10618600.2021.1960850

Liu Z, Guo W (2012) Functional mixed effects models. Wiley Interdiscip Rev Comput Stat 4(6):527–534. https://doi.org/10.1002/wics.1226

Mann R, Malisoux L, Nührenbörger C, Urhausen A, Meijer K, Theisen D (2015) Association of previous injury and speed with running style and stride-to-stride fluctuations. Scand J Med Sci Sports 25(6):638–645. https://doi.org/10.1111/sms.12397

Matabuena M, Karas M, Riazati S, Caplan N, Hayes PR (2023) Estimating knee movement patterns of recreational runners across training sessions using multilevel functional regression models. Am Stat 77(2):169–181. https://doi.org/10.1080/00031305.2022.2105950

Messier SP, Martin DF, Mihalko SL, Ip E, DeVita P, Cannon DW, Love M, Beringer D, Saldana S, Fellin RE, Seay JF (2018) A 2-year prospective cohort study of overuse running injuries: the runners and injury longitudinal study (trails). Am J Sports Med 46(9):2211–2221. https://doi.org/10.1177/0363546518773755

Morris JS (2015) Functional regression. Annu Rev Stat Appl 2:321–359. https://doi.org/10.1146/annurev-statistics-010814-020413

Morris JS (2017) Comparison and contrast of two general functional regression modelling frameworks. Stat Modell 17(1–2):59–85. https://doi.org/10.1177/1471082X16681875

Morris JS, Carroll RJ (2006) Wavelet-based functional mixed models. J R Stat Soc Series B Stat Methodol 68(2):179–199. https://doi.org/10.1111/j.1467-9868.2006.00539.x

Orendurff MS, Kobayashi T, Tulchin-Francis K, Tullock AMH, Villarosa C, Chan C, Kraus E, Strike S (2018) A little bit faster: lower extremity joint kinematics and kinetics as recreational runners achieve faster speeds. J Biomech 71:167–175. https://doi.org/10.1016/j.jbiomech.2018.02.010

Park SY, Staicu A-M (2015) Longitudinal functional data analysis. Stat 4(1):212–226. https://doi.org/10.1002/sta4.89

Park SY, Staicu A-M, Xiao L, Crainiceanu CM (2018) Simple fixed-effects inference for complex functional models. Biostatistics 19(2):137–152. https://doi.org/10.1093/biostatistics/kxx026

Pataky TC, Vanrenterghem J, Robinson MA (2015) Zero- vs. one-dimensional, parametric vs. non-parametric, and confidence interval vs. hypothesis testing procedures in one-dimensional biomechanical trajectory analysis. J Biomech 48(7):1277–1285. https://doi.org/10.1016/j.jbiomech.2015.02.051

Pinheiro JS, Bates D (2000) Mixed-Effects Models in S And S-PLUS. Springer, New York. https://doi.org/10.1007/b98882

R Core Team: R (2022) A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/

Ramsay JO, Silverman BW (2005) Functional Data Analysis, 2nd edn. Springer Series in Statistics. Springer, New York. https://doi.org/10.1007/b98888

Ruppert D, Wand MP, Carroll RJ (2003) Semiparametric Regression. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9780511755453

Ryan W, Harrison A, Hayes K (2006) Functional data analysis of knee joint kinematics in the vertical jump. Sports Biomech 5(1):121–138. https://doi.org/10.1080/14763141.2006.9628228

Saragiotto BT, Yamato TP, Hespanhol Junior LC, Rainbow MJ, Davis IS, Lopes AD (2014) What are the main risk factors for running-related injuries? Sports Med 44(8):1153–1163. https://doi.org/10.1007/s40279-014-0194-6

Scheipl F, Staicu A-M, Greven S (2015) Functional additive mixed models. J Comput Graph Stat 24(2):477–501. https://doi.org/10.1080/10618600.2014.901914

Taylor WR, Kornaropoulos EI, Duda GN, Kratzenstein S, Ehrig RM, Arampatzis A, Heller MO (2010) Repeatability and reproducibility of OSSCA, a functional approach for assessing the kinematics of the lower limb. Gait Posture 32(2):231–236. https://doi.org/10.1016/j.gaitpost.2010.05.005

Trounson KM, Busch A, Collier NF, Robertson S (2020) Effects of acute wearable resistance loading on overground running lower body kinematics. PLoS ONE 15(12):0244361. https://doi.org/10.1371/journal.pone.0244361

Volkmann A (2021) multifamm: Multivariate Functional Additive Mixed Models. R package version 0.1.1.. https://CRAN.R-project.org/package=multifamm

Volkmann A, Stöcker A, Scheipl F, Greven S (2021) Multivariate functional additive mixed models. Stat Modell. https://doi.org/10.1177/1471082X211056158

Warmenhoven J, Bargary N, Liebl D, Harrison AJ, Robinson MA, Gunning E, Hooker G (2021) PCA of waveforms and functional PCA: a primer for biomechanics. J Biomech 116:110106. https://doi.org/10.1016/j.jbiomech.2020.110106

Willwacher S, Kurz M, Robbin J, Thelen M, Hamill J, Kelly L, Mai P (2022) Running-related biomechanical risk factors for overuse injuries in distance runners: a systematic review considering injury specificity and the potentials for future research. Sports Med 52(8):1863–1877. https://doi.org/10.1007/s40279-022-01666-3

Winter DA (1979) Biomechanics of Human Movement. Wiley, New Jersey

Wood SN (2011) Fast stable restricted maximum likelihood and marginal likelihood estimation of semi-iparametric generalized linear models. J R Stat Soc Series B Stat Methodol 73(1):3–36. https://doi.org/10.1111/j.1467-9868.2010.00749.x

Wood SN (2017) Generalized Additive Models: An Introduction with R, 2nd edn. Chapman & Hall/CRC, Boca Raton. https://doi.org/10.1201/9781315370279

Wu PP-Y, Sterkenburg N, Everett K, Chapman DW, White N, Mengersen K (2019) Predicting fatigue using countermovement jump force-time signatures: PCA can distinguish neuromuscular versus metabolic fatigue. PLoS ONE 14(7):0219295. https://doi.org/10.1371/journal.pone.0219295

Zhang B, Twycross-Lewis R, Großmann H, Morrissey D (2017) Testing gait with ankle-foot orthoses in children with cerebral palsy by using functional mixed-effects analysis of variance. Sci Rep 7(1):11081. https://doi.org/10.1038/s41598-017-11282-1

Zhu H, Morris JS, Wei F, Cox DD (2017) Multivariate functional response regression, with application to fluorescence spectroscopy in a cervical pre-cancer study. Comput Stat Data Anal 111:88–101. https://doi.org/10.1016/j.csda.2017.02.004

Zhu T, Zhang J-T, Cheng M-Y (2022) One-way MANOVA for functional data via Lawley-Hotelling trace test. J Multivar Anal 192:105095. https://doi.org/10.1016/j.jmva.2022.105095

Zohner YEM (2021) Feature Learning and Bayesian Functional Regression for High-Dimensional Complex Data. Rice University, United States - Texas, Ph.D