

Structured functional additive regression in reproducing kernel Hilbert spaces

Hongxiao Zhu,

Virginia Tech, Blacksburg, USA

Fang Yao

University of Toronto, Canada

Hao Helen Zhang

University of Arizona, Tucson, USA

[Received July 2012. Revised May 2013]

Summary. Functional additive models provide a flexible yet simple framework for regressions involving functional predictors. The utilization of a data-driven basis in an additive rather than linear structure naturally extends the classical functional linear model. However, the critical issue of selecting non-linear additive components has been less studied. In this work, we propose a new regularization framework for structure estimation in the context of reproducing kernel Hilbert spaces. The approach proposed takes advantage of functional principal components which greatly facilitates implementation and theoretical analysis. The selection and estimation are achieved by penalized least squares using a penalty which encourages the sparse structure of the additive components. Theoretical properties such as the rate of convergence are investigated. The empirical performance is demonstrated through simulation studies and a real data application.

Keywords: Additive models; Component selection; Functional data analysis; Principal components; Reproducing kernel Hilbert space; Smoothing spline

1. Introduction

Large complex data collected in modern science and technology impose tremendous challenges on traditional statistical methods because of their high dimensionality, massive volume and complicated structures. Emerging as a promising field, functional data analysis employs random functions as model units and is designed to model data distributed over continua such as time, space and wavelength; see Ramsay and Silverman (2005) for a comprehensive introduction. Such data may be viewed as realizations of latent or observed stochastic processes and are commonly encountered in many fields, e.g. longitudinal studies, microarray experiments and brain images.

Regression models involving functional objects play a major role in the functional data analysis literature. The most widely used is the functional linear model, in which a scalar response Y is regressed on a functional predictor X through a linear operator

$$E(Y|X) = \int_{\mathcal{T}} X(t) \beta(t) dt, \quad (1)$$

Address for correspondence: Hongxiao Zhu, Department of Statistics, Virginia Tech, 406-A Hutcheson Hall, Blacksburg, VA 24061-0439, USA.
E-mail: hongxiao@vt.edu

where $X(t)$ is often assumed to be a smooth and square integrable random function defined on a compact domain \mathcal{T} , and $\beta(t)$ is the regression parameter function which is also assumed to be smooth and square integrable. A commonly adopted approach for fitting model (1) is through basis expansion, i.e. representing the functional predictor as linear combinations of a basis $\{\phi_k\}$: $X(t) = \mu(t) + \sum_{k=1}^{\infty} \xi_k \phi_k(t)$, where $\mu(t) = E\{X(t)\}$. Model (1) is then transformed to a linear form of the coefficients $\{\xi_k, k = 1, 2, \dots\}$: $E(Y|X) = b_0 + \sum_{k=1}^{\infty} \xi_k b_k$, where $b_0 = \int_{\mathcal{T}} \beta(t) \mu(t) dt$ and $b_k = \int_{\mathcal{T}} \beta(t) \phi_k(t) dt$. More references on functional linear regression can be found in Cardot *et al.* (1999, 2003), Fan and Zhang (2000), etc. Extensions to generalized functional linear models were proposed by James (2002), Müller and Stadtmüller (2005) and Li *et al.* (2010). The basis set $\{\phi_k\}$ can be either predetermined (e.g. Fourier basis, wavelets or B -splines), or data driven. One convenient choice for the latter is the eigenbasis of the autocovariance operator of X , in which case the random coefficients $\{\xi_k\}$ are called functional principal component (FPC) scores. The FPC scores have zero means and variances equal to the corresponding eigenvalues $\{\lambda_k, k = 1, 2, \dots\}$. This isomorphic representation of X is referred to as Karhunen–Loève expansion, and related methods are often called functional principal component analysis (FPCA) (Rice and Silverman, 1991; Yao *et al.*, 2005; Hall *et al.*, 2006; Hall and Hosseini-Nasab, 2006; Yao, 2007). Owing to the rapid decay of the eigenvalues, the orthonormal eigenbasis provides a more parsimonious and efficient representation compared with other bases. Furthermore, FPC scores are mutually uncorrelated, which can considerably simplify model fitting and theoretical analysis. We focus on the FPC representation of the functional regression throughout this paper; nevertheless, the proposal is also applicable to other prespecified bases.

Although widely used, the linear relationship can be restrictive for general applications. This linear assumption was then relaxed by Müller and Yao (2008) who proposed the functional additive model (FAM). The FAM provides a flexible yet practical framework that accommodates non-linear associations and at the same time avoids the curse of dimensionality that is encountered in high dimensional non-parametric regression problems (Hastie and Tibshirani, 1990). In the case of scalar response, the linear structure was replaced by the sum of non-linear functional components, i.e.

$$E(Y|X) = b_0 + \sum_{k=1}^{\infty} f_k(\xi_k), \quad (2)$$

where $\{f_k(\cdot)\}$ are unknown smooth functions. In Müller and Yao (2008), the FAM was fitted by estimating $\{\xi_k\}$ using FPCA (Yao *et al.*, 2005) and estimating $\{f_k\}$ by using local polynomial smoothing.

Apparently regularizing equation (2) is necessary. In Müller and Yao (2008) the regularization was achieved by truncating the eigensequence to the first K leading components, where K was chosen to explain the majority of the total variation in predictor X . Despite its simplicity, this naive truncation procedure can be inadequate in many complex problems. First, the effect of FPCs on the response does not necessarily coincide with their magnitudes specified by the autocovariance operator of the predictor process alone. For instance, some higher order FPCs may contribute to the regression significantly more than the leading FPCs. This was discussed by Hadi and Ling (1998) in the principal component regression context and later was observed in real examples of high dimensional data (Bair *et al.*, 2006) and functional data (Zhu *et al.*, 2007). Second, although a small number of leading FPCs might be able to capture the major variability in X due to the rapidly decaying eigenvalues, one often needs to include more components for better regression performance, especially for the prediction purpose as observed in Yao and Müller (2010). However, retaining more than needed FPCs brings the risk of overfitting, which is caused by including components that contribute little to the regression

but introduce noise. Therefore a desirable strategy is to identify ‘important’ components out of a sufficiently large number of candidates, whereas to shrink those ‘unimportant’ ones to 0.

With these considerations, we seek an entirely new regularization and estimation framework for identifying the sparse structure of the FAM. Model selection that encourages sparse structure has received substantial attention in the last decade mostly due to the rapidly emerging high dimensional data. In the context of linear regression, the seminal works include the lasso (Tibshirani, 1996), the adaptive lasso (Zou, 2006), the smoothly clipped absolute deviation estimator (Fan and Li, 2001) and the references therein. Traditional additive models were considered by Lin and Zhang (2006), Meier *et al.* (2009) and Ravikumar *et al.* (2009); and extensions to generalized additive models were studied by Wood (2006) and Marra and Wood (2011). In comparison with these works, sparse estimation in functional regression has been much less explored. To our knowledge, most existing works are for functional linear models with sparse penalty (James *et al.*, 2009; Zhu *et al.*, 2010) or L^2 -type penalty (Goldsmith *et al.*, 2011). Relevant research for additive structures is scant in the literature. In this paper, we consider selection and estimation of the additive components in FAMs that encourage a sparse structure, in the framework of a reproducing kernel Hilbert space (RKHS). Unlike in standard additive models, the FPC scores are not directly observed in FAMs. They need to be firstly estimated from the functional covariates and then plugged into the additive model. The estimated scores are random variables, which creates a major challenge to the theoretical exploration. It is necessary to take into account the influence of the unobservable FPC scores on the resulting estimator properly. Furthermore, the functional curve X is not fully observed either. We typically collect repeated and irregularly spaced sample points, which are subject to measurement errors. Measurement error in data adds extra difficulty for model implementation and inference. All of these issues are tackled in this paper. We propose a two-step estimation procedure to achieve the desired sparse structure estimation in FAMs. For the regularization, we adopt the COSSO (Lin and Zhang, 2006) penalty because of its direct shrinkage effect on functions in the RKHS. On the practical side, the method proposed is easy to implement, by taking advantage of existing algorithms of FPCA.

The rest of the paper is organized as follows. In Section 2, we present the proposed approach and algorithm, as well as the theoretical properties of the resulting estimator. Simulation results in comparison with existing methods are included in Section 3. We apply the proposed method to the Tecator data in Section 4, studying the regression of protein content on the absorbance spectrum. Concluding remarks are provided in Section 5, whereas details of the estimation procedure and technical proofs are deferred to the appendices.

The data that are analysed in the paper and the programs that were used to analyse them can be obtained from

<http://wileyonlinelibrary.com/journal/rss-datasets>

2. Structured functional additive model regression

Let Y be a scalar response associated with a functional predictor $X(t)$, $t \in \mathcal{T}$, and let $\{y_i, x_i(\cdot)\}_{i=1}^n$ be independent, identically distributed (IID) realizations of the pair $\{Y, X(\cdot)\}$. The trajectories $\{x_i(t) : t \in \mathcal{T}\}$ are observed intermittently on possibly irregular grids $\mathbf{t}_i = (t_{i1}, \dots, t_{iN_i})^T$. Denote the discretized $x_i(t)$ in vector form by $\mathbf{x}_i = (x_{i1}, \dots, x_{iN_i})^T$. We also assume that the trajectories are subject to IID measurement error, i.e. $x_{ij} = x_i(t_{ij}) + e_{ij}$ with $E(e_{ij}) = 0$ and $\text{var}(e_{ij}) = \nu^2$. Following the FPCA of Yao *et al.* (2005) and Yao (2007), denote by $\boldsymbol{\xi}_{i,\infty} = (\xi_{i1}, \xi_{i2}, \dots)^T$ the sequence of FPC scores of x_i , which is associated with eigenvalues $\{\lambda_1, \lambda_2, \dots\}$ with $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$.

2.1. Proposed methodology

As discussed in Section 1, the theory of FPCA enables isomorphic transformation of random functions to their FPC scores, which brings tremendous convenience to model fitting and theoretical development in functional linear regression. To establish a framework for non-linear and non-parametric regression, we consider regressing the scalar responses $\{y_i\}$ directly on the sequences of FPC scores $\{\xi_{i,\infty}\}$ of $\{x_i\}$. For the convenience of model regularization, we would like to restrict the predictor variables (i.e. FPC scores) to taking values in a closed and bounded subset of the real line, e.g. $[0, 1]$ without loss of generality. This is easy to achieve by taking a transformation of the FPC scores through a monotonic function $\Psi: \mathfrak{R} \rightarrow [0, 1]$, for all $\{\xi_{ik}\}$. In fact the choice of Ψ is quite flexible. A wide range of cumulative distribution functions (CDFs) can be used; see assumption 2 in Section 2.2 for the regularity condition. Additionally one may choose Ψ so that the transformed variables have similar or the same variations. This can be achieved by allowing $\Psi(\cdot)$ to depend on the eigenvalues $\{\lambda_k\}$, where $\{\lambda_k\}$ serve as scaling variables. For simplicity, in what follows we use a suitable CDF (e.g. normal), denoted by $\Psi(\cdot, \lambda_k)$, from a location–scale family with zero mean and variance λ_k . It is obvious that, if ξ_{ik} s are normally distributed, the normal CDF leads to uniformly distributed transformed variables on $[0, 1]$.

Denoting the transformed variable of ξ_{ik} by ζ_{ik} , i.e. $\zeta_{ik} = \Psi(\xi_{ik}, \lambda_k)$, and denoting $\zeta_{i,\infty} = (\zeta_{i1}, \zeta_{i2}, \dots)^T$, we propose an additive model as follows:

$$y_i = b_0 + \sum_{k=1}^{\infty} f_{0k}(\zeta_{ik}) + \varepsilon_i, \quad (3)$$

where $\{\varepsilon_i\}$ are independent errors with zero mean and variance σ^2 , and $f_0(\zeta_{i,\infty}) = b_0 + \sum_{k=1}^{\infty} f_{0k}(\zeta_{ik})$ is a smooth function. For each k , let H^k be the l th-order Sobolev Hilbert space on $[0, 1]$, defined by

$$H^k([0, 1]) = \{g | g^{(\nu)} \text{ is absolutely continuous for } \nu = 0, 1, \dots, l-1; g^{(l)} \in L^2\}.$$

One can show that H^k is an RKHS equipped with the norm

$$\|g\|^2 = \sum_{\nu=0}^{l-1} \left\{ \int_0^1 g^{(\nu)}(t) dt \right\}^2 + \int_0^1 g^{(l)}(t)^2 dt.$$

See Wahba (1990) and Lin and Zhang (2006) for more details. Note that H^k has the orthogonal decomposition $H^k = \{1\} \oplus \bar{H}^k$. Then the additive function f_0 corresponds to \mathcal{F} which is a direct sum of subspaces, i.e. $\mathcal{F} = \{1\} \oplus \sum_{k=1}^{\infty} \bar{H}^k$ with $f_{0k} \in \bar{H}^k$, for all k . It is easy to check that, for any $f = b + \sum_k f_k \in \mathcal{F}$, we have $\|f\|^2 = b^2 + \sum_{k=1}^{\infty} \|f_k\|^2$. In this paper, we take $l=2$ but the results can be extended to other cases straightforwardly. To distinguish the Sobolev norm from the L^2 -norm, we write $\|\cdot\|$ for the former and $\|\cdot\|_{L^2}$ for the latter.

As motivated in Section 1, it is desirable to impose some type of regularization condition on model (3) to select important components. An important assumption that is commonly made in high dimensional linear regression is the sparse structure of the underlying true model. This assumption is also critical in the context of functional data analysis, which enables us to develop a more systematic strategy than the heuristic truncation that retains the leading FPCs. Although widely adopted, retaining the leading FPCs is a strategy that is guided solely by the covariance operator of the predictor X , and therefore it fails to take into account the response Y . To be more flexible, we assume that the number of important functional additive components that contribute to the response is finite, but not necessarily restricted to the leading terms. In particular, we denote \mathcal{I} the index set of the important components and assume that $|\mathcal{I}| < \infty$, where $|\cdot|$ denotes the cardinality of a set. In other words, there is a sufficiently large s such that $\mathcal{I} \subseteq \{1, \dots, s\}$, which implies that $f_k \equiv 0$ as long as $k > s$. The FAM is thus equivalent to

$$y_i = b_0 + \sum_{k=1}^s f_{0k}(\zeta_{ik}) + \varepsilon_i. \quad (4)$$

It is noted that the initial truncation s merely controls the total number of additive components to be considered, which is different from the heuristic truncation that was suggested by Yao *et al.* (2005) and Müller and Yao (2008) based on model selection criteria such as cross-validation, the Akaike information criterion AIC or the fraction of variation explained. In practice we suggest choosing s large so that nearly 100% of the total variation is explained. This often leads to more than 10 FPCs in most empirical cases.

With this assumption, the regression function $f_0(\zeta) = b_0 + \sum_{k=1}^s f_{0k}(\zeta_k)$ lies in the truncated subspace $\mathcal{F}^s = \{1\} \oplus \sum_{k=1}^s \bar{H}^k$ of \mathcal{F} , where ζ is the truncated version of ζ_∞ , i.e. $\zeta = (\zeta_1, \dots, \zeta_s)^\top$ with the dependence on s suppressed if no confusion arises. To regularize the unknown smooth functions $\{f_{0k}\}$ non-parametrically, we employ the COSSO regularization defined for function estimation in RKHS and estimate f_0 by finding $f \in \mathcal{F}^s$ that minimizes

$$Q(f|\zeta_i) = \frac{1}{n} \sum_{i=1}^n \{y_i - f(\zeta_i)\}^2 + \tau_n^2 J(f), \quad J(f) = \sum_{k=1}^s \|P^k f\|, \quad (5)$$

where $P^k f$ is the orthogonal projection of f onto \bar{H}^k . Here τ_n is the only smoothing parameter that requires tuning, whereas the common smoothing spline approach involves multiple smoothing parameters. The penalty $J(f)$ is a convex functional and is a pseudonorm in \mathcal{F}^s . One interesting connection between COSSO and the lasso is that, when $f_{0k}(\zeta_k) = \zeta_k \beta_{0k}$, the penalty in expression (5) reduces to $\sum_{k=1}^s |\zeta_k \beta_{0k}|$, which becomes the adaptive lasso penalty (Zou, 2006).

Different from the standard additive regression models, the transformed FPC scores $\{\zeta_i\}$ serving as predictor variables in expression (5) cannot be observed. Therefore we need to estimate the FPC scores first before the estimation and structure selection of f . A simple two-step algorithm is given as follows.

Step 1: implement FPCA to estimate the FPC scores $\{\xi_{i1}, \dots, \xi_{is}\}$ of x_i , and then the transformed variables $\hat{\zeta}_{ik} = \Psi(\hat{\xi}_{ik}, \hat{\lambda}_k)$, where $\hat{\lambda}_k$ is the estimated eigenvalue, and s is chosen to explain nearly 100% of the total variation.

Step 2: implement the COSSO algorithm of Lin and Zhang (2006) to solve

$$\min_{f \in \mathcal{F}^s} Q(f|\hat{\zeta}_i) = \min_{f \in \mathcal{F}^s} \frac{1}{n} \sum_{i=1}^n \{y_i - f(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f), \quad J(f) = \sum_{k=1}^s \|P^k f\|. \quad (6)$$

We refer to Appendix A for details in the case of densely or sparsely observed predictor trajectories. We call the proposed method the component selection and estimation for functional additive model (CSEFAM).

2.2. Theoretical properties

We focus on consistency of the resulting estimator of the CSEFAM for the case when $\{x_i(t)\}$ are densely observed in this subsection, where the rate of convergence is assessed by using the empirical norm. In particular, we introduce the empirical norm and the entropy of \mathcal{F}^s as follows. Let $g \in \mathcal{F}^s$; the empirical norm of g is defined as $\|g\|_n = \sqrt{\{(1/n) \sum_{i=1}^n g(\zeta_i)\}^2}$. The empirical inner product of the error term ε and g is defined as $(\varepsilon, g)_n = (1/n) \sum_{i=1}^n \varepsilon_i g(\zeta_i)$. Similarly, the empirical inner product of f and g in \mathcal{F}^s is $(f, g)_n = (1/n) \sum_{i=1}^n f(\zeta_i) g(\zeta_i)$.

The assumptions on the regression function f and the transformation $\Psi(\cdot, \cdot)$ are listed below in assumptions 1 and 2, whereas the commonly adopted regularity conditions on the functional

predictors $\{x_i(t)\}$, the dense design, and the smoothing procedures are deferred to conditions 1–3 in Appendix B.

Assumption 1. For any $f \in \mathcal{F}^s$, there are independent $\{B_i\}_{i=1}^n$ with $E(B_i^2) < \infty$, such that, with probability 1,

$$\left| \frac{\partial f(\zeta_i)}{\partial \zeta_{ik}} \right| \leq B_i \|f\|_{L^2}.$$

Assumption 2. The transformation function $\Psi(\xi, \lambda)$ is differentiable at ξ and λ , and satisfies that $|\partial \Psi(\xi, \lambda)/\partial \xi| \leq C\lambda^\gamma$ and $|\partial \Psi(\xi, \lambda)/\partial \lambda| \leq C\lambda^\gamma |\xi|$ for some constant C and γ ($\gamma < 0$).

Assumption 1 is a regularization condition that controls the amount of fluctuation in f relative to its L^2 -norm. For assumption 2, one can easily verify that, if choosing $\Psi(\cdot, \cdot)$ to be the normal CDF with zero mean and variance λ , then $C = 1$ and $\gamma = -\frac{1}{2}$ (when $\lambda \geq 1$) or $\gamma = -\frac{3}{2}$ (when $0 < \lambda < 1$). One can also choose the CDF of Student's t - or other distributions with variances λ .

For brevity of the presentation, the technical lemmas and proofs are deferred to Appendix B. It is noted that the existence of the minimizer for criterion (5) is guaranteed by analogy with theorem 1 of Lin and Zhang (2006), by considering a design conditional on the input $\{y_i, \zeta_{i1}, \dots, \zeta_{is}\}$, $i = 1, \dots, n$, where s is the initial truncation parameter.

Theorem 1. Consider the regression model (4) with $\zeta_{ik} = \Psi(\xi_{ik}, \lambda_k)$, where $\{\xi_{ik}\}_{k=1}^s$ are FPC scores of $x_i(t)$ based on densely observed trajectories, and $\{\lambda_k\}_{k=1}^s$ are the corresponding eigenvalues. Let \hat{f} be the minimizer of the target function (6) over $f \in \mathcal{F}^s$, and let τ_n be the tuning parameter in function (6). Assume that assumptions 1 and 2 and conditions 1–3 hold. If $J(f_0) > 0$ and

$$\tau_n^{-1} = n^{2/5} J^{3/10}(f_0), \quad (7)$$

then $\|\hat{f} - f_0\|_n = O_p(n^{-2/5}) J^{1/5}(f_0)$ and $J(\hat{f}) = J(f_0) O_p(1)$. If $J(f_0) = 0$ and

$$\tau_n^{-1} = n^{1/4}, \quad (8)$$

then $\|\hat{f} - f_0\|_n = O_p(n^{-1/2})$ and $J(\hat{f}) = O_p(n^{-1/2})$.

It is worth mentioning that the technical difficulty arises from the unobserved variables ζ_i , and major effort has been devoted to tackle the influence of the estimated quantities $\hat{\zeta}_i$ on the resulting estimator by utilizing the analytical tools from the spectral decomposition of the autocovariance operator of X . Theorem 1 suggests that, if the repeated measures that are observed for all individuals are sufficiently dense and $J(f_0)$ is bounded, the resulting estimator \hat{f} obtained from expression (6) has rate of convergence $n^{-2/5}$, which is the same as the rate when $\{\zeta_i\}$ are directly observed.

3. Simulation studies

To demonstrate the performance of the CSEFAM approach proposed, we conduct simulation studies under different settings. In Sections 3.1 and 3.2, we study the performance of the CSEFAM for dense and sparse functional data respectively, assuming that the underlying true model contains both ‘important’ and ‘unimportant’ additive components. We compare the CSEFAM approach with the FAM-type methods and the multivariate adaptive regression spline (MARS) approach. The FAM-type methods are implemented in three different ways, two of which are the ‘oracle’ methods, FAM_{O1} and FAM_{O2} , both assuming full knowledge of the underlying model structure. In particular, the FAM_{O1} method serves as the gold standard, in which both

the true values of $\{\zeta_{ik}\}$ and the true non-vanishing additive components are used. The $\text{FAM}_{\text{O}2}$ method is another type of oracle, in which the values of $\{\zeta_{ik}\}$ are estimated through FPCA, but the true non-vanishing additive components are used. In Section 3.3, we study the performance of the CSEFAM when the underlying true model is actually non-sparse, and we compare the results with the saturated and truncated FAMs. For each setting, we perform 100 Monte Carlo simulations and present the model selection and prediction results for the methods under comparison.

3.1. Dense functional data

We generate 1000 IID trajectories by using 20 eigenfunctions, among which $n = 200$ are randomly allocated to the training set and the other 800 form the test set. The functional predictors $x_i(t)$, $t \in [0, 10]$, are measured over a grid with 100 equally spaced points, with independent measurement error $e_{ij} \sim N(0, v^2)$, $v^2 = 0.2$. The eigenvalues of $x_i(t)$ are generated by $\lambda_k = ab^{k-1}$ with $a = 45.25$ and $b = 0.64$. The true FPC scores $\{\xi_{ik}\}$ are generated from $N(0, \lambda_k)$, and the eigenbasis $\{\phi_k(\cdot)\}$ is taken to be the first 20 Fourier basis functions on $[0, 10]$. The mean curve is set to be $\mu_x(t) = t + \sin(t)$. We use the normal CDF to obtain the transformed variables: $\zeta_k = \Psi(\xi_k; 0, \lambda_k)$, $k = 1, \dots, 20$. The values of y_i are then generated by $y_i = f_0(\zeta_i) + \varepsilon_i$, where $\varepsilon_i \sim N(0, \sigma^2)$ and $\sigma^2 = 1$. We assume that f_0 depends on only three non-zero additive components: the first, the second and the fourth, i.e. $f_0(\zeta_i) = b_0 + f_{01}(\zeta_{i1}) + f_{02}(\zeta_{i2}) + f_{04}(\zeta_{i4})$, $\mathcal{I} = \{1, 2, 4\}$. Here we take $b_0 = 1.4$, $f_{01}(\zeta_1) = 3\zeta_1 - \frac{3}{2}$, $f_{02}(\zeta_2) = \sin\{2\pi(\zeta_2 - \frac{1}{2})\}$, $f_{04}(\zeta_4) = 8(\zeta_4 - \frac{1}{3})^2 - 8/9$ and $f_{0k}(\zeta_k) \equiv 0$ for $k \notin \mathcal{I}$. This gives the signal-to-noise ratio (SNR) 2.2, where the SNR is defined as $\text{SNR} = \text{var}\{f_0(\zeta)\} / \text{var}(\varepsilon)$, and $\text{var}\{f_0(\zeta)\} = \sum_{k \in \mathcal{I}} \int_0^1 f_{0k}^2(\zeta_k) d\zeta_k = 2.2$ given that $\zeta_k \sim U[0, 1]$.

We apply the proposed CSEFAM algorithm to the training data, following the FPCA and COSSO steps described in Section 2.1 and Appendix A. For illustration, we pick one Monte Carlo simulation and display the component selection and estimation results in Fig. 1. In FPCA, the initial truncation is $s = 18$, accounting for nearly 100% of the total variation, and is passed to the COSSO step. The component selection is then achieved by tuning the regularization parameters λ_0 in expression (9) with generalized cross-validation and M in expression (10) with the Bayesian information criterion (BIC), illustrated in Figs 1(a) and 1(b) whereas the empirical L_1 -norms of \hat{f}_k (computed by $n^{-1} \sum_{i=1}^n |\hat{f}_k(\zeta_{ik})|$ at different M) are shown in Fig. 1(c). In Figs 1(d)–1(f), the resulting estimates of f_k , $k = 1, 2, 4$, are displayed, and $\{\hat{f}_k, k \neq 1, 2, 4\}$ are shrunk to 0 as desired.

The model selection and prediction results are presented in the top panel of Table 1. We implement the FAM procedure in a different manner from that in Müller and Yao (2008). Instead of using local polynomial smoothing for estimating each f_k separately, we perform a more general additive fitting, the generalized additive model, on the transformed FPC scores which allows backfitting and also provides a p -value for each additive component. The only reason for doing so is that the generalized additive model algorithm shows more numerical stability especially when the number of additive components is large. Owing to the use of the true model structure, both oracle methods $\text{FAM}_{\text{O}1}$ and $\text{FAM}_{\text{O}2}$ are expected to outperform the rest. Because of the estimation error that is induced in the FPCA step, $\text{FAM}_{\text{O}2}$ is expected to sacrifice certain estimation accuracy and prediction power compared with $\text{FAM}_{\text{O}1}$. The FAM_{S} model is the saturated model based on the estimated FPC scores and the leading s terms used in the CSEFAM. No model selection is performed in FAM_{S} . The s -values vary from 17 to 19 which take into account nearly 100% of the total variation of $\{x_i(t)\}$. The MARS method is based on Hastie *et al.* (2001).

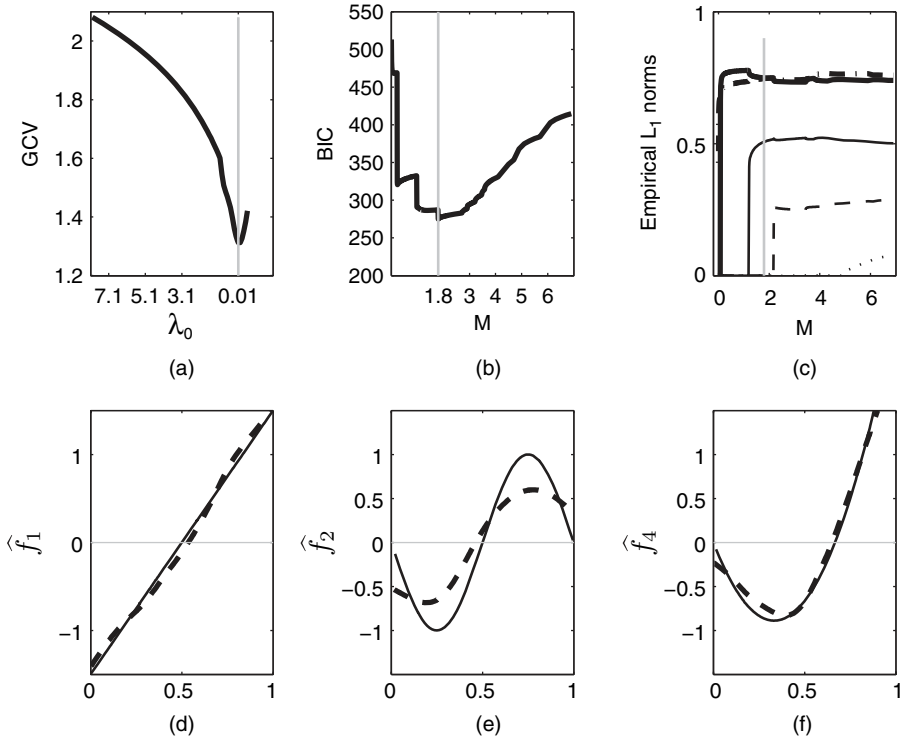


Fig. 1. Plots of component selection and estimation from one simulation: (a) generalized cross-validation versus λ_0 ; (b) BIC versus M ; (c) empirical L_1 -norms at various M -values (—, f_1 ; — —, f_2 ; - · - ·, f_3 ; · · · ·, f_4 ; · · · · ·, f_5) (|, tuning parameter chosen in (a)–(c)); (d)–(f) estimated \hat{f}_k s (— —) versus true f_k s (—) for $k = 1, 2, 4$

It is noted that the subjective truncation based on the explained variation in X is suboptimal for regression purpose (for conciseness the results are not reported). Therefore, in Table 1, we report (under the ‘counts for the following model sizes’ columns) the counts of selected numbers of non-vanishing additive components in the CSEFAM, and the counts of the number of significantly non-zero additive components in FAM, FAM_{O1} and FAM_{O2}. For convenience of display, only the counts for model size up to 8 are reported. The ‘selection frequencies for the following components’ columns of Table 1 record the number of times that each additive component is estimated to be non-zero for the first eight components. For the MARS method, if the j th component \hat{f}_j is selected in one or more basis functions, we counted it as 1 and 0 otherwise. Regarding the prediction error (PE), we use the population estimate from the training set (e.g. the mean, covariance and eigenbasis) to obtain the FPC scores for both training and test set; then we apply the $\{\hat{f}_k\}$ estimated from the training set to obtain predictions for $\{y_i\}$ in the test set. The PEs are calculated by $n^{-1} \sum_{i=1}^n (y_i - \hat{y}_i)^2$. From the top panel of Table 1, we see that, under the dense design, the CSEFAM chooses the correct models (with model size equal to 3) 61% of the time whereas the FAM_S method always overselects ($\alpha = 0.05$ is used to retain significant additive components). The PE of CSEFAM is the smallest among the three non-oracle models. Compared with the oracle methods, the CSEFAM has less prediction power than FAM_{O2} (slightly) and FAM_{O1}, which can be regarded as the price paid by both estimating the ζ and selecting the additive components.

Table 1. Summary of the model selection and prediction in 100 Monte Carlo simulations under the dense and sparse design

Data	Model	Counts for the following model sizes:								Selection frequencies for the following components:								PE
		1	2	3	4	5	6	7	8	\hat{f}_1	\hat{f}_2	\hat{f}_3	\hat{f}_4	\hat{f}_5	\hat{f}_6	\hat{f}_7	\hat{f}_8	
Dense design	CSEFAM	0	5	61	29	5	0	0	0	100	94	22	100	7	3	0	1	1.30 (0.13)
	FAMS	0	0	10	32	21	21	8	4	100	98	51	100	32	14	12	8	1.50 (0.17)
	MARS	—	—	—	—	—	—	—	—	100	99	60	100	41	23	25	18	1.46 (0.16)
	FAM02	0	1	99	—	—	—	—	—	100	99	—	100	—	—	—	—	1.28 (0.12)
	FAM01	0	0	100	—	—	—	—	—	100	100	—	100	—	—	—	—	1.07 (0.06)
	CSEFAM	0	22	61	13	4	0	0	0	100	78	10	82	12	9	7	1	2.07 (0.16)
Sparse design	FAMS	0	0	14	30	25	20	9	2	100	98	41	96	35	17	9	12	2.17 (0.16)
	MARS	—	—	—	—	—	—	—	—	100	98	58	98	56	30	20	23	2.11 (0.14)
	FAM02	0	4	96	—	—	—	—	—	100	98	—	98	—	—	—	—	2.01 (0.14)
	FAM01	0	0	100	—	—	—	—	—	100	100	—	100	—	—	—	—	1.05 (0.05)

Table 2. AISE for 100 Monte Carlo simulations under the dense and sparse design

<i>Data</i>	<i>Model</i>	<i>AISEs for the following functions:</i>								
		f_1	f_2	f_3	f_4	f_5	f_6	f_7	f_8	f
Dense design	CSEFAM	0.038	0.117	0.022	0.038	0.005	0.001	0.000	0.001	0.226
	FAM _S	0.030	0.095	0.050	0.047	0.031	0.018	0.016	0.015	0.476
	FAM _{O2}	0.027	0.090	—	0.042	—	—	—	—	0.158
	FAM _{O1}	0.007	0.028	—	0.019	—	—	—	—	0.054
Sparse design	CSEFAM	0.033	0.22	0.036	0.298	0.055	0.040	0.045	0.001	0.720
	FAM _S	0.016	0.118	0.032	0.159	0.102	0.121	0.399	2.64	$> 10^3$
	FAM _{O2}	0.026	0.129	—	0.220	—	—	—	—	0.376
	FAM _{O1}	0.007	0.016	—	0.013	—	—	—	—	0.036

To assess the accuracy of estimation, the averaged integrated squared errors (AISEs) for the first eight additive components and for the overall function f are presented in the top panel of Table 2, where ISE is defined by

$$ISE(f_k) = E_{\zeta_k} \{ \hat{f}_k(\zeta_k) - f_k(\zeta_k) \}^2 = \int_0^1 \{ \hat{f}_k(t) - f_k(t) \}^2 dt.$$

From Table 2, we see that the CSEFAM provides considerably smaller AISE for the truly zero components ($f_j, j = 3, 5, 6, 7, 8$) than the FAM_S method. For the non-zero components, the CSEFAM, FAM_S and FAM_{O2} have comparable AISE values.

3.2. Sparse functional data

To compare with the dense case, we also conducted a simulation to examine the performance of the CSEFAM for sparse functional data. We generated 1200 IID trajectories, with 300 in the training set and 900 in the test set. In each trajectory, there are 5–10 repeated observations uniformly located in $[0, 10]$, with the number of points chosen from 5 to 10 with equal probabilities. The other settings are the same as in the dense design. The summaries of the model selection, prediction and estimation results are presented in the bottom panel of Table 1 and Table 2. We observe a similar pattern to that in the dense design case. Moreover, Table 2 suggests that, for the sparse design, the FAM_S estimate of f_k becomes quite unstable for higher order components (e.g. $k > 7$). The AISE increases rapidly owing to the influence of outlying estimates. This is not a surprise, because under the sparse design the high order eigenfunctions and FPC scores are difficult to estimate accurately owing to the sparseness of the data and the moderate sample size, which lead to inaccurate f_k -estimates when the saturated model FAM_S is used. In this situation, we see that the proposed CSEFAM still performs quite stably, since the COSSO penalty has the effect of automatically downweighting the ‘unimportant’ components. This provides further support for the proposed CSEFAM approach.

3.3. Non-sparse underlying additive components

To show the model performance when the true additive components are actually non-sparse, we conduct an additional simulation with two settings (study I and study II) for the dense design, and we compare the CSEFAM with two versions of the FAM: the saturated model FAM_S as defined in Section 3.1, and the truncated model FAM_T with a truncation chosen to retain 99% of the total variation. In study I, the true model contains three ‘larger’ additive components $\{f_{01}, f_{02}, f_{04}\}$, taking the same form as in Section 3.1 except being rescaled by a constant $\frac{1}{2}$. The

Table 3. Additional simulation for cases with non-sparse additive components†

Type	Model	Counts for the following model sizes:								Selection frequencies for the following components:								PE	AISE of f	
		1	2	3	4	5	6	7	8	\hat{f}_1	\hat{f}_2	\hat{f}_3	\hat{f}_4	\hat{f}_5	\hat{f}_6	\hat{f}_7	\hat{f}_8			
I	CSEFAM	0	3	20	34	26	9	6	2	100	88	12	100	17	12	10	10	10	1.19 (0.08)	0.17
	FAM _S	0	4	16	18	29	17	7	8	100	92	16	99	20	17	19	16	16	1.33 (0.12)	0.33
	FAM _T	0	4	39	35	15	6	0	1	100	91	15	100	19	15	9	9	9	1.22 (0.08)	0.18
II	CSEFAM	1	2	4	12	13	20	26	13	46	42	33	42	36	42	38	44	44	1.25 (0.07)	0.12
	FAM _S	1	6	8	25	14	13	13	10	42	45	29	37	29	38	36	36	36	1.38 (0.11)	0.42
	FAM _T	13	30	22	20	6	6	2	0	34	35	20	31	25	38	34	30	30	1.32 (0.08)	0.20

†I, the true model contains both 'larger' and 'smaller' additive components; II, the true model contains only small additive components.

rest are ‘smaller’ additive components, each randomly selected from $\{f_{01}, f_{02}, f_{04}\}$ with equal probability and rescaled by a smaller constant uniformly chosen from $[1/17, 1/14]$. The data generated have a lower (more challenging) SNR around 0.60, among which 8.7% are from the ‘smaller’ components. The results are listed in the top panel of Table 3, which shows that the CSEFAM tends to favour smaller model size than FAM_S . We also observe that the model size of FAM_T tends to be smaller than for the CSEFAM since FAM_T adopts more truncation with the 99% threshold. It is important to note that the CSEFAM in fact yields PE and AISE that are substantially smaller than the FAM_S method, and the results of the CSEFAM are comparable with that of FAM_T . In study II, we replace the three larger components by the smaller ones; therefore all additive components have roughly equal small contributions. We select the scaling constant uniformly from $[\frac{1}{8}, \frac{1}{6}]$ so that the total SNR is 0.30 on average. The results listed in the bottom panel of Table 3 suggest that the CSEFAM now tends to select more components (i.e. to produce non-sparse fits) and again yields smaller PE and AISE than both the FAM_S and the FAM_T methods. Overall, this simulation suggests that the proposed CSEFAM is still a reasonable option even if the underlying true model is non-sparse. It is also worth mentioning that the gain of the CSEFAM is more apparent in the challenging settings with low SNR.

4. Real data application

We demonstrate the performance of the proposed method through the regression of protein




Fig. 2. (a) Near infrared absorbance spectral curves and (b) the first five estimated eigenfunctions (—, $\phi_1(t)$; — — —, $\phi_2(t)$; - - -, $\phi_3(t)$; - · - ·, $\phi_4(t)$; ·····, $\phi_5(t)$)

content on the near infrared absorbance spectrum measured over 240 meat samples. The data set is collected by the Tecator company and is publicly available on the StatLib Web site (<http://lib.stat.cmu.edu>). The measurements were made through a spectrometer named the Tecator Infratec Food and Feed Analyzer. The spectral curves were recorded at wavelengths ranging from 850 nm to 1050 nm. For each meat sample the data consist of a 100-channel spectrum of absorbances (100 grid points) as well as the contents of moisture (water), fat and protein. The absorbance is the negative common logarithm of the transmittance measured by the spectrometer. The three contents, measured in percentages, are determined by analytic chemistry. Of primary interest is to predict the protein content by using the spectral trajectories. The 240 meat samples were randomly split into a training set (with 185 samples) and a test set (with 55 samples). We aim to predict the content of protein in the test set by using the training data. Fig. 2 illustrates the spectral curves and the first five eigenfunctions estimated by using FPCA.

We initially retain the first 20 FPCs which take into account nearly 100% of the total variation. The proposed CSEFAM is then applied for component selection and estimation. The

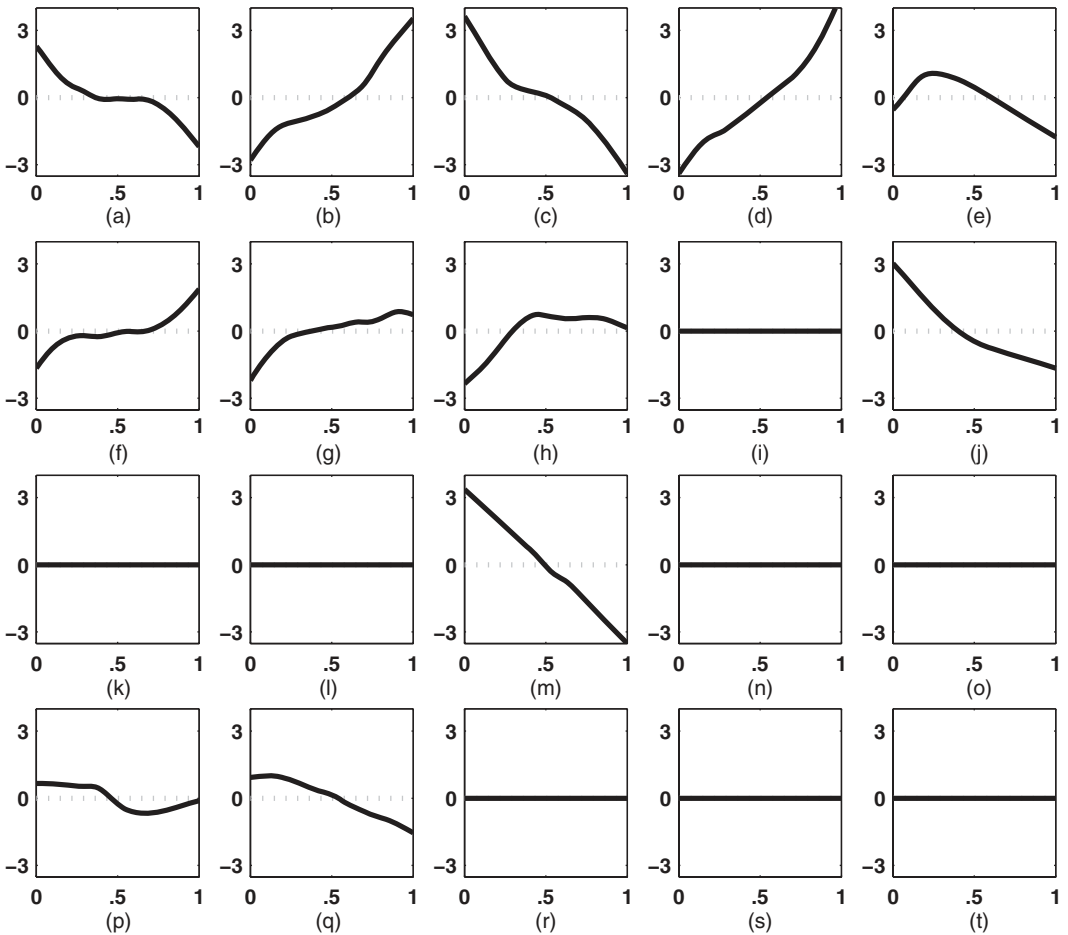


Fig. 3. Plots of the estimated additive components: (a) \hat{f}_1 ; (b) \hat{f}_2 ; (c) \hat{f}_3 ; (d) \hat{f}_4 ; (e) \hat{f}_5 ; (f) \hat{f}_6 ; (g) \hat{f}_7 ; (h) \hat{f}_8 ; (i) \hat{f}_9 ; (j) \hat{f}_{10} ; (k) \hat{f}_{11} ; (l) \hat{f}_{12} ; (m) \hat{f}_{13} ; (n) \hat{f}_{14} ; (o) \hat{f}_{15} ; (p) \hat{f}_{16} ; (q) \hat{f}_{17} ; (r) \hat{f}_{18} ; (s) \hat{f}_{19} ; (t) \hat{f}_{20}

Table 4. Prediction results on the test set compared with several other methods†

Results for the following methods:								
	CSEFAM		FAM			MARS	Partial least squares, PLD20	Functional linear model, AIC7
	$s=10$	$s=20$	PC5	PC10	PC20	PC20		
PE	2.22	0.72	3.98	2.13	0.84	0.77	1.02	1.50
R_Q^2	0.82	0.94	0.68	0.83	0.93	0.93	0.92	0.88

†PC10 indicates that 10 FPC scores are used. PLD20 indicates that the number of partial least squares directions used is 20. AIC7 indicates that seven FPC scores are used based on the Akaike information criterion.

determination of the tuning parameters in the COSSO step is guided by the generalized cross-validation criterion for λ_0 , which gives $\lambda_0 = 0.0013$, and by tenfold cross-validation for M , which gives $M = 10.0$. The estimated additive components are plotted in Fig. 3, from which we see that the CSEFAM selects 12 out of the 20 components, $\{\hat{f}_1, \dots, \hat{f}_8, \hat{f}_{10}, \hat{f}_{13}, \hat{f}_{16}, \hat{f}_{17}\}$, and the other components are estimated to be 0. To assess the performance of the method proposed, we report the PE on the test set in Table 4, where the PE is calculated in the same way as in Section 3. We also report the quasi- R^2 for the test set, which is defined as

$$R_Q^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y}_i)^2}.$$

To show the influence of the initial truncation, we also use a smaller value of s , $s = 10$ in the CSEFAM, which gives suboptimal results. This suggests that we shall use a sufficiently large s to begin with. The FAM is carried out with the leading five, 10 and 20 FPCs. An interesting phenomenon is that, though the high order FPCs (over 10) explain very little variation of the functional predictor (less than 1%), their contribution to the prediction is surprisingly substantial. Such phenomena are also observed for the MARS method and partial least squares (which is a popular approach in chemometrics; see Xu *et al.* (2007) and the references therein). One more comparison is with the classical functional linear model with the estimated leading FPCs served as predictors, where a heuristic AIC is used to choose the first seven components.

From Table 4, we see that, when the initial truncation is set at 10, the proposed CSEFAM is not obviously advantageous compared with the FAM. As the number of FPCs increases to 20, the method proposed provides a much smaller PE and higher R_Q^2 than all other methods. A sensible explanation is that, for these data, most of the first 10 FPCs (except the ninth) have non-zero contributions to the response (shown in Fig. 3); therefore penalizing these components does not help to improve the prediction. However, as the number of FPC scores increases, more redundant terms come into play, so the penalized method the CSEFAM gains more prediction power. We have repeated this analysis for different random splits of the training and test sets, and the conclusions stay virtually the same.

5. Discussion

We proposed a structure estimation method for functional data regression where a scalar response is regressed on a functional predictor. The model is constructed in the framework of FAMs, where the additive components are functions of the scaled FPC scores. The selection

and estimation of the additive components are performed through penalized least squares using the COSSO penalty in the context of RKHS. The method proposed allows for more general non-parametric relationships between the response and predictors and therefore serves as an important extension of functional linear regression. Through the adoption of the additive structure, it avoids the curse of dimensionality that is caused by the infinite dimensional predictor process. The method proposed provides a way to select the important features of the predictor process and to shrink the unimportant ones to 0 simultaneously. This selection scenario takes into account not only the explained variation of the predictor process, but also its contribution to the response. The theoretical result shows that, under the dense design, the non-parametric rate from component selection and estimation will dominate the discrepancy due to the unobservable FPC scores.

A concern raised is whether the sparsity is necessary in the FAM framework. The sparseness assumption in general helps to balance the trade-off between variance and bias, which may lead to improved model performance. This can be particularly useful when part of the predictor has negligible contribution to the regression. Even if the underlying model is in fact non-sparse and we care only about estimation and prediction, the proposed CSEFAM is still a reasonable option, as illustrated by the simulation in Section 3.3. We also point out that, when all non-zero additive components are linear, the COSSO penalty reduces to the adaptive lasso penalty. An additional simulation (which for conciseness is not reported) has shown that the method proposed produces estimation and prediction results that are comparable with those of the adaptive lasso. Moreover, the COSSO penalty requires that $s < n$, which does not conflict with the requirement that the initial truncation s is chosen sufficiently large to include all important features. In practice the number of FPCs accounting for nearly 100% predictor variation is often far less than the sample size n owing to the fast decay of the eigenvalues. Finally both simulated and real examples indicate that the model performance is not sensitive to s as long as it is chosen to be sufficiently large.

On the computation side, our algorithm takes advantage of both FPCA and COSSO. On a desktop with Intel(R) Core(TM) i5-2400 central processor unit with a 3.10-GHz processor and 8 Gbytes random-access memory each Monte Carlo sample in Section 3.1 takes 30 s and the real data analysis takes about 10 s. As far as the dimensionality is concerned, the capacity and speed depend on the particular FPCA algorithm used. We have used the principal component analysis by conditional expectation algorithm PACE which can deal with fairly large data (<http://anson.ucdavis.edu/~ntyang/PACE/>). For dense functional data with 5000 or more dimensions, pre-binning is suggested to accelerate the computation. An FPCA algorithm geared towards extremely large dimensions (with an identical time grid for all subjects) is also available; for instance, Zipunnikov *et al.* (2011) considered functional magnetic resonance imaging data with dimension of the order of $O(10^7)$ through partitioning the original data matrix to blocks and performing singular value decomposition using blockwise operation.

Although we have focused on the FPC-based analysis in this work, the CSEFAM framework is generally applicable to other basis structures, e.g. splines and wavelets, where the additive components are functions of the corresponding basis coefficients of the predictor process. It may also work for non-parametric penalties other than COSSO, such as the sparsity smoothness penalty that was proposed in Meier *et al.* (2009). The method proposed may be further extended to accommodate categorical responses, where an appropriate link function can be chosen to associate the mean response with the additive structure. Another possible extension is regression with multiple functional predictors, where component selection can be performed for selecting functional predictors. In this case the additive components that are associated with each functional predictor need to be selected in a group manner.

Acknowledgements

This work was conducted through the ‘Analysis of object data’ programme at the Statistical and Applied Mathematical Sciences Institute, USA. Fang Yao’s research was partially supported by an individual discovery grant and discovery accelerator supplement from the Natural Sciences and Engineering Research Council, Canada. Hao Helen Zhang was supported by US National Institutes of Health grant R01 CA-085848 and National Science Foundation grant DMS-0645293.

Appendix A: The estimation procedure

To estimate ζ_i , we assume that the functional predictors are observed with measurement error on a grid of \mathcal{T} . We adopt two different procedures for functional data that are either densely or sparsely observed.

- (a) *Obtain $\hat{\zeta}_i$ in the dense design.* If $\{x_i(t)\}$ are observed on a sufficiently dense grid for each subject, we apply local linear smoothing to the data $\{t_{ij}, x_{ij}\}_{j=1, \dots, N_i}$ individually, which gives the smooth approximation $\hat{x}_i(t)$. The mean and covariance function are obtained by $\hat{\mu}(t) = (1/n) \sum_{i=1}^n \hat{x}_i(t)$ and

$$\hat{G}(s, t) = (1/n) \sum_{i=1}^n \{\hat{x}_i(s) - \hat{\mu}(s)\} \{\hat{x}_i(t) - \hat{\mu}(t)\}$$

respectively. The eigenvalues and eigenfunctions are estimated by solving the equation

$$\int_{\mathcal{T}} \hat{G}(s, t) \phi_k(s) \, ds = \lambda_k \phi_k(t)$$

for λ_k and $\phi_k(\cdot)$, subject to $\int_{\mathcal{T}} \phi_k^2(t) \, dt = 1$ and $\int_{\mathcal{T}} \phi_m(t) \phi_k(t) \, dt = 0$ for $m \neq k$, $k, m = 1, \dots, s$. The FPC scores are obtained by $\hat{\xi}_{ik} = \int_{\mathcal{T}} \{\hat{x}_i(t) - \hat{\mu}(t)\} \phi_k(t) \, dt$. Finally CDF transformation yields $\hat{\zeta}_{ik} = \Psi(\hat{\xi}_{ik}; 0, \hat{\lambda}_k)$.

- (b) *Obtain $\hat{\zeta}_i$ in the sparse design.* We adopt the principal component analysis through the PACE algorithm that was proposed by Yao *et al.* (2005), where the mean estimate $\hat{\mu}(t)$ is obtained by using local linear smoothers based on the pooled data of all individuals. In particular,

$$\hat{\mu}(t) = \sum_{i=1}^n \sum_{j=1}^{N_i} K\{(t_{ij} - t)/b\} \{x_{ij} - \beta_0 - \beta_1(t - t_{ij})\}^2$$

with $K(\cdot)$ a kernel function and b a bandwidth. For the covariance estimation, denote $G_{ijl} = \{x_{ij} - \hat{\mu}(t_{ij})\} \{x_{il} - \hat{\mu}(t_{il})\}$ and let $K_h^*(\cdot, \cdot)$ be a bivariate kernel function with a bandwidth h . One minimizes

$$\sum_{i=1}^n \sum_{j \neq l} K^*\{(t_{ij} - s)/h, (t_{il} - t)/h\} \{G_{ijl} - \beta_{00} - \beta_{11}(s - t_{ij}) - \beta_{12}(t - t_{il})\}^2.$$

One may estimate the noise variance ν^2 by taking the difference between the diagonal of the surface estimate $\hat{G}(t, t)$ and the local polynomial estimate obtained from the raw variances $\{(t_{ij}, G_{ijj}) : j = 1, \dots, N_i; i = 1, \dots, n\}$. The eigenvalues or eigenfunctions are obtained as in the dense case. To estimate the FPC scores, denote $\mathbf{x}_i = (x_{i1}, \dots, x_{iN_i})^T$, the PACE estimate is given by $\hat{\xi}_{ik} = \hat{\lambda}_k \hat{\phi}_{ik} \sum_{x_i}^{-1} (\mathbf{x}_i - \hat{\mu}_i)$, which leads to $\hat{\zeta}_{ik} = \Psi(\hat{\xi}_{ik}; 0, \hat{\lambda}_k)$, $k = 1, \dots, s$. Here $\hat{\phi}_{ik} = (\phi_k(t_{i1}), \dots, \phi_k(t_{iN_i}))^T$, $\mu_i = (\mu(t_{i1}), \dots, \mu(t_{iN_i}))^T$, and the (j, l) th element $(\Sigma_{x_i})_{j,l} = G(t_{ij}, t_{il}) + \nu^2 \delta_{jl}$ with $\delta_{jl} = 1$ if $j = l$ and $\delta_{jl} = 0$ otherwise, and ‘ \cdot ’ is generic notation for the estimated parameters.

We next estimate $f_0 \in \mathcal{F}^s$ by minimizing expression (6), following the COSSO procedure conditional on the estimated values $\hat{\zeta}_i$. It is important to note that the target function (6) is equivalent to

$$(1/n) \sum_{i=1}^n \{y_i - f(\hat{\zeta}_i)\}^2 + \lambda_0 \sum_{k=1}^s \theta_k^{-1} \|P^k f\|^2 + \lambda \sum_{k=1}^s \theta_k,$$

subject to $\theta_k \geq 0$ (Lin and Zhang, 2006), which enables a two-step iterative algorithm. Specifically, one first finds $\mathbf{c} \in \mathcal{R}^n$ and $b \in \mathcal{R}$ by minimizing

$$(\mathbf{y} - \mathbf{R}_\theta \mathbf{c} - b \mathbf{1}_n)^T (\mathbf{y} - \mathbf{R}_\theta \mathbf{c} - b \mathbf{1}_n) + n \lambda_0 \mathbf{c}^T \mathbf{R}_\theta \mathbf{c}, \quad (9)$$

with fixed $\theta = (\theta_1, \dots, \theta_s)^T$, where $\mathbf{y} = (y_1, \dots, y_n)^T$, λ_0 is the smoothing parameter, $\mathbf{1}_n$ is the $n \times 1$ vector of

is, $\mathbf{R}_\theta = \sum_{k=1}^s \theta_k \mathbf{R}_k$ and \mathbf{R}_k is the reproducing kernel of \bar{H}^k , i.e. $\mathbf{R}_k = \{R_k(\hat{\zeta}_{ik}, \hat{\zeta}_{jk})\}_{1 \leq i, j \leq n}$. This optimization is exactly a smoothing spline problem. We then fix \mathbf{c} and b , and find θ by minimizing

$$(\mathbf{z} - \mathbf{Q}\theta)^\top (\mathbf{z} - \mathbf{Q}\theta) \quad \text{subject to } \theta_k \geq 0; \quad \sum_{k=1}^s \theta_k \leq M, \quad (10)$$

where $\mathbf{z} = \mathbf{y} - \frac{1}{2}n\lambda_0\mathbf{c} - b\mathbf{1}_n$ and \mathbf{Q} is an $n \times s$ matrix with the k th column being $R_k\mathbf{c}$. This step is the same as calculating the non-negative garrotte estimate by using M as the tuning parameter. On convergence, the final estimation of f is then given by $\hat{f}(\zeta) = \sum_{i=1}^n \hat{c}_i R_\theta(\zeta_i, \zeta) + \hat{b}$.

Regarding the choice of tuning parameters, besides the sufficiently large initial truncation s , the most relevant are λ_0 and M in the COSSO step, whereas the bandwidths in the smoothing step of FPCA are chosen by traditional cross-validation or its generalized approximation. For more details, see Fan and Gijbels (1996) for the dense case and Yao *et al.* (2005) for the sparse case. We suggest selecting λ_0 by using generalized cross-validation, i.e. $\text{GCV}(\lambda_0) = (\hat{\mathbf{y}} - \mathbf{y})^\top (\hat{\mathbf{y}} - \mathbf{y}) / \{n^{-1} \text{tr}(I - A)\}^2$ with $\hat{\mathbf{y}} = A\mathbf{y}$. For choosing M , we adopt the Bayesian information criterion BIC, i.e. $\text{BIC}(M) = (\hat{\mathbf{y}} - \mathbf{y})^\top (\hat{\mathbf{y}} - \mathbf{y}) / \sigma^2 + \log(n) \text{df}$ where df is the degree of freedom in problem (10), whereas an alternative is cross-validation which requires more computation.

Appendix B: Technical assumptions and proofs

We first lay out the commonly adopted regularity conditions on the functional predictor process X for the dense design. Recall that $\{t_{ij}, j = 1, \dots, N_i; i = 1, \dots, n\}$ is the grid on the support \mathcal{T} over which the functional predictor $x_i(t)$ is observed. Without loss of generality, let $\mathcal{T} = [0, a]$. Denote $t_{i0} = 0, t_{iN_i} = a$ and $\mathcal{T}_d = [-d, a + d]$ for some $d > 0$. Denote the bandwidth that is used for individually smoothing the i th trajectory as b_i .

Condition 1. Assume that the second derivative $X^{(2)}(t)$ is continuous on \mathcal{T}_d with probability 1 and $\int E\{[X^{(k)}(t)]^4\} dt < \infty$ with probability 1 for $k = 0, 2$. Also assume that $E(e_{ij}^4) < \infty$, where e_{ij} is the IID measurement error of the observed trajectory x_i .

Condition 2. Assume that there exists $m \equiv m(n) \rightarrow \infty$, such that $\min_i N_i \geq m$ as $n \rightarrow \infty$. Denoting $\Delta_i = \max\{t_{ij} - t_{i,j-1} : j = 1, \dots, N_i + 1\}$, assume that $\max_i \Delta_i = O(m^{-1})$.

Condition 3. Assume that there is a sequence $b = b(n)$, such that $cb \leq \min_i b_i \leq \max_i b_i \leq Cb$ for some $C \geq c > 0$. Furthermore, $b \rightarrow 0$ and $m \rightarrow \infty$ as $n \rightarrow \infty$ in rates such that $(mb)^{-1} + b^4 + m^{-2} = O(n^{-1})$, e.g. $b = O(n^{-1/2})$ and $m = O(n^{3/2})$. Also assume that the kernel function $K(\cdot)$ is compactly supported and Lipschitz continuous.

Denote the operator that is associated with the covariance function $G(s, t)$ by G , and define $\|G\|_S^2 = \int_{\mathcal{T}} \int_{\mathcal{T}} G^2(s, t) ds dt$. Denote the smoothed trajectory of $X_i(t)$ by using local linear smoothing with bandwidth b_i by \hat{X}_i and the estimated eigenvalue, eigenfunction and FPC score in the dense design by $\hat{\lambda}_k, \hat{\phi}_k$ and $\hat{\xi}_{ik}$ respectively. Since the decay of eigenvalues plays an important role, define $\delta_1 = \lambda_1 - \lambda_2$ and $\delta_k = \min_{j \leq k} (\lambda_{j-1} - \lambda_j, \lambda_j - \lambda_{j+1})$ for $k \geq 2$.

Lemma 1. Under conditions 1–3 we have

$$E(\|\hat{X}_i - X_i\|_{L^2}^2) = O(n^{-1}), \quad \|\hat{\mu} - \mu\|_{L^2} = O_p(n^{-1/2}), \quad \|\hat{G} - G\|_S = O_p(n^{-1/2}), \quad (11)$$

$$|\hat{\lambda}_k - \lambda_k| \leq \|\hat{G} - G\|_S, \quad \|\hat{\phi}_k - \phi_k\|_{L^2} \leq 2\sqrt{2}\delta_k^{-1} \|\hat{G} - G\|_S, \quad (12)$$

$$|\hat{\xi}_{ik} - \xi_{ik}| = O_p(\|\hat{X}_i - X_i\|_{L^2} + \delta_k^{-1} \|X_i\|_{L^2} \|\hat{G} - G\|_S), \quad (13)$$

where $O(\cdot)$ and $O_p(\cdot)$ are uniform over $1 \leq i \leq n$.

Note that the measurement error e_{ij} is independent of the process X_i , which makes it possible to factor the probability space $\Omega = \Omega_X \times \Omega_e$ and to characterize the individual smoothing and cross-sectional averaging separately. Then equation (11) can be shown by using standard techniques with local polynomial smoothing (which are for conciseness not elaborated); see Hall *et al.* (2006) for more details of this type of argument. Consequently equations (12) and (13) follow immediately by the classical perturbation result provided in lemma 4.3 of Bosq (2000). We see from lemma 1 that, when the measurements are sufficiently dense for each subject satisfying condition 3, the effect due to individual smoothing on the estimated population quantities (e.g. mean, covariance, eigenvalues and eigenfunctions) are negligible.

The following lemma characterizes the discrepancy between the underlying and estimated transformed variables ζ_{ik} , as well as the boundedness of the derivative of the resulting estimate \hat{f} .

Lemma 2. Under assumption 2 in Section 2.2 and condition 1–3, we have

$$|\hat{\zeta}_{ik} - \zeta_{ik}| = O_p[\lambda_k^\gamma \{\|\hat{X}_i - X_i\|_{L^2} + (\delta_k^{-1}\|X_i\|_{L^2} + |\xi_{ik}|)\|\hat{G} - G\|_S\}], \quad (14)$$

$$\frac{1}{n} \sum_{i=1}^n \left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}| \right)^2 = O_p(n^{-1}). \quad (15)$$

Additionally, if assumption 1 holds, let \hat{f} be the estimate of f_0 obtained by minimizing expression (6). Then there is a constant $\rho > 0$, such that

$$\left| \frac{\partial \hat{f}(\zeta_i)}{\partial \zeta_{ik}} \right| \leq \rho, \quad (16)$$

uniformly over $1 \leq k \leq s$ and $1 \leq i \leq n$.

B.1. Proof of lemma 2

From lemma 1 and assumptions 2, we have

$$\begin{aligned} |\hat{\zeta}_{ik} - \zeta_{ik}| &= \left| (\hat{\xi}_{ik} - \xi_{ik}) \frac{\partial}{\partial \xi_{ik}} \Psi(\xi_{ik}, \lambda_k) + (\hat{\lambda}_k - \lambda_k) \frac{\partial}{\partial \lambda_k} \Psi(\xi_{ik}, \lambda_k) + o_p(|\hat{\xi}_{ik} - \xi_{ik}| + |\hat{\lambda}_k - \lambda_k|) \right| \\ &\leq |\hat{\xi}_{ik} - \xi_{ik}| \left| \frac{\partial}{\partial \xi_{ik}} \Psi(\xi_{ik}, \lambda_k) \right| + |\hat{\lambda}_k - \lambda_k| \left| \frac{\partial}{\partial \lambda_k} \Psi(\xi_{ik}, \lambda_k) \right| + o_p(|\hat{\xi}_{ik} - \xi_{ik}| + (|\hat{\lambda}_k - \lambda_k|)) \\ &= O_p[\lambda_k^\gamma \{\|\hat{X}_i - X_i\|_{L^2} + (\delta_k^{-1}\|X_i\|_{L^2} + |\xi_{ik}|)\|\hat{G} - G\|_S\}]. \end{aligned}$$

Abbreviate $\sum_{i=1}^n$ to Σ_i , $\sum_{k=1}^s$ to Σ_k and $O_p(\cdot)$ to ‘ \sim ’. Since $E\|\hat{X}_i - X_i\|_{L^2} \leq E(\|\hat{X}_i - X_i\|_{L^2}^2)^{1/2} = O(n^{-1/2})$, it is easy to see that $E(n^{-1} \sum_i \|\hat{X}_i - X_i\|_{L^2}) = E\|\hat{X}_i - X_i\|_{L^2} = O(n^{-1/2})$. To show result (15) for any fixed s , note that

$$n^{-1} \sum_i \left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}| \right)^2 \leq sn^{-1} \sum_i \sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|^2.$$

Then

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^s (\hat{\zeta}_{ik} - \zeta_{ik})^2 &\sim \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^s \lambda_k^{2\gamma} \{\|\hat{X}_i - X_i\|_{L^2} + (\delta_k^{-1}\|X_i\|_{L^2} + |\xi_{ik}|)\|\hat{G} - G\|_S\}^2 \\ &\sim \frac{1}{n} \sum_i \sum_k \lambda_k^{2\gamma} \|\hat{X}_i - X_i\|_{L^2}^2 + \frac{1}{n} \sum_i \sum_k \lambda_k^{2\gamma} \delta_k^{-2} \|X_i\|_{L^2}^2 \|\hat{G} - G\|_S^2 \\ &\quad + \frac{1}{n} \sum_i \sum_k \lambda_k^{2\gamma} |\xi_{ik}|^2 \|\hat{G} - G\|_S^2 + \frac{1}{n} \sum_i \sum_k \lambda_k^{2\gamma} \|\hat{X}_i - X_i\|_{L^2} \delta_k^{-1} \|X_i\|_{L^2} \|\hat{G} - G\|_S \\ &\quad + \frac{1}{n} \sum_i \sum_k \lambda_k^{2\gamma} \|\hat{X}_i - X_i\|_{L^2} |\xi_{ik}| \|\hat{G} - G\|_S + \frac{1}{n} \sum_i \sum_k \lambda_k^{2\gamma} \delta_k^{-1} |\xi_{ik}| \|X_i\|_{L^2} \|\hat{G} - G\|_S^2. \end{aligned}$$

Denoting the additive terms in this formula E_1 – E_6 , we have

$$\begin{aligned} E_1 &= \left(\sum_k \lambda_k^{2\gamma} \right) \left(n^{-1} \sum_i \|\hat{X}_i - X_i\|_{L^2}^2 \right) = O_p(n^{-1}), \\ E_2 &= \|\hat{G} - G\|_S^2 \left(\sum_k \lambda_k^{2\gamma} \delta_k^{-2} \right) \left(n^{-1} \sum_i \|X_i\|_{L^2}^2 \right) = O_p(n^{-1}), \\ E_3 &= \|\hat{G} - G\|_S^2 \left\{ (1/n) \sum_i \sum_k \lambda_k^{2\gamma} |\xi_{ik}|^2 \right\} = O_p(n^{-1}), \end{aligned}$$

as

$$E \left(n^{-1} \sum_{i=1}^n \sum_{k=1}^s \lambda_k^{2\gamma} |\xi_{ik}|^2 \right) = \sum_k \lambda_k^{2\gamma+1} = O(1).$$

For E_4 , applying the Cauchy–Schwarz inequality,

$$E_4 \sim \|\hat{G} - G\|_S \left(\sum_{k=1}^s \lambda_k^{2\gamma} \delta_k^{-1} \right) \left(\frac{1}{n} \sum_{i=1}^n \|\hat{X}_i - X_i\|_{L^2} \|X_i\|_{L^2} \right)$$

$$\begin{aligned} &\leq 2C\|\hat{G} - G\|_s \left(\sum_{k=1}^s \lambda_k^{2\gamma} \delta_k^{-1} \right) \sqrt{\left\{ \left(\frac{1}{n} \sum_{i=1}^n \|\hat{X}_i - X_i\|_{L^2}^2 \right) \left(\frac{1}{n} \sum_{i=1}^n \|X_i\|_{L^2}^2 \right) \right\}} \\ &= O_p(n^{-1/2}) O(1) O_p(n^{-1/2}) O_p(1) = O_p(n^{-1}). \end{aligned}$$

Similarly, we have $E_5 = O_p(n^{-1})$ and $E_6 = O_p(n^{-1})$, using the facts that $E\{(\sum_{k=1}^s \lambda_k^{2\gamma} |\xi_{ik}|)^2\} \leq s \sum_{k=1}^s \lambda_k^{4\gamma+1}$ $= O(1)$ and $E(\sum_{k=1}^s \lambda_k^{2\gamma} \delta_k^{-1} |\xi_{ik}|)^2 \leq s \sum_{k=1}^s \lambda_k^{4\gamma+1} \delta_k^{-2} = O(1)$. This proves result (15).

We now turn to inequality (16). For any $f \in \mathcal{F}^s$, we have

$$f(\zeta_i) = \langle f(\cdot), R(\zeta_i, \cdot) \rangle_{\mathcal{F}^s} \leq \|f\| \langle R(\zeta_i, \cdot), R(\zeta_i, \cdot) \rangle_{\mathcal{F}^s}^{1/2} = \|f\| R^{1/2}(\zeta_i, \zeta_i),$$

where $R(\cdot, \cdot)$ is the reproducing kernel of space \mathcal{F}^s and $\langle \cdot, \cdot \rangle_{\mathcal{F}^s}$ is the corresponding inner product. Therefore,

$$\frac{\partial f(\zeta_i)}{\partial \zeta_{ik}} = \left\langle f(\cdot), \frac{\partial R(\zeta_i, \cdot)}{\partial \zeta_{ik}} \right\rangle_{\mathcal{F}^s} \leq \|f\| \left\langle \frac{\partial R(\zeta_i, \cdot)}{\partial \zeta_{ik}}, \frac{\partial R(\zeta_i, \cdot)}{\partial \zeta_{ik}} \right\rangle_{\mathcal{F}^s}^{1/2}.$$

Since $J(f)$ is a convex functional and a pseudonorm, we have

$$\sum_{k=1}^s \|P^k f\|^2 \leq J^2(f) \leq s \sum_{k=1}^s \|P^k f\|^2. \quad (17)$$

We first claim that $\|f\| \leq J(f)$, because $\|f\|^2 = b^2 + \sum_{k=1}^s \|P^k f\|^2$. If $b=0$, inequality (17) implies that $\|f\| \leq J(f)$. If $b \neq 0$, we can write $\tilde{J}(f) = b + J(f) = b + \sum_{k=1}^s \|P^k f\|$. For minimizing expression (5), it is equivalent to substitute $J(f)$ with $\tilde{J}(f)$, and inequality (17) implies that $\|f\|^2 = b^2 + \sum_{k=1}^s \|P^k f\|^2 \leq b^2 + J^2(f) \leq \tilde{J}^2(f)$. Therefore we have $\|f\| \leq J(f)$ in general. Secondly, owing to the orthogonality of $\{\tilde{H}^k\}$, we can write $R(\mathbf{u}, \mathbf{v}) = R_1(u_1, v_1) + R_2(u_2, v_2) + \dots + R_s(u_s, v_s)$ by theorem 5 in Berlinet and Thomas-agnan (2004), where $R_k(\cdot, \cdot)$ is the reproducing kernel of the subspace \tilde{H}^k . For \tilde{H}^k being a second-order Sobolev Hilbert space, we have $R_k(s, t) = h_1(s)h_1(t) + h_2(s)h_2(t) - h_4(|s-t|)$, with $h_1(t) = t - \frac{1}{2}$, $h_2(t) = \{h_1^2(t) - 1/12\}/2$ and $h_4(t) = \{h_1^4(t) - h_1^2(t)/2 + 7/240\}/24$. Therefore $R_k(s, t)$ is continuous and differentiable over $[0, 1]^2$ and we can find constants a_k and b_k such that

$$\langle R_k(u, \cdot), R_k(u, \cdot) \rangle_{\mathcal{F}^s} < a_k,$$

$$\left\langle \frac{\partial R_k(u, \cdot)}{\partial u}, \frac{\partial R_k(u, \cdot)}{\partial u} \right\rangle_{\mathcal{F}^s} \leq b_k,$$

for $k=1, \dots, s$. One can find a uniform bound c with $\langle \partial R(\zeta_i, \cdot) / \partial \zeta_{ik}, \partial R(\zeta_i, \cdot) / \partial \zeta_{ik} \rangle_{\mathcal{F}^s} \leq c$. However, an \hat{f} minimizing expression (6) is equivalent to minimizing $n^{-1} \sum_i \{y_i - f(\zeta_i)\}^2$ under the constraint that $J(f) \leq \tilde{c}$ for some $\tilde{c} > 0$. Therefore let $\rho = c^{1/2} \tilde{c}$; we have

$$\left| \frac{\partial \hat{f}(\zeta_i)}{\partial \zeta_{ik}} \right| \leq \|\hat{f}\| \left\langle \frac{\partial R(\zeta_i, \cdot)}{\partial \zeta_{ik}}, \frac{\partial R(\zeta_i, \cdot)}{\partial \zeta_{ik}} \right\rangle_{\mathcal{F}^s}^{1/2} \leq J(\hat{f}) c^{1/2} \leq \tilde{c} c^{1/2} = \rho. \quad \square$$

Before stating lemma 3, we define the entropy of \mathcal{F}^s with respect to the $\|\cdot\|_n$ metric. For each $\omega > 0$, one can find a collection of functions $\{g_1, g_2, \dots, g_N\}$ in \mathcal{F}^s such that, for each $g \in \mathcal{F}^s$, there is a $j = j(g) \in \{1, 2, \dots, N\}$ satisfying $\|g - g_j\|_n \leq \omega$. Let $\mathbb{N}(\omega, \mathcal{F}^s, \|\cdot\|_n)$ be the smallest value of N for which such a cover of balls with radius ω and centres g_1, g_2, \dots, g_N exists. Then $H(\omega, \mathcal{F}^s, \|\cdot\|_n) = \log\{\mathbb{N}(\omega, \mathcal{F}^s, \|\cdot\|_n)\}$ is called the ω -entropy of \mathcal{F}^s .

Lemma 3. Assume that $\mathcal{F}^s = \{1\} \oplus \sum_{k=1}^s \tilde{H}^k$, where \tilde{H}^k is the second-order Sobolev space. Denote the ω -entropy of $\{f \in \mathcal{F}^s : J(f) \leq 1\}$ by $H(\omega, \{f \in \mathcal{F}^s : J(f) \leq 1\}, \|\cdot\|_n)$. Then

$$H(\omega, \{f \in \mathcal{F}^s : J(f) \leq 1\}, \|\cdot\|_n) \leq A\omega^{-1/2}, \quad (18)$$

for all $\omega > 0$, $n \geq 1$, and for some constants $A > 0$. Furthermore, for $\{\varepsilon_i\}_{i=1}^n$ independent with finite variance and $J(f_0) > 0$,

$$\sup_{f \in \mathcal{F}^s} \frac{|(\varepsilon, f - f_0)_n|}{\|f - f_0\|_n^{3/4} \{J(f) + J(f_0)\}^{1/4}} = O_p(n^{-1/2}). \quad (19)$$

Inequality (18) is implied by lemma A.1 of Lin and Zhang (2006). As the $\{\varepsilon_i\}$ satisfy the sub-Gaussian error assumption, the same argument as in Van de Geer (2000) (page 168) leads to result (19). We are now ready to present the proof of the main theorem.

B.2. Proof of theorem 1

We first centre the functions as in the proof of theorem 2 in Lin and Zhang (2006) so that results (18) and (19) hold. Write $f(\zeta) = c + f_1(\zeta_1) + \dots + f_s(\zeta_s) = c + \tilde{f}(\zeta)$, such that $\sum_{i=1}^n f_k(\zeta_{i,k}) = 0$, and write $f_0(\zeta) = c_0 + f_{01}(\zeta_1) + \dots + f_{0s}(\zeta_s) = c_0 + \tilde{f}_0(\zeta)$ such that $\sum_{i=1}^n f_{0k}(\zeta_{i,k}) = 0$ and $\hat{f}(\zeta) = \hat{c} + \hat{f}_1(\zeta_1) + \dots + \hat{f}_s(\zeta_s)$. Since the target function can be written as

$$\begin{aligned} Q(f|\hat{\zeta}_i) &= \frac{1}{n} \sum_{i=1}^n \{y_i - f(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f) = \frac{1}{n} \sum_{i=1}^n \{c_0 + \tilde{f}_0(\zeta_i) + \varepsilon_i - c - \tilde{f}(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f) \\ &= (c_0 - c)^2 + \frac{2}{n} (c_0 - c) \sum_i \varepsilon_i + \frac{1}{n} \sum_{i=1}^n \{\tilde{f}_0(\zeta_i) + \varepsilon_i - \tilde{f}(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f), \end{aligned}$$

we must have that \hat{c} minimize $\{(c_0 - c)^2 + 2n^{-1}(c_0 - c)\sum_i \varepsilon_i\}$ and the additive parts of \hat{f} minimize the rest. Therefore we have $\hat{c} - c_0 = n^{-1}\sum_i \varepsilon_i$, implying that $|\hat{c} - c_0| = O_p(n^{-1/2})$. Denote

$$\tilde{Q}(\tilde{f}|\hat{\zeta}_i) = \frac{1}{n} \sum_{i=1}^n \{\tilde{f}_0(\zeta_i) + \varepsilon_i - \tilde{f}(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f). \quad (20)$$

We can substitute $\tau_n^2 J(f)$ with $\tau_n^2 J(\tilde{f})$ in equation (20). In the rest of the proof, we suppress the tilde notation of \tilde{f}_0 and \tilde{f} for convenience. Since $\hat{f} = \arg \min_{f \in \mathcal{F}^s} \tilde{Q}(f|\{\hat{\zeta}_i\})$, we have $\tilde{Q}(\hat{f}|\{\hat{\zeta}_i\}) \leq \tilde{Q}(f_0|\{\hat{\zeta}_i\})$, which implies that

$$\frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) + \varepsilon_i - \hat{f}(\hat{\zeta}_i)\}^2 + \tau_n^2 J(\hat{f}) \leq \frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) + \varepsilon_i - f_0(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f_0).$$

Simplification of this inequality gives

$$\frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) - \hat{f}(\hat{\zeta}_i)\}^2 + \tau_n^2 J(\hat{f}) \leq \frac{2}{n} \sum_{i=1}^n \varepsilon_i \{\hat{f}(\hat{\zeta}_i) - f_0(\hat{\zeta}_i)\} + \frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) - f_0(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f_0). \quad (21)$$

Let $g(\cdot) = \hat{f}(\cdot) - f_0(\cdot)$. Since both \hat{f} and f_0 are in \mathcal{F}^s , $g \in \mathcal{F}^s$. Taylor series expansion of $g(\cdot)$ gives $g(\hat{\zeta}) = g(\zeta) + Dg(\zeta)(\hat{\zeta} - \zeta) + o_p(\sum_{k=1}^s |\hat{\zeta}_k - \zeta_k|)$, for all $\zeta \in (0, 1)^s$, where $Dg(\zeta)(\hat{\zeta} - \zeta) = \sum_{k=1}^s \{\partial g(\zeta)/\partial \zeta_k\}(\hat{\zeta}_k - \zeta_k)$. Then we have

$$\frac{2}{n} \sum_{i=1}^n \varepsilon_i g(\hat{\zeta}_i) = \frac{2}{n} \sum_{i=1}^n \varepsilon_i g(\zeta_i) + \frac{2}{n} \sum_{i=1}^n \varepsilon_i \left\{ Dg(\zeta_i)(\hat{\zeta}_i - \zeta_i) + o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \right\},$$

and we plug it into the right-hand side of inequality (21), leading to the upper bound

$$\begin{aligned} \frac{2}{n} \sum_{i=1}^n \varepsilon_i \{\hat{f}(\hat{\zeta}_i) - f_0(\zeta_i)\} + \frac{2}{n} \sum_{i=1}^n \varepsilon_i \left[\{D\hat{f}(\zeta_i) - Df_0(\zeta_i)\}(\hat{\zeta}_i - \zeta_i) + o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \right] \\ + \frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) - f_0(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f_0). \end{aligned} \quad (22)$$

Applying lemma 3, we can bound the first term in expression (22) as follows:

$$\frac{2}{n} \sum_{i=1}^n \varepsilon_i \{\hat{f}(\hat{\zeta}_i) - f_0(\zeta_i)\} = 2(\varepsilon, \hat{f} - f_0)_n \leq O_p(n^{-1/2}) \|\hat{f} - f_0\|_n^{3/4} \{J(\hat{f}) + J(f_0)\}^{1/4}.$$

For the left-hand side of inequality (21), applying the Taylor series expansion, $\hat{f}(\hat{\zeta}_i) = \hat{f}(\zeta_i) + D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i) + o_p(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|)$, to the first term

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) - \hat{f}(\hat{\zeta}_i)\}^2 &= \frac{1}{n} \sum_{i=1}^n \left\{ f_0(\zeta_i) - \hat{f}(\zeta_i) - D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i) - o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \right\}^2 \\ &= \frac{1}{n} \sum_{i=1}^n [\{f_0(\zeta_i) - \hat{f}(\zeta_i)\}^2 + \{D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i)\}^2 - 2\{f_0(\zeta_i) - \hat{f}(\zeta_i)\}D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i) + R_i], \end{aligned}$$

where

$$R_i = \left\{ o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \right\}^2 - o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \{f_0(\zeta_i) - \hat{f}(\zeta_i) - D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i)\}.$$

Substituting the terms on both sides of inequality (21), we obtain

$$\begin{aligned} \|\hat{f} - f_0\|_n^2 + \frac{1}{n} \sum_{i=1}^n [\{D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i)\}^2 + 2\{\hat{f}(\zeta_i) - f_0(\zeta_i)\}D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i) + R_i] + \tau_n^2 J(\hat{f}) \\ \leq O_p(n^{-1/2})\|\hat{f} - f_0\|_n^{3/4}\{J(\hat{f}) + J(f_0)\}^{1/4} + \frac{1}{n} \sum_{i=1}^n \{f_0(\zeta_i) - f_0(\hat{\zeta}_i)\}^2 + \tau_n^2 J(f_0) \\ + \frac{2}{n} \sum_{i=1}^n \varepsilon_i \left[\{D\hat{f}(\zeta_i) - Df_0(\zeta_i)\}(\hat{\zeta}_i - \zeta_i) + o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \right]. \end{aligned}$$

Dropping the positive term $n^{-1} \sum_i \{D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i)\}^2$ on the left-hand side and rearranging the terms,

$$\begin{aligned} \|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq O_p(n^{-1/2})\|\hat{f} - f_0\|_n^{3/4}\{J(\hat{f}) + J(f_0)\}^{1/4} + \tau_n^2 J(f_0) + T_1 + T_2 + T_3 + \frac{2}{n} \sum_{i=1}^n \varepsilon_i \tilde{R}_{2i} \\ + \frac{1}{n} \sum_{i=1}^n \tilde{R}_{1i} \end{aligned} \quad (23)$$

where

$$T_1 = -2n^{-1} \sum_i \{\hat{f}(\zeta_i) - f_0(\zeta_i)\}D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i),$$

$$T_2 = 2n^{-1} \sum_i \varepsilon_i [\{D\hat{f}(\zeta_i) - Df_0(\zeta_i)\}(\hat{\zeta}_i - \zeta_i)],$$

$$T_3 = n^{-1} \sum_i \{f_0(\zeta_i) - f_0(\hat{\zeta}_i)\}^2,$$

$$\tilde{R}_{1i} = o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \{f_0(\zeta_i) - \hat{f}(\zeta_i) - D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i)\}$$

and $\tilde{R}_{2i} = o_p(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|)$.

For T_1 , by the Cauchy–Schwarz inequality and lemma 2, we have that $T_1 \leq 2\sqrt{(\|\hat{f} - f_0\|_n^2 A)}$, where

$$A = \frac{1}{n} \sum_{i=1}^n \left\{ \sum_{k=1}^s \frac{\partial \hat{f}(\zeta_i)}{\partial \zeta_{ik}} (\hat{\zeta}_{ik} - \zeta_{ik}) \right\}^2 \leq \frac{\rho^2}{n} \sum_{i=1}^n \left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}| \right)^2 = O_p(n^{-1}),$$

i.e. $T_1 = \|\hat{f} - f_0\|_n O_p(n^{-1/2})$. From assumption 1 and result (16) of lemma 2, there are independent random variables $\{B_i\}$ with $E(B_i^2) < \infty$ such that $\max_k \{|\partial \hat{f}(\zeta_i)/\partial \zeta_{ik} - \partial f_0(\zeta_i)/\partial \zeta_{ik}|\} \leq B_i \|\hat{f} - f_0\|_{L^2}$. Also note that $\|g\|_n \rightarrow \|g\|_{L^2}$ almost surely by the strong law of large numbers. Therefore we have, for some constant c ,

$$\begin{aligned} T_2 \leq \frac{2}{n} \sum_{i=1}^n |\varepsilon_i| \sum_{k=1}^s B_i \|\hat{f} - f_0\|_{L^2} |\hat{\zeta}_{ik} - \zeta_{ik}| = 2\|\hat{f} - f_0\|_{L^2} \left(\frac{1}{n} \sum_{i=1}^n |\varepsilon_i B_i| \sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}| \right) \\ \leq c\|\hat{f} - f_0\|_n \sqrt{\left[\left(\frac{1}{n} \sum_{i=1}^n \varepsilon_i^2 B_i^2 \right) \left\{ \frac{1}{n} \sum_{i=1}^n \left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}| \right)^2 \right\} \right]} = \|\hat{f} - f_0\|_n O_p(n^{-1/2}), \end{aligned}$$

$$T_3 = \frac{1}{n} \sum_{i=1}^n \left[\sum_{k=1}^s \frac{\partial f_0(\zeta_i)}{\partial \zeta_{ik}} (\hat{\zeta}_{ik} - \zeta_{ik}) + o_p(|\hat{\zeta}_{ik} - \zeta_{ik}|) \right]^2 \leq \frac{c}{n} \sum_{i=1}^n \left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}| \right)^2 = O_p(n^{-1}).$$

For the remaining terms, $n^{-1} \sum_{i=1}^n \varepsilon_i \tilde{R}_{2i} = o_p(T_2)$, and

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \tilde{R}_{1i} = \frac{1}{n} \sum_{i=1}^n o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \{f_0(\zeta_i) - \hat{f}(\zeta_i)\} - \frac{1}{n} \sum_{i=1}^n o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i) \\ \leq o_p(T_1) + \left[\frac{1}{n} \sum_{i=1}^n \left\{ o_p\left(\sum_{k=1}^s |\hat{\zeta}_{ik} - \zeta_{ik}|\right) \right\}^2 \frac{1}{n} \sum_{i=1}^n \{D\hat{f}(\zeta_i)(\hat{\zeta}_i - \zeta_i)\}^2 \right]^{-1/2} = o_p(T_1) + o_p(n^{-1}). \end{aligned}$$

We can now simplify inequality (23) as follows:

$$\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq O_p(n^{-1/2})\|\hat{f} - f_0\|_n^{3/4}\{J(\hat{f}) + J(f_0)\}^{1/4} + \|\hat{f} - f_0\|_n O_p(n^{-1/2}) + O_p(n^{-1}) + \tau_n^2 J(f_0).$$

If $O_p(n^{-1/2})\|\hat{f} - f_0\|_n^{3/4}\{J(\hat{f}) + J(f_0)\}^{1/4} \geq \|\hat{f} - f_0\|_n O_p(n^{-1/2}) + O_p(n^{-1}) + \tau_n^2 J(f_0)$, we have

$$\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq O_p(n^{-1/2})\|\hat{f} - f_0\|_n^{3/4}\{J(\hat{f}) + J(f_0)\}^{1/4}; \quad (24)$$

otherwise,

$$\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq \|\hat{f} - f_0\|_n O_p(n^{-1/2}) + O_p(n^{-1}) + 2\tau_n^2 J(f_0). \quad (25)$$

The proof will be completed by solving them separately. For the case of inequality (24), there are two possibilities.

(a) If $J(\hat{f}) \geq J(f_0)$, inequality (24) implies that $\tau_n^2 J^{3/4}(\hat{f}) \leq O_p(n^{-1/2}) \|\hat{f} - f_0\|_n^{3/4}$, and

$$J^{1/4}(\hat{f}) \leq \{\tau_n^{-2} O_p(n^{-1/2}) \|\hat{f} - f_0\|_n^{3/4}\}^{1/3} = O_p(n^{-1/6}) \|\hat{f} - f_0\|_n^{1/4} \tau_n^{-2/3}.$$

Therefore,

$$\|\hat{f} - f_0\|_n^2 \leq O_p(n^{-1/2}) \|\hat{f} - f_0\|_n^{3/4} J^{1/4}(\hat{f}) \leq O_p(n^{-2/3}) \|\hat{f} - f_0\|_n \tau_n^{-2/3},$$

i.e.

$$\begin{aligned} \|\hat{f} - f_0\|_n &= O_p(n^{-2/3}) O_p(\tau_n^{-2/3}), \\ J(\hat{f}) &= O_p(n^{-4/3}) O_p(\tau_n^{-10/3}). \end{aligned} \quad (26)$$

(b) If $J(\hat{f}) < J(f_0)$, then $J(\hat{f}) = O_p\{J(f_0)\} O_p(1)$, and inequality (24) implies that

$$\|\hat{f} - f_0\|_n^2 \leq O_p(n^{-1/2}) \|\hat{f} - f_0\|_n^{3/4} J^{1/4}(f_0),$$

which leads to

$$\begin{aligned} \|\hat{f} - f_0\|_n &= O_p(n^{-2/5}) J^{1/5}(f_0), \\ J(\hat{f}) &= J(f_0) O_p(1). \end{aligned} \quad (27)$$

Note that results (26) and (27) are equivalent under condition (7).

For the case of inequality (25), if $\|\hat{f} - f_0\|_n O_p(n^{-1/2}) > O_p(n^{-1}) + 2\tau_n^2 J(f_0)$, we have $\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq \|\hat{f} - f_0\|_n O_p(n^{-1/2})$; otherwise $\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq O_p(n^{-1}) + 4\tau_n^2 J(f_0)$. The first inequality implies that

$$\begin{aligned} \|\hat{f} - f_0\|_n &= O_p(n^{-1/2}), \\ J(\hat{f}) &= O_p(n^{-1}) O_p(\tau_n^{-2}). \end{aligned} \quad (28)$$

For the second inequality, if $O_p(n^{-1}) < 4\tau_n^2 J(f_0)$, we have $\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq 8\tau_n^2 J(f_0)$, implying that

$$\begin{aligned} \|\hat{f} - f_0\|_n &= O_p(\tau_n) J^{1/2}(f_0), \\ J(\hat{f}) &= J(f_0) O_p(1). \end{aligned} \quad (29)$$

If $O_p(n^{-1}) \geq 4\tau_n^2 J(f_0)$ and $\|\hat{f} - f_0\|_n^2 + \tau_n^2 J(\hat{f}) \leq O_p(n^{-1})$, then

$$\begin{aligned} \|\hat{f} - f_0\|_n &= O_p(n^{-1/2}), \\ J(\hat{f}) &= O_p(n^{-1}) O_p(\tau_n^{-2}). \end{aligned} \quad (30)$$

When $J(f_0) > 0$, given condition (7), the rates of $\|\hat{f} - f_0\|_n$ and $J(\hat{f})$ from expressions (29), (26) and (27) are the same, and dominate those of expressions (28) and (30). Therefore we have $\|\hat{f} - f_0\|_n = O_p(n^{-2/5}) J^{1/5}(f_0)$ and $J(\hat{f}) = J(f_0) O_p(1)$. When $J(f_0) = 0$, then inequality (24) implies expression (26), whereas inequality (25) implies expressions (28) and (30). The possibility (b) of inequality (24) does not exist; nor does the result in expression (29). Under condition (8), the result of expression (26) is the same as those of expressions (28) and (30). Therefore $\|\hat{f} - f_0\|_n = O_p(n^{-1/2})$ and $J(\hat{f}) = O_p(n^{-1/2})$.

References

- Bair, E., Hastie, T., Paul, D. and Tibshirani, R. (2006) Prediction by supervised principal components. *J. Am. Statist. Ass.*, **101**, 119–137.
- Berlinet, A. and Thomas-agnan, C. (2004) *Reproducing Kernel Hilbert Space in Probability and Statistics*. Norwell: Kluwer.
- Bosq, D. (2000) *Linear Processes in Function Spaces: Theory and Applications*. New York: Springer.

- Cardot, H., Ferraty, F., Mas, A. and Sarda, P. (2003) Testing hypotheses in the functional linear model. *Scand. J. Statist.*, **30**, 241–255.
- Cardot, H., Ferraty, F. and Sarda, P. (1999) Functional linear model. *Statist. Probab. Lett.*, **45**, 11–22.
- Fan, J. and Gijbels, I. (1996) *Local Polynomial Modelling and Its Applications*. London: Chapman and Hall.
- Fan, J. and Li, R. (2001) Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Statist. Ass.*, **96**, 1348–1360.
- Fan, J. and Zhang, J. (2000) Two-step estimation of functional linear models with applications to longitudinal data. *J. R. Statist. Soc. B*, **62**, 303–322.
- Goldsmith, J., Bobb, J., Crainiceanu, C. M., Caffo, B. and Reich, D. (2011) Penalized functional regression. *J. Computnl Graph. Statist.*, **20**, 830–851.
- Hadi, A. S. and Ling, R. F. (1998) Some cautionary notes on the use of principal components regression. *Am. Statistn*, **52**, 15–19.
- Hall, P. and Hosseini-Nasab, M. (2006) On properties of functional principal components analysis. *J. R. Statist. Soc. B*, **68**, 109–126.
- Hall, P., Müller, H. and Wang, J. (2006) Properties of principle component methods for functional and longitudinal data analysis. *Ann. Statist.*, **34**, 1493–1517.
- Hastie, T. and Tibshirani, R. J. (1990) *Generalized Additive Models*. London: Chapman and Hall.
- Hastie, T., Tibshirani, R. and Friedman, J. (2001) *The Elements of Statistical Learning*. New York: Springer.
- James, G. M. (2002) Generalized linear models with functional predictors. *J. R. Statist. Soc. B*, **64**, 411–432.
- James, G. M., Wang, J. and Zhu, J. (2009) Functional linear regression that's interpretable. *Ann. Statist.*, **37**, 2083–2108.
- Li, Y., Wang, N. and Carroll, R. (2010) Generalized functional linear models with semi-parametric single-index interactions. *J. Am. Statist. Ass.*, **105**, 621–633.
- Lin, Y. and Zhang, H. (2006) Component selection and smoothing in multivariate nonparametric regression. *Ann. Statist.*, **34**, 2272–2297.
- Marra, G. and Wood, S. N. (2011) Practical variable selection for generalized additive models. *J. Computnl Graph. Statist.*, **55**, 2372–2387.
- Meier, L., Van de Geer, S. and Bühlmann, P. (2009) High-dimensional additive modeling. *Ann. Statist.*, **37**, 3779–3821.
- Müller, H. and Stadtmüller, U. (2005) Generalized functional linear models. *Ann. Statist.*, **33**, 774–805.
- Müller, H. and Yao, F. (2008) Functional additive models. *J. Am. Statist. Ass.*, **103**, 1534–1544.
- Ramsay, J. O. and Silverman, B. W. (2005) *Functional Data Analysis*, 2nd edn. New York: Springer.
- Ravikumar, P., Lafferty, J., Liu, H. and Wasserman, L. (2009) Sparse additive models. *J. R. Statist. Soc. B*, **71**, 1009–1030.
- Rice, J. A. and Silverman, B. W. (1991) Estimating the mean and covariance structure nonparametrically when the data are curves. *J. R. Statist. Soc. B*, **53**, 233–243.
- Tibshirani, R. (1996) Regression shrinkage and selection via the lasso. *J. R. Statist. Soc. B*, **58**, 267–288.
- Van de Geer, S. (2000) *Empirical Processes in M-estimation*. Cambridge: Cambridge University Press.
- Wahba, G. (1990) *Spline Models for Observational Data*. Philadelphia: Society for Industrial and Applied Mathematics.
- Wood, S. N. (2006) *Generalized Additive Models: an Introduction with R*. New York: Chapman and Hall.
- Xu, L., Jiang, J., Wu, H., Shen, G. and Yu, R. (2007) Variable-weighted PLS. *Chemometr. Intell. Lab. Syst.*, **85**, 140–143.
- Yao, F. (2007) Asymptotic distributions of nonparametric regression estimators for longitudinal or functional data. *J. Multiv. Anal.*, **98**, 40–56.
- Yao, F. and Müller, H. G. (2010) Functional quadratic regression. *Biometrika*, **97**, 49–64.
- Yao, F., Müller, H. G. and Wang, J. L. (2005) Functional data analysis for sparse longitudinal data. *J. Am. Statist. Ass.*, **100**, 577–590.
- Zhu, H., Vannucci, M. and Cox, D. D. (2007) Functional data classification in cervical pre-cancer diagnosis—a bayesian variable selection model. In *Proc. Jt Statist. Meet.*
- Zhu, H., Vannucci, M. and Cox, D. D. (2010) A Bayesian hierarchical model for classification with selection of functional predictors. *Biometrics*, **66**, 463–473.
- Zipunnikov, V., Caffo, B., Yousem, D. M., Davatzikos, C., Schwartz, B. S. and Crainiceanu, C. (2011) Functional principal component model for high-dimensional brain imaging. *NeuroImage*, **58**, 772–784.
- Zou, H. (2006) The adaptive lasso and its oracle properties. *J. Am. Statist. Ass.*, **101**, 1418–1429.