

Identifiability in penalized function-on-function regression models

Fabian Scheipl and Sonja Greven

*Institut für Statistik
Ludwig-Maximilians-Universität München
Ludwigstrasse 33
Munich, Germany*

e-mail: fabian.scheipl@stat.uni-muenchen.de; sonja.greven@stat.uni-muenchen.de

Abstract: Regression models with functional responses and covariates constitute a powerful and increasingly important model class. However, regression with functional data poses well known and challenging problems of non-identifiability. This non-identifiability can manifest itself in arbitrarily large errors for coefficient surface estimates despite accurate predictions of the responses, thus invalidating substantial interpretations of the fitted models. We offer an accessible rephrasing of these identifiability issues in realistic applications of penalized linear function-on-function-regression and delimit the set of circumstances under which they are likely to occur in practice. Specifically, non-identifiability that persists under smoothness assumptions on the coefficient surface can occur if the functional covariate’s empirical covariance has a kernel which overlaps that of the roughness penalty of the spline estimator. Extensive simulation studies validate the theoretical insights, explore the extent of the problem and allow us to evaluate their practical consequences under varying assumptions about the data generating processes. A case study illustrates the practical significance of the problem. Based on theoretical considerations and our empirical evaluation, we provide immediately applicable diagnostics for lack of identifiability and give recommendations for avoiding estimation artifacts in practice.

MSC 2010 subject classifications: Primary 62J07; secondary 62J20.

Keywords and phrases: Functional data, penalized splines.

Received June 2015.

1. Introduction

The last two decades have seen rapid progress in regression methodology for high-dimensional data, largely driven by applications to genomic data in the “small n , large p ” paradigm. In regression models with functional predictors, similar problems arise from the fact that covariate information comes in the shape of high-dimensional, strongly auto-correlated vectors of function evaluations. Whenever the number of regression parameters to estimate exceeds the number of observations, estimates are not unique and the resulting model is not identifiable. To overcome this lack of identifiability, it then becomes necessary to use heuristics or prior knowledge to impose additional structural constraints like sparsity or smoothness. Results then inherently depend – at least to some

degree – on the assumptions underlying the chosen regularization method. In the following, we present a detailed analysis of the way in which smoothness assumptions combined with properties of the data generating process affect estimation results for function-on-function regression. We differentiate between two kinds of non-identifiability: First, *simple non-identifiability* arising from low information content in the functional data. This can be cured by imposing structural assumptions like smoothness or sparsity on the estimators, i.e., by regularization of the estimators. Fig. 1 shows an example for 4 estimates under different structural assumptions all yielding identical fits in such a scenario. Second, *persistent non-identifiability* that remains despite regularization for certain combinations of models and data. Figures 4 and 8 show examples for the latter on synthetic and real data, respectively.

The problem of – especially *persistent* – non-identifiability is as yet underappreciated in the functional data literature and analyzed here in depth for the first time. As software capable of fitting increasingly complex models with functional data becomes available (e.g. `fda`, Ramsay et al. (2014); `fda.usc`, Febrero-Bande and Oviedo de la Fuente (2012); `refund`, Huang et al. (2015); `PACE`, Yang et al. (2012); `WFMM`, Herrick (2013)), investigating the practical relevance of identifiability issues arising in these models is both timely and important in this rapidly developing field. The present work aims to perform such an investigation for the model class described in Scheipl et al. (2015) and implemented in `refund`'s `pffr` function, while results carry over to other penalized function-on-function regression approaches such as those implemented in the `fda` package.

A popular approach in regression for functional data restricts the functional coefficients to lie in the span of the first $K < n$ estimated eigenfunctions of a functional covariate's covariance operator with the largest eigenvalues, (see Cardot et al., 1999, 2003; Yao et al., 2005; Reiss and Ogden, 2007; Yao and Müller, 2010; Wu et al., 2010, for example). This functional principal component regression (FPCR) approach solves the problem of overparameterization (i.e., non-identifiability of the functional effect) by a – usually drastic – dimension reduction. The main challenges in this approach then become 1) achieving good estimates of the covariance's eigenfunctions (“functional principal components” or FPCs), eigenvalues, and FPC scores from observed functional data and 2) choosing the regularization parameter K . In practice, the effect of the functional covariate is estimated by using the first K estimated FPC scores as synthetic covariates. However, the critical assumption that the true coefficient lies in the span of the first few empirical eigenfunctions of a suitable (cross-)covariance operator estimate is impossible to verify empirically. Due to the often wiggly and unsmooth nature of eigenfunctions of real data this assumption can also lead to estimates that are difficult to interpret or implausible to practitioners.

An alternative approach is to make assumptions on the functional coefficients informed by insights into the problem at hand, e.g., sparsity or smoothness of functional coefficients, and to estimate these functional coefficients subject to an appropriate penalty (e.g., LASSO or smoothness penalties). In this work, we will focus on smooth spline-based penalized regression models for functional responses with functional covariates as described in Ivanescu et al. (2015) and

Scheipl et al. (2015), which constitute a powerful and flexible model class able to deal with the wealth of functional data increasingly collected in many fields of science. Nevertheless, our considerations carry over to other approaches to estimate smooth coefficient functions, such as approaches using derivative-based penalties as advocated by Ramsay and Silverman (2005). This paper describes the data settings in which penalized models remain unidentifiable despite the penalty in Section 3 and develops and evaluates suitable diagnostics and modified penalties for such settings.

Identifiability issues in functional regression have previously been discussed in Cardot et al. (2003) in the context of functional regression models for scalar responses and also, briefly, in the context of models with both functional responses and functional predictors by He et al. (2000), Chiou et al. (2004) and Prchal and Sarda (2007). While results therein provide conditions for the theoretical existence and unicity of solutions based on functional analysis arguments, they do not yield empirically verifiable criteria to determine whether the conditions for unicity are violated for a given data set. They also always assume that the true coefficient surface lies in the space spanned by the eigenfunctions of a (cross-)covariance operator. As far as we are aware, case studies in the previous literature have implicitly assumed that this assumption and the necessary conditions based on it will be satisfied for observed data. This is problematic since 1) the theoretical conditions found in the previous literature are impossible to satisfy, or at least verify, on finite samples of functional data in finite resolution, and 2) our experience with applications of functional regression models as well as results from simulation studies indicate that persistent non-identifiability leading to spurious coefficient estimates may occur quite regularly. This is obviously a concern for applied statisticians desiring *interpretable* regression models associating functional covariates and (functional) responses.

Instead of relying on the functional analysis arguments suitable for investigations of asymptotic properties of the theoretical model, we use simple linear algebra to derive a condition for unicity of coefficient surface estimates in realistic, finite sample data settings in which functional covariates are observed with finite resolution in Section 3. This allows us to give a necessary and sufficient condition for persistent non-identifiability in penalized function-on-function regression that is empirically verifiable and thus applicable in realistic problems. The criterion is based on the amount of overlap between the kernel of the penalty matrix and the kernel of the design matrix for the functional effect. Simulation studies indicate that, in practice, severe errors due to non-identifiability are strongly associated with our criterion; thus, this criterion is the first one that can be used in a wide variety of applications to assess identifiability or lack thereof. Our analyses also indicate that many widely used preprocessing techniques for functional data which replace observed curves with spline-based or FPC-based low-rank approximations (i.e., pre-smoothing) or the centering of individual observed curves will considerably increase the likelihood of identifiability issues in many settings.

We use the well-known DTI data set – publicly available in R-package `refund` (Huang et al., 2015) – to illustrate the issues we discuss here. These data contain

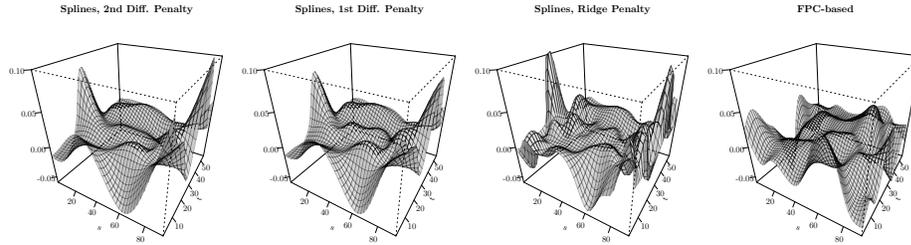


FIG 1. Four coefficient surfaces for regressing RCST-FA on pre-smoothed CCA-FA truncated to its first 6 empirical FPCs. All lead to identical fitted values. First [second] panel: Tensor product spline-based fits (12 marginal cubic B-spline basis functions) with second [first] order difference penalty. Third panel: Tensor product spline-based fit with a Ridge penalty. Fourth panel: FPC regression result (6 FPCs). All fits performed with `pfpr()`.

spatially indexed, i.e., *functional*, measurements of fractional anisotropy (FA), a proxy variable for neuronal health, along 3 cerebral white matter tracts (WMTs) of multiple sclerosis (MS) patients. To illustrate how strongly different structural assumptions about the regression coefficient surface can affect the results, we fit simple univariate functional linear models $E(Y(t)) = \beta_0(t) + \int X(s)\beta(s, t)ds$ regressing FA along the right cortico-spinal WMT (RCST-FA, $Y(t)$) on pre-smoothed FA along the *corpus callosum* WMT (CCA-FA, $X(s)$). Figure 1 shows estimated coefficient surfaces $\hat{\beta}(s, t)$ for penalized spline based fits with second or first order difference penalties (two leftmost panels) or a ridge penalty (third panel from left), as well as the coefficient surface implied by a functional principal component (FPC) based fit (right). Even though the surfaces are quite different, they result in (practically) identical fitted values in this example for a setting with simple non-identifiability. The first and third panels from the left in the top row of Figure 7 show the data used in this example. Graphical examples of the second kind of non-identifiability that persists even under penalization are shown in Figures 4 and 8 for synthetic and real data, respectively.

Our paper is structured as follows: Section 2 defines the model and data structure under discussion. We present an accessible rephrasing of the fundamental issue (Section 3) and derive necessary and sufficient conditions for settings in which $\beta(s, t)$ is or is not identifiable in Section 3.2. Section 4 follows up with an analysis of the scope of the problem based on simulated data, while Section 5 describes a real-world example of the issue. The main conclusions we draw from our analysis are that the complexity of observed functional covariates puts hard limits on the identifiability of coefficient surfaces in a number of ways, that these limits can be diagnosed based on the data at hand, and that pre-processing of functional covariates can often exacerbate identifiability issues.

2. Model and data structure

In the following, bold symbols denote vectors and matrices, and calligraphic letters denote function domains, function spaces, or sets of functions or vectors.

Define a simple function-on-function regression model as

$$Y_i(t) = \int_{\mathcal{S}} X_i(s)\beta(s, t)ds + \varepsilon_{it}, \tag{1}$$

where $Y_i(t)$ and $X_i(s)$, $i = 1, \dots, N$, are functional responses and covariates on closed intervals \mathcal{T} and \mathcal{S} in \mathbb{R} , respectively, and assume that they are realizations of zero-mean square integrable stochastic processes $Y(t) \in L^2[\mathcal{T}]$ and $X(s) \in L^2[\mathcal{S}]$ with continuous covariance functions, respectively. To simplify notation and exposition but without loss of generality, we assume that $E(Y_i(t)) \equiv 0$ and $E(X_i(s)) \equiv 0$. Propositions 3.1 to 3.3 in Section 3 directly carry over to any model with an additive predictor that includes terms like $\int_{\mathcal{S}} X_i(s)\beta(s, t)ds$. The further development in Section 3 leading up to Proposition 3.4 assumes that the estimate for $\beta(s, t)$ minimizes a (penalized) quadratic loss function, which is equivalent to maximizing the likelihood of a model with i.i.d. Gaussian errors ε_{it} , but does not strictly speaking depend on distributional assumptions about ε_{it} and is also very similar to the system of equations solved in each iteration of the penalized iteratively re-weighted least squares algorithm (P-IWLS; Wood, 2000) used to fit additive models like (1) for non-Gaussian responses.

Due to the assumptions on the functional covariates, they can be represented by a Karhunen-Loève expansion

$$X_i(s) = \sum_{m=1}^{\infty} \xi_{im}\phi_m(s), \tag{2}$$

with orthonormal $\phi_m(s)$, $\int_{\mathcal{S}} \phi_m(s)\phi_{m'}(s)ds = \delta_{mm'}$, and uncorrelated zero-mean FPC scores ξ_{im} with variances $\nu_1 \geq \nu_2 \geq \dots \geq 0, m \in \mathbb{N}$. The ν_m and $\phi_m(s), m = 1, \dots, M$, are the ordered eigenvalues and eigenfunctions of the covariance operator K^X of $X(s)$, respectively, with the covariance function given by Mercer’s theorem as

$$k^X(s, s') = E(X(s)X(s')) = \sum_{m=1}^{\infty} \nu_m\phi_m(s)\phi_m(s'). \tag{3}$$

Since estimating $\beta(s, t)$ is an inverse problem, some kind of regularization is required. Functional principal component based approaches like, for example, Yao et al. (2005) restrict $\beta(s, t)$ to lie in the span of the first K estimated eigenfunctions $\hat{\phi}_m(s)$, $m = 1, \dots, K$ for all t . The number of eigenfunctions K that is used serves as the (discrete) regularization parameter. In contrast, we will discuss and analyze a penalized approach. The underlying assumption is that $\beta(s, t)$ is a smooth function that can be well represented as a linear combination of suitable basis functions defined on $\mathcal{S} \times \mathcal{T}$.

In practice, functional responses $Y_i(t)$ and functional covariates $X_i(s)$ are observed on grid points $\mathbf{s}_i = (s_{i1}, \dots, s_{iS_i})$ and $\mathbf{t}_i = (t_{i1}, \dots, t_{iT_i})$. For simplicity, we assume those to be identical vectors \mathbf{s}, \mathbf{t} with lengths S and T , respectively, for each observation i . In the following, expressions $a(\mathbf{s})$ or $a(\mathbf{t})$ with a bold

argument denote the vector of evaluations of $a(\cdot)$ on the respective grid, e.g., $a(\mathbf{s}) = (a(s_1), \dots, a(s_S))^\top$.

Model (1) can then be approximated for observed data as

$$Y_i(\mathbf{t})^\top \approx \left(\mathbf{w} \cdot X_i(\mathbf{s}) \right)^\top \beta(\mathbf{s}, \mathbf{t}) + \boldsymbol{\varepsilon}_i, \quad (4)$$

$1 \times T$ $S \times 1$ $S \times 1$ $S \times T$ $1 \times T$

with $\beta(\mathbf{s}, \mathbf{t}) = [\beta(s_j, t_k)]_{\substack{j=1, \dots, S \\ k=1, \dots, T}}$ and $\boldsymbol{\varepsilon}_i = (\varepsilon_{it_1}, \dots, \varepsilon_{it_T})^\top$. We also define a weight vector \mathbf{w} for numerical integration, e.g. $\mathbf{w} = (w_j)_{j=1, \dots, S}$ for simple quadrature via Riemann sums, with w_j the length of the sub-interval of \mathcal{S} represented by s_j . The symbol \cdot denotes element-wise multiplication. The coefficient surface $\beta(\mathbf{s}, \mathbf{t})$ is represented using a tensor product spline basis

$$\beta(\mathbf{s}, \mathbf{t}) = \mathbf{B}_s \boldsymbol{\Theta} \mathbf{B}_t^\top, \quad (5)$$

$S \times K_s$ $K_s \times K_t$ $K_t \times T$

with basis matrices \mathbf{B}_s and \mathbf{B}_t of K_s and K_t basis functions evaluated in \mathbf{s} and \mathbf{t} , respectively, and spline coefficient matrix $\boldsymbol{\Theta}$. The roughness penalty matrix for the surface is given by $\mathbf{P} \equiv \mathbf{P}(\lambda_s, \lambda_t) = \lambda_s(\mathbf{P}_s \otimes \mathbf{I}_{K_t}) + \lambda_t(\mathbf{I}_{K_s} \otimes \mathbf{P}_t)$ (Wood, 2006), where λ_s, λ_t are smoothing parameters to be estimated from the data and \mathbf{P}_s and \mathbf{P}_t are the fixed and known marginal penalty matrices for the s - and t -directions, respectively.

Estimation and inference is described in more detail in Ivanescu et al. (2015) and Scheipl et al. (2015). In the following, our considerations are not limited to simple models such as model (1), but carry over to more general models $\tilde{Y}_i(t) = \eta_i(t) + \int_{\mathcal{S}} X_i(s) \beta(s, t) ds + \varepsilon_{it}$ by using $Y_i(t) = \tilde{Y}_i(t) - \eta_i(t)$. The additive predictor $\eta_i(t)$ represents the sum of other terms in the model such as a global functional intercept $\beta_0(t)$, index-varying linear or smooth effects of scalar covariates x like $x_i \beta(t)$ or $f(x_i, t)$, scalar or functional random effects, etc. Scheipl et al. (2015) contains methods and applied examples for this class of flexible additive functional regression models. Of course, these more general models may suffer from additional identifiability problems caused by collinearity or concurrency of the terms in the additive predictor which are outside the scope of this paper. While we focus our discussion on a spline-based approach, the representation in (5) also accommodates other choices of basis functions and penalties.

3. Identifiability

In this section, we discuss potential sources of non-identifiability in model (1). The first subsection restates some known results on these issues for the theoretical model (1) with truly functional observations $X_i(s)$ and $Y_i(t)$. Subsection 3.2 then discusses identifiability for the finite resolution vector data available in practice.

3.1. Identifiability in the theoretical model

It is well known (e.g. Prchal and Sarda (2007, c.f. p. 5), He et al. (2000, Th. 4.3. c)) that coefficient surfaces are identifiable only up to the addition

of functions in the kernel of K^X , i.e., if $\beta(\cdot, t)$ fulfills model (1), so does $\beta(\cdot, t) + \beta_K(\cdot, t)$ for any $\beta_K(\cdot, t)$ with $\int_{\mathcal{S}} k^X(s, v)\beta_K(v, t)dv = 0$ for all s, t . Thus, we have identifiability only when the kernel is trivial.

Proposition 3.1. *The coefficient surface $\beta(s, t)$ in (1) is identifiable if and only if $\text{ke}(K^X) = \{0\}$.*

An important secondary consequence is that large changes in $\beta(s, t)$ can leave the predicted responses $\int_{\mathcal{S}} X_i(s)\beta(s, t)ds$ entirely unaffected. Thus, strategies for detection of identifiability problems cannot be based on predictive performance in cross-validation, bootstrapping and related methods.

Non-identifiability in Proposition 3.1 occurs when $\text{ke}(K^X)$ is non-trivial, when the eigenfunctions in (2) with non-zero eigenvalues ν_m do not span the $L^2[\mathcal{S}]$. While it is possible to assume a trivial kernel in theory (e.g. Prchal and Sarda, 2007, equation (4)), in practice, functional covariates are observed on a finite number of grid points S and the empirical covariance for N observations thus can have at most $\min(N, S)$ non-zero eigenvalues. As is exploited in functional principal component analysis (e.g. Ramsay and Silverman, 2005), functional observations are often simple enough to be represented accurately by a relatively small number of eigenfunctions, with eigenvalues of higher order small compared to noise or measurement error. It is also wide-spread practice to use pre-smoothed versions of observed functional covariates as inputs for models like (1) (e.g. James, 2002; Ramsay and Silverman, 2005), and these will have a non-empty kernel since they are represented as linear combinations of a limited number of basis functions. Basis function representations of $X(s)$ are also used when sparsely or incompletely observed functional covariates have to be imputed on a grid of s -values to be used as inputs for model (1) (e.g. Goldsmith et al., 2011).

In the following section, we thus investigate identifiability problems for finite-sample finite-resolution functional data and the interplay between the rank of the observed covariance, the rank of the basis used to represent $\beta(s, t)$ in s -direction and the penalty used in the penalized estimation approach for $\beta(s, t)$ introduced in Section 2.

3.2. Identifiability in practice

Rank-deficient design matrix

In practice, $\beta(s, t)$ is represented as a linear combination of a finite number $K_s K_t$ of basis functions, see (5). For the following, we will assume that the corresponding approximation error is negligible and that a suitably flexible basis has been chosen for $\beta(s, t)$. Combining (4) and (5), we can write the model as

$$\mathbf{Y} = \mathbf{X} \mathbf{W} \mathbf{B}_s \mathbf{\Theta} \mathbf{B}_t^T + \boldsymbol{\varepsilon}, \tag{6}$$

$\begin{matrix} N \times T & N \times S & S \times S & S \times K_s & K_s \times K_t & K_t \times T & N \times T \end{matrix}$

where $\mathbf{Y} = [Y_i(t_j)]_{\substack{i=1, \dots, N \\ j=1, \dots, T}}$, $\mathbf{X} = [X_i(s_l)]_{\substack{i=1, \dots, N \\ l=1, \dots, S}}$, $\mathbf{W} = \text{diag}(\mathbf{w})$ and $\boldsymbol{\varepsilon} = [\varepsilon_{it_j}]_{\substack{i=1, \dots, N \\ j=1, \dots, T}}$. Using $\text{vec}(\mathbf{ACB}) = (\mathbf{B}^T \otimes \mathbf{A}) \text{vec}(\mathbf{C})$ yields

$$\text{vec}(\mathbf{Y}) = \begin{bmatrix} \mathbf{B}_t & (\mathbf{X} & \mathbf{W} & \mathbf{B}_s) \\ NT \times 1 & T \times K_t & N \times S & S \times S & S \times K_s \\ & & & & K_t K_s \times 1 \end{bmatrix} \text{vec}(\boldsymbol{\Theta}) + \text{vec}(\boldsymbol{\varepsilon}). \quad (7)$$

In the linear regression model (7) for $\mathbf{y} = \text{vec}(\mathbf{Y})$, the parameter vector $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta})$ is identifiable if and only if the design matrix $\mathbf{D} = \mathbf{B}_t \otimes (\mathbf{X}\mathbf{W}\mathbf{B}_s)$ is of full column-rank.

The rank of \mathbf{D} is equal to $\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{B}_t) \text{rank}(\mathbf{X}\mathbf{W}\mathbf{B}_s)$. \mathbf{B}_t will typically be of full rank K_t as long as $K_t \leq T$, as the K_t spline functions form a basis and the columns of \mathbf{B}_t are thus linearly independent for non-pathological cases. For \mathbf{X} , let $\mathbf{X} = \boldsymbol{\Xi}\boldsymbol{\Phi}$ be the empirical version of the Karhunen-Loève expansion (2), where $\mathbf{X}^\top \mathbf{X} = \boldsymbol{\Phi}^\top \boldsymbol{\Lambda} \boldsymbol{\Phi}$ with $\boldsymbol{\Phi}^\top$ an $S \times M$ orthonormal matrix of eigenvectors, $M = \text{rank}(\mathbf{X})$, $\boldsymbol{\Lambda} = \text{diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_M)$ a diagonal matrix of ordered positive eigenvalues and $\boldsymbol{\Xi}$ containing M columns of estimated scores with empirical variances $\hat{\lambda}_1, \dots, \hat{\lambda}_M$. Then, by construction, the matrix $\mathbf{X}\mathbf{W}\mathbf{B}_s$ is at most of rank $\min(N, M, K_s, S) = \min(M, K_s)$, since $M \leq \min(N, S)$.

We then have the following proposition:

Proposition 3.2. *Assume that \mathbf{B}_t is of full rank K_t . Then, the design matrix $\mathbf{D} = \mathbf{B}_t \otimes (\mathbf{X}\mathbf{W}\mathbf{B}_s)$ in model (7) is rank-deficient if and only if*

- a) $M < K_s$ or
- b) if $M \geq K_s$, but $\text{rank}(\boldsymbol{\Phi}\mathbf{W}\mathbf{B}_s) < K_s$.

Proof. See Appendix A.1. □

Proposition 3.2 yields a direct criterion to check for rank-deficient design matrices. Case a) corresponds to a low-rank covariance for the X -process. In this case, the functional predictor does not carry enough information, as measured by the number of eigenfunctions $\phi_m(s)$ with non-zero eigenvalues, compared to the number of parameters to estimate. Case b) means that even if $M \geq K_s$, non-identifiability can occur if the span of the basis used for $\beta(s, t)$ in s -direction contains functions in $\text{ke}(K^X)$, as measured by numerical integration using the integration weights \mathbf{w} . More intuitively, this means that the basis for the s -direction of $\beta(s, t)$ accommodates modes of variation orthogonal to those of the $X(s)$ -process.

Note that identifiability is determined by the interplay between the complexity of the $X(s)$ and the coefficient basis for $\beta(s, t)$. Thus, more data will not necessarily resolve identifiability issues: A finer grid \mathbf{s} will only eliminate identifiability problems present on a coarser grid if there is sufficient small-scale structure in $X(s)$ that is also present in the basis used for $\beta(s, t)$ in s -direction. Increasing the sample size N will likewise eliminate problems with identifiability only if the low-rank of the covariate's covariance is due to small sample size.

The effect of the penalty

In cases of non-identifiability, the best we can hope for is partial identifiability of the parameters in a parameter subspace, i.e., identifiability under additional assumptions on the parameters. In this vein, functional principal component

regression (e.g. Yao et al., 2005) restricts $\beta(s, t)$ to lie in the span of the first K eigenfunctions $\phi_m(s)$, $m = 1, \dots, K$, for all t . Remaining problems then include the fact that the $\hat{\phi}_m(s)$ are estimated quantities in practice, with corresponding measurement error, and the choice of K , which can strongly affect the shape of the resulting function estimate (Crainiceanu et al., 2009). Also, this approach couples assumptions on the shape of $\beta(s, t)$ to properties of the space spanned by the retained $\phi_m(s)$. In particular, their smoothness determines that of $\beta(s, t)$, and inclusion of higher-order eigenfunctions often leads to wiggly surface estimates that are hard to interpret and unstable under replication.

Here, we focus on a penalized approach assuming smoothness of $\beta(s, t)$ and investigate the effects of the penalty on identifiability. While the use of a penalized approach is well-known to avoid identifiability problems due to high correlation between observations at neighboring grid points (e.g. Ramsay and Silverman, 2005, Ch. 15.2), the full interplay between penalty and identifiability is, we believe, not fully understood and under-appreciated.

Consider again the design matrix $\mathbf{D} = \mathbf{B}_t \otimes (\mathbf{X}\mathbf{W}\mathbf{B}_s)$ of rank d with the singular value decomposition $\mathbf{D} = \mathbf{V}\boldsymbol{\Sigma}\mathbf{U}^\top = (\mathbf{V}_t \otimes \mathbf{V}_s)(\boldsymbol{\Sigma}_t \otimes \boldsymbol{\Sigma}_s)(\mathbf{U}_t^\top \otimes \mathbf{U}_s^\top)$, with $\mathbf{V}_t\boldsymbol{\Sigma}_t\mathbf{U}_t^\top$ and $\mathbf{V}_s\boldsymbol{\Sigma}_s\mathbf{U}_s^\top$ the singular value decompositions of \mathbf{B}_t and $\mathbf{D}_s := \mathbf{X}\mathbf{W}\mathbf{B}_s$, respectively. Let indices $_+$ and $_0$ denote the corresponding sub-matrices obtained by removing columns and/or rows corresponding to zero and non-zero singular values, respectively. We assume in the following that \mathbf{B}_t is of full rank K_t . Then, $\mathbf{D} = \mathbf{V}_+\boldsymbol{\Sigma}_+\mathbf{U}_+^\top = (\mathbf{V}_t \otimes \mathbf{V}_{s+})(\boldsymbol{\Sigma}_t \otimes \boldsymbol{\Sigma}_{s+})(\mathbf{U}_t^\top \otimes \mathbf{U}_{s+}^\top)$.

Thus, for any given $\boldsymbol{\theta}_*$ with $\mathbf{D}\boldsymbol{\theta}_* =: \mathbf{f}$, there exists a linear subspace $\mathcal{H}_f \subset \mathbb{R}^{K_s K_t}$ of dimension $(K_s K_t - d)$ given by

$$\mathcal{H}_f = \{\boldsymbol{\theta} \in \mathbb{R}^{K_s K_t} : \mathbf{D}\boldsymbol{\theta} = \mathbf{f}\} = \{\boldsymbol{\theta}_* + \boldsymbol{\theta}_0 : \boldsymbol{\theta}_0 \in \text{im}(\mathbf{U}_0)\}$$

that yields identical fits \mathbf{f} .

If we assume our parameter function to come from a space of smooth functions, we can select the smoothest solution on a given hyperplane \mathcal{H}_f by minimizing $\boldsymbol{\theta}^\top \mathbf{P}\boldsymbol{\theta}$ for a suitable penalty matrix \mathbf{P} that penalizes roughness of the function parameterized by $\boldsymbol{\theta}$. *Simple* non-identifiability occurs if this minimum is unique, *persistent* non-identifiability occurs if it is not. We have the following proposition regarding uniqueness of the corresponding minimum.

Proposition 3.3. *Let $\mathbf{P} = \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s})$, with \mathbf{P}_s and \mathbf{P}_t positive semi-definite matrices. Assume that \mathbf{B}_t is of full rank K_t , that $\text{rank}(\mathbf{P}_t) < K_t$ and that $\lambda_s > 0, \lambda_t \geq 0$. Then, for any $\mathbf{f} \in \text{im}(\mathbf{D})$ there is a unique minimum $\min_{\boldsymbol{\theta} \in \mathcal{H}_f} \{\boldsymbol{\theta}^\top \mathbf{P}\boldsymbol{\theta}\}$ if and only if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$.*

Proof. See Appendix A.2. □

The assumption that \mathbf{P}_t is of less-than-full rank is natural in the context of derivative-based penalties and excludes cases like the ridge penalty $\mathbf{P}_t = \mathbf{I}_{K_t}$, which would have the same effect as a full-rank penalty $\mathbf{P}_s = \mathbf{I}_{K_s}$, even in cases of a kernel overlap between $\mathbf{D}_s^\top \mathbf{D}_s$ and \mathbf{P}_s . A potentially full-rank \mathbf{P}_t would change the ‘if and only if’ in Propositions 3.3 and 3.4 below to ‘if’.

Proposition 3.3 shows that in the case of a kernel overlap $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) \neq \{\mathbf{0}\}$, the additional side condition $\boldsymbol{\theta}^\top \mathbf{P} \boldsymbol{\theta} \rightarrow \min$ does not yield a unique smoothest point on the hyperplane and the model remains unidentifiable. On the other hand, if there is no kernel overlap, there is a unique smoothest point $\boldsymbol{\theta}_f \in \mathcal{H}_f$ and this unique point has the form of a projection of $\boldsymbol{\theta}$ along the hyperplane. Note that for the ridge penalty $\mathbf{P} = \lambda \mathbf{I}_{K_s K_t}$, one obtains the projection onto the image $\text{im}(\mathbf{D})$, which sets the part in the kernel of $\mathbf{D}^\top \mathbf{D}$ to zero. More generally, a smoothness penalty \mathbf{P} generates a projection that may have a non-zero component in the kernel of $\mathbf{D}^\top \mathbf{D}$ if this yields a smaller overall penalty value. In this case of no kernel overlap, we thus have a weak form of identifiability, which guarantees that there is a unique smoothest representative on any hyperplane of parameters giving the same conditional distribution for \mathbf{Y} .

This characterization, which only requires checking of design matrix and penalty in s -direction and not for the full model, carries over to the penalized maximum likelihood or least squares estimation problem

$$\min_{\boldsymbol{\theta}} \{ \|\mathbf{y} - \mathbf{D}\boldsymbol{\theta}\|^2 + \lambda_s \boldsymbol{\theta}^\top (\mathbf{I}_{K_t} \otimes \mathbf{P}_s) \boldsymbol{\theta} + \lambda_t \boldsymbol{\theta}^\top (\mathbf{P}_t \otimes \mathbf{I}_{K_s}) \boldsymbol{\theta} \} \quad (8)$$

for some $\lambda_s > 0, \lambda_t \geq 0$.

Proposition 3.4. *Assume that \mathbf{B}_t is of full rank K_t , that $\text{rank}(\mathbf{P}_t) < K_t$ and that $\lambda_s > 0, \lambda_t \geq 0$. Then, there is a unique penalized least squares solution for (8) if and only if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$.*

Proof. See Appendix A.3. □

Proposition 3.4 gives a criterion for the uniqueness of the penalized least squares solution. We can show how the penalty achieves this uniqueness by writing (8) as a nested minimization problem. Here, the outer minimization finds the $\mathbf{f} = \mathbf{D}\boldsymbol{\theta}$ with optimal fit to the data, and the inner minimization minimizes the penalty term over \mathcal{H}_f to obtain the smoothest solution for a given level of residual variation.

$$\begin{aligned} & \min_{\boldsymbol{\theta}} \{ \|\mathbf{y} - \mathbf{D}\boldsymbol{\theta}\|^2 + \lambda_s \boldsymbol{\theta}^\top (\mathbf{I}_{K_t} \otimes \mathbf{P}_s) \boldsymbol{\theta} + \lambda_t \boldsymbol{\theta}^\top (\mathbf{P}_t \otimes \mathbf{I}_{K_s}) \boldsymbol{\theta} \} \\ &= \min_{\mathbf{f} \in \text{im}(\mathbf{D})} \min_{\boldsymbol{\theta}} \{ \|\mathbf{y} - \mathbf{f}\|^2 + \boldsymbol{\theta}^\top \mathbf{P} \boldsymbol{\theta} \quad \text{s.t.} \quad \mathbf{D}\boldsymbol{\theta} = \mathbf{f} \} \\ &= \min_{\mathbf{f} \in \text{im}(\mathbf{D})} \{ \|\mathbf{y} - \mathbf{D}\boldsymbol{\theta}_f\|^2 + \boldsymbol{\theta}_f^\top \mathbf{P} \boldsymbol{\theta}_f \} \\ &= \min_{\mathbf{v}_+ \in \mathbb{R}^d} \{ \|\mathbf{y} - \mathbf{V}_+ \boldsymbol{\Sigma}_+ \mathbf{v}_+\|^2 + \mathbf{v}_+^\top \mathbf{U}_+^\top \mathbf{H}^\top \mathbf{P} \mathbf{H} \mathbf{U}_+ \mathbf{v}_+ \} \end{aligned}$$

where $\boldsymbol{\theta}_f = \mathbf{H} \mathbf{U}_+ \mathbf{v}_+$ for a given \mathbf{f} is uniquely defined as in (11), with $\mathbf{v}_+ = \boldsymbol{\Sigma}_+^{-1} \mathbf{V}_+^\top \mathbf{f}$ and $\mathbf{H} = (\mathbf{I}_{K_s K_t} - \mathbf{U}_0 (\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0)^{-1} \mathbf{U}_0^\top \mathbf{P})$, if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$.

As $\mathbf{V}_+ \boldsymbol{\Sigma}_+$ is a matrix of full column rank d , this minimization problem has a unique solution that, for given λ_s, λ_t , balances the fit to the data and smoothness.

To summarize, in the case of no kernel overlap, i.e., $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$, we obtain a weaker form of identifiability even when the design matrix \mathbf{D} is not of full rank, which guarantees that there is a unique smoothest representative on any hyperplane of parameters giving the same conditional distribution for \mathbf{y} . Then, there will also be a unique solution to the penalized estimation problem, which is the smoothest representative on the hyperplane of possible solutions with equally good fit.

In practice, λ_s and λ_t are estimated from the data. We here do not investigate the more complex case when λ_s, λ_t are not fixed. It should also be noted that $(\mathbf{D}^\top \mathbf{D} + \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s}))$ may still be close to singular even in cases of no kernel overlap if smoothing parameters are very small, with corresponding reduced stability in estimation.

3.3. Diagnostics, practical recommendations and countermeasures

In order to safeguard against misleading coefficient estimates in practical applications of functional regression, it is necessary 1) to develop empirical criteria for diagnosing problematic data settings in which the coefficient function is not identifiable based on the available data, or where only the penalty ensures unique estimates, 2) to avoid pre-processing protocols and tuning parameters that increase the likelihood of identifiability issues and 3) to develop improved algorithms for estimating function-on-function effects that are less prone to severely misleading estimates in problematic settings.

Diagnostics We are interested in identifying settings with *simple* non-identifiability in which only the penalty term guarantees the existence of a unique solution. Following Proposition 3.2, the most direct approach to do so is to compute the condition number of $\mathbf{D}_s^\top \mathbf{D}_s = (\mathbf{X}\mathbf{W}\mathbf{B}_s)^\top \mathbf{X}\mathbf{W}\mathbf{B}_s$ and choose a suitable cut-off (10^6 , in the following) for numeric rank deficiency.

In addition, propositions 3.3 and 3.4 indicate that a measure of the degree of overlap between the spans of $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s)$ and $\text{ke}(\mathbf{P}_s)$ can be used to detect *persistent* non-identifiability that remains despite the penalization. In our empirical evaluation of such measures, we found that a measure for the distance between the spans of two matrices introduced in Larsson and Villani (2001), when modified for our setting, showed the most promise as the resulting measure is free of tuning parameters and can be computed quickly from the data.

In particular, we modify the original definition of Larsson-Villani in order to accommodate two matrices of unequal column numbers. We then define the amount of overlap \bigcap_{LV} between the span of two matrices $\mathbf{A} \in \mathbb{R}^{n \times p_A}$, $\mathbf{B} \in \mathbb{R}^{n \times p_B}$, $n > p_A, p_B$, by

$$\bigcap_{LV}(\mathbf{A}, \mathbf{B}) = \text{trace}(\mathbf{V}_B^\top \mathbf{V}_A \mathbf{V}_A^\top \mathbf{V}_B).$$

Here, \mathbf{V}_Z is a matrix containing the left singular vectors of the matrix \mathbf{Z} and is thus an orthogonal matrix spanning the same column space as \mathbf{Z} , $\mathbf{Z} \in \{\mathbf{A}, \mathbf{B}\}$. It is easy to see that this measure is symmetric, $\bigcap_{LV}(\mathbf{A}, \mathbf{B}) = \bigcap_{LV}(\mathbf{B}, \mathbf{A})$.

Similarly to Theorem 2 in Larsson and Villani (2001), one can also show that $\bigcap_{LV}(\mathbf{A}, \mathbf{B}) \in [0, \min(p_A, p_B)]$, with the overlap assuming its maximum of $\min(p_A, p_B)$ iff $\mathbf{A} \in \text{im}(\mathbf{B})$ or $\mathbf{B} \in \text{im}(\mathbf{A})$ and its minimum of 0 iff $\mathbf{A} \in \text{im}(\mathbf{B}_\perp)$ or $\mathbf{B}_\perp \in \text{im}(\mathbf{A})$, where the $n \times (n - p_Z)$ matrix \mathbf{Z}_\perp is the orthonormal complement of \mathbf{Z} , $\mathbf{Z} \in \{\mathbf{A}, \mathbf{B}\}$.

To measure the degree of overlap between the kernels of $\mathbf{D}_s^\top \mathbf{D}_s$ and of \mathbf{P}_s , we could use $\bigcap_{LV}((\mathbf{D}_s^\top \mathbf{D}_s)_\perp, \mathbf{P}_{s\perp})$, as the span of $(\mathbf{D}_s^\top \mathbf{D}_s)_\perp$ corresponds to the kernel of $\mathbf{D}_s^\top \mathbf{D}_s$ and the span of $\mathbf{P}_{s\perp}$ corresponds to the kernel of \mathbf{P}_s . In the following, we will however use

$$\bigcap_{X_\perp \mathcal{P}_\perp} = \bigcap_{LV}((\mathbf{X}^\top \mathbf{X})_\perp, \mathbf{W} \mathbf{B}_s \mathbf{P}_{s\perp}) = \bigcap_{LV}((\mathbf{X}^\top)_\perp, \mathbf{W} \mathbf{B}_s \mathbf{P}_{s\perp}), \quad (9)$$

as this choice obtained slightly better sensitivity and specificity for the detection of problematic settings in our simulations in Section 4. Moreover, one can easily show that

$$\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\} \quad \Leftrightarrow \quad \text{ke}(\mathbf{X}^\top \mathbf{X}) \cap \{\mathbf{W} \mathbf{B}_s \mathbf{x} \mid \mathbf{x} \in \text{ke}(\mathbf{P}_s)\} = \{\mathbf{0}\}$$

such that the two formulations address the same question. For $\mathbf{W} = \mathbf{I}_S$, this measure has the interpretation of the overlap between the empirical null-space of the observed $X(s)$ process or $\text{ke}(K^X)$ and $\mathcal{P}_{s\perp}$, the space of functions not penalized by the penalty defined by \mathbf{P}_s , evaluated on the grid given by \mathbf{s} . It can be determined quickly and accurately before the model is fit. Problematic cases are indicated by overlap measures ≥ 1 , as this is indicative of an at least one-dimensional sub-space of functions contained in the kernel overlap.

Practical recommendations The theoretical results suggest several recommendations for pre-processing of functional covariates and choice of the penalty in practice.

1. Pre-smoothing of functional covariates is commonly done to remove measurement error and/or obtain functions on a common grid (e.g. James, 2002; Ramsay and Silverman, 2005; Goldsmith et al., 2011). If the resulting (effective) rank of the smoothed covariate process drops below K_s , this will lead to models that are only identifiable through the penalty term. We thus recommend to use a sufficiently large number of FPCs and/or spline basis functions if such pre-processing is required. As a peculiar consequence of this point it may be preferable in some cases to accept a small amount of measurement error-induced attenuation in $\hat{\beta}(s, t)$ based on noisy, unprocessed $X(s)$ in order to avoid a potentially much larger non-identifiability-induced error in $\hat{\beta}(s, t)$ based on low-rank, pre-processed $X(s)$.
2. Curve-wise centering of functional covariates such that $\sum_{l=1}^S X_i(s_l) = 0$ for all i is sometimes used e.g. in the context of spectroscopy data to remove the optical offset (c.f. Fuchs et al., 2015). Then, constant functions lie in the kernel of K^X , $\text{ke}(K^X)$. This is not recommended if a penalty is used that does not penalize constant functions (as most difference or derivative-based penalties do) to avoid non-identifiability.

3. Penalties with larger null-spaces increase the likelihood of a kernel overlap and resulting non-identifiability problems. For difference or derivative-based penalties, for example, a penalty penalizing deviations from constant functions (first order differences or derivatives) would thus be preferable in this sense to higher-order differences/derivatives. Constant coefficient functions, which then span the penalty null-space, correspond to models with the mean over the functions as covariate – which are often used by practitioners – and thus also lend themselves to intuitive interpretations. In particular, for penalties where constant coefficient functions span the penalty nullspace, it is straightforward to see that only X-processes that are centered curve-wise will result in a kernel overlap (unless smoothing parameters are estimated to be very small). Unless curve-wise centering is performed as discussed in 2., such processes will typically occur rarely and using first-order difference/derivative penalties should thus guard against many if not most serious identifiability issues in practice.

Countermeasures The third point above suggests modifications of the penalty null-space to avoid non-identifiability caused by potential overlap between the penalty null-space $\mathcal{P}_{s\perp}$ and $\text{ke}(K^X)$. Alternatively, the estimated coefficient surface can be constrained to be orthogonal to functions in the overlap of $\mathcal{P}_{s\perp}$ and $\text{ke}(K^X)$. We describe three approaches using penalties with empty null-spaces and one constraint-based approach; a systematic comparison of their performance is given in Section 4.2.

1. The simplest approach is the use of a simple ridge penalty $\mathbf{P}_s = \mathbf{I}_{K_s}$. However, the resulting estimates will typically not have good smoothness properties as the ridge penalty is not a roughness penalty in the conventional sense and its bias towards small absolute values of $\beta(s, t)$ may increase estimation error, cf. Figures 1, 4 and 5.
2. A second approach uses a modified marginal penalty matrix without null-space along the lines of the so-called “shrinkage approach” described in Marra and Wood (2011) and originally developed for the purpose of variable selection in generalized additive models. Marra and Wood (2011) replace the marginal penalty $\mathbf{P}_s = \mathbf{\Gamma} \text{diag}(\rho_1, \dots, \rho_{K_s}) \mathbf{\Gamma}^T$, with eigenvectors contained in $\mathbf{\Gamma}$, eigenvalues $\rho_1, \dots, \rho_{K_s}$ and rank $\tilde{K}_s < K_s$, by a full rank marginal penalty

$$\tilde{\mathbf{P}}_s = \mathbf{\Gamma} \text{diag}(\rho_1, \dots, \rho_{\tilde{K}_s}, \epsilon \rho_{\tilde{K}_s}, \dots, \epsilon \rho_{\tilde{K}_s}) \mathbf{\Gamma}^T.$$

This substitutes the zero eigenvalues $\rho_{\tilde{K}_s+1}, \dots, \rho_{K_s}$ with $\epsilon \rho_{\tilde{K}_s}$ for all $k = \tilde{K}_s + 1, \dots, K_s$ using $0 < \epsilon \ll 1$, and thereby adds a small amount of penalization to parameter vectors in the null space of the original penalty. In the following, we will refer to such a modified penalty as a full-rank penalty. By imposing a small degree of regularization on functions in $\mathcal{P}_{s\perp}$ one can in principle preserve the attractive smoothing properties of the original penalty while still avoiding large artifacts due to non-identifiability resulting from a kernel overlap. We use $\epsilon = 0.1$ as suggested in Marra and

Wood (2011), but results show that this choice is not always effective in removing artifacts if the estimated smoothing parameter is very small and overall penalization of the fit is weak.

3. In a similar effort to avoid spurious estimates in scalar-on-function regression, James and Silverman (2005, their eq. (16)) suggested using the empirical FPCs of $X(s)$ scaled by their inverse eigenvalues as a penalty. This penalizes coefficient functions with large variability in directions in which $X(s)$ varies very little or not at all (i.e., in $\text{ke}(K^X)$). We adapted this approach to our function-on-function setting by replacing the conventional difference operator based B-spline penalty matrix with

$$\mathbf{P}_s = \mathbf{B}_s^T \sum_{m=1}^{\min(N,S)} \hat{\nu}_m^{-1} \text{diag}(\mathbf{w} \cdot \hat{\phi}_m(s)^2) \mathbf{B}_s$$

with estimated FPCs and eigenvalues $\hat{\phi}_m(s)$ and $\hat{\nu}_m$, $m = 1, \dots, \min(N, S)$. For a given t_0 , this penalty matrix approximates the marginal penalty term $\sum_{m=1}^{\min(N,S)} \int (\hat{\nu}_m^{-1/2} \hat{\phi}_m(s) \beta(s, t_0))^2 ds$ suggested by James and Silverman (2005).

The empirical $\hat{\phi}_m(s)$ and $\hat{\nu}_m$ have to be estimated from a singular value decomposition of \mathbf{X} . It is unclear, however, how to compute the inverse eigenvalues if \mathbf{X} is of low rank, i.e., if some of the $\hat{\nu}_m$ are (numerically) zero, which is of course precisely the setting in which this penalty might yield more stable estimates. In our experiments, we tried replacement of the zero eigenvalues with the smallest non-zero eigenvalue and a variety of other replacement schemes, but results from this approach seem to be fairly sensitive to both the chosen replacement scheme for zero eigenvalues and to the estimated FPCs themselves.

4. Instead of the augmented or alternative penalization schemes for components of the coefficient surface that lie in the overlap $\text{ke}(K^X) \cap \mathcal{P}_{s\perp}$, we can also directly constrain these components to be zero. As these components are not estimable from the data, such constraints are a plausible default but need to be taken into account for the interpretation of the estimated surface. To implement them, we compute a basis of the overlap and constrain the coefficient surface evaluated on the observed grid to be orthogonal to vectors in the span of that basis. Let $\mathbf{X}_\perp^\top = (\mathbf{X}^\top)_\perp$ and the $S \times K_s$ matrix \mathbf{B}_s containing the marginal basis functions over \mathcal{S} evaluated on \mathbf{s} . A basis spanning the overlap is then defined by the left singular vectors $\mathbf{V}_{C_{s+}}^\top$ of the matrix $\mathbf{X}_\perp^\top ((\mathbf{X}_\perp^\top)^\top \mathbf{X}_\perp^\top)^{-1} (\mathbf{X}_\perp^\top)^\top \text{diag}(\mathbf{w}) \mathbf{B}_s \mathbf{P}_{s\perp}$ that have positive singular values. For intuition, consider that if $\mathbf{w} = \mathbf{1}$, the expression above projects $\mathbf{B}_s \mathbf{P}_{s\perp}$, i.e., a basis for $\mathcal{P}_{s\perp}$, into the span of \mathbf{X}_\perp^\top , i.e., a basis for $\text{ke}(K^X)$. This yields a basis for the intersection of the spans of $\mathbf{B}_s \mathbf{P}_{s\perp}$ and \mathbf{X}_\perp^\top . The expression above is cheap to compute in a numerically stable way using the QR decomposition of \mathbf{X}_\perp^\top . We then estimate Θ under the constraints $\mathbf{V}_{C_{s+}}^\top \text{diag}(\mathbf{w}) \mathbf{B}_s \Theta \stackrel{!}{=} \mathbf{0}$. To give an example, the constrained coefficient surface estimate for a curve-wise

centered functional covariate that does not contain any constant components will be centered around zero and not be offset in any direction if the penalty null-space includes constant functions. Note that, different from the augmented or alternative penalties described above, these constraints are empty if $\text{ke}(K^X) \cap \mathcal{P}_{s\perp}$ is empty and penalties with constraints then reduce to the usual penalties without constraints. The constraint definition above is quite general and can be used for any basis with a quadratic penalty, it is not limited to the difference penalties we focus on. Specifically, it will also be applicable for derivative based penalties as well as some of the PDE-based penalties introduced by Ramsay et al. (2007) whenever identifiability issues arise. By default, the `pffr` function uses the diagnostic criterion (9) combined with a check for a rank-deficient design matrix to determine the presence of persistent non-identifiability. If persistent non-identifiability is found, a corresponding warning is issued and the constraints developed here are applied to the fit. Interpretation of the constrained coefficient surface estimates requires careful attention. For example, if constrained surface estimates are centered around zero because $\text{ke}(K^X) \cap \mathcal{P}_{s\perp}$ contains constant functions, the signs of different regions of $\hat{\beta}(s, t)$ are not interpretable due to the estimated absolute level of $\hat{\beta}(s, t)$ being essentially arbitrary.

4. Simulation study

This section presents results on the practical consequences of the theoretical development in Section 3. Specifically, we investigate the performance of the tensor product spline-based approach given in Section 2 on artificial data of varying complexity and noise levels in terms of estimation accuracy and use the simulation results to validate the diagnostics for problematic settings we have developed. Subsequently, we present results for the modified full-rank penalties of Marra and Wood (2011), the FPC-based penalty of James and Silverman (2005) and the constrained estimates described in Section 3.3. All models were fitted with the `pffr()` function available in the `refund` package, which estimates the smoothing parameters using restricted maximum likelihood (REML).

4.1. Simulation setup

We simulate data from data generating process (1), with $n = 50$ subjects, $\mathcal{T} = \mathcal{S} = [0, 1]$ and $S = 100$ grid-points for $X_i(s) = \sum_{m=1}^M \xi_{im} \phi_m(s)$. The effect surface $\beta(s, t)$ is estimated using tensor product cubic B-splines. We set $K_t = 10$ and use a marginal first order difference penalty for the t -direction. Test runs showed that results are insensitive to the number of grid-points for the response; we used $T = 50$ grid-points for $Y(t)$. Twenty replications were simulated for all sensible combinations of the following parameters (144000 replicates in total). For the fitting algorithm, we vary

- the **marginal roughness penalties** for the spline coefficients: either second order difference penalties (“ Δ^2 ”) or first order difference penalties (“ Δ^1 ”). For the second order differences, coefficient vectors in the penalty’s null-space $\text{ke}(\mathbf{P}(\lambda_s, \lambda_t))$ parameterize surfaces that are constant or linear in both directions. For the first order differences, coefficient vectors in $\text{ke}(\mathbf{P}(\lambda_s, \lambda_t))$ parameterize constant surfaces. Note that rank deficiencies can increase if associated smoothing parameters become sufficiently small. In that case, the penalty null-space effectively becomes the entire span of the associated basis functions.
- and the **number of basis functions** over \mathcal{S} : $K_s \in \{5, 8, 12\}$.

For the data generating process, we vary the following parameters:

- **number of eigenfunctions for the $X(s)$ -trajectories** with non-zero eigenvalues:
 $M \in \{3, 5, 8, 12, 20\}$. This means we have settings with $M \leq K_s$ and $M > K_s$ for most M . Note that the effective numerical rank of a simulated \mathbf{X} can be (much) lower than M depending on the speed with which the eigenvalues decrease.
- **signal-to-noise ratio**: $\text{SNR}_\varepsilon = \frac{\text{sd}(\int_{\mathcal{S}} X_i(s)\beta(s,t)ds)}{\text{sd}(\varepsilon_{it})} \in \{2, 10, 1000\}$, where $\text{sd}(x)$ is the empirical standard deviation of x . This corresponds to high and intermediate noise levels for realistic scenarios as well as settings with almost no noise to check the theoretical properties.
- **FPC systems** for $X_i(s) = \sum_{m=1}^M \phi_m(s)\xi_{im}$ with $\xi_{im} \sim N(0, \nu_m)$, with different patterns of decrease in the eigenvalues ν_m of the covariance operators: either a linear decrease ($\nu_m = \frac{M+1-m}{M}$) or an exponential decrease ($\nu_m = \exp(-\frac{m-1}{2})$). Some of these processes are constructed so that their covariance operator kernels $\text{ke}(K^X)$ include functions in the penalty null-space $\mathcal{P}_{s\perp}$.
 - *Poly*: eigenfunctions are orthogonal polynomials of degree 0 to $M-1$ with linear or exponentially decreasing eigenvalues (*Poly, Lin* and *Poly, Exp*, respectively). For *Poly*, $\text{ke}(K^X)$ is disjunct from $\mathcal{P}_{s\perp}$, since the first and second eigenfunctions are constant and linear polynomials.
 - *Fourier*: eigenfunctions are those of a standard Fourier basis. Although a complete Fourier basis is a basis for all square-integrable functions, in practice the kernel $\text{ke}(K^X)$ of a truncated Fourier basis contains functions that are very close to the constant since no linear combination of a finite set of Fourier basis functions yields an exactly constant function, so $\text{ke}(K^X)$ is not disjunct from $\mathcal{P}_{s\perp}$. We used this basis with constant ($\nu_m \equiv 1$ for *Fourier, Const*) or exponentially decreasing (for *Fourier, Exp*) eigenvalues ν_m .
 - *Wiener*: eigenfunctions and eigenvalues are those of the standard Wiener process on $[0, 1]$:

$$\phi_m(s) = \sqrt{2} \sin(\pi(m-0.5)s); \nu_m = \left(\frac{\pi}{2}(2m+1)\right)^{-2}.$$

The $\text{ke}(K^X)$ is close to $\mathcal{P}_{s\perp}$ in this case as no linear combination of a finite set of these basis functions yields an exactly constant or linear function.

- *BrownBridge*: eigenfunctions and eigenvalues are those of the standard Brownian bridge on $[0, 1]$: $\phi_m(s) = \sqrt{2} \sin(\pi m s)$; $\nu_m = \frac{1}{\pi m}$. The $\text{ke}(K^X)$ is close to $\mathcal{P}_{s\perp}$ in this case as no linear combination of a finite set of these basis functions yields an exactly constant or linear function.
- *Poly(1+), Poly(2+), Poly(-1)*: eigenfunctions are orthogonal polynomials of degree 1 [2] to M [$M + 1$] for *Poly(1+)* [*Poly(2+)*], so that $\text{ke}(K^X)$ includes the (complete) null-space of the rank-deficient penalties, i.e., the constant [and linear] functions. *Poly(-1)* has polynomial eigenfunctions of degree $\{0, 2, 3, 4, \dots, M + 1\}$ so that $\text{ke}(K^X)$ overlaps the null-space of the second differences penalty but not the first differences penalty. All three processes are associated with linearly decreasing ν_m .

From top to bottom, these processes become increasingly more “antagonistic” in the sense that 1) the kernels of these eigenfunction systems move increasingly closer to the kernels of the penalties we consider and 2) more quickly decreasing eigenvalues result in lower effective rank of the observed $X(s)$.

- **coefficient functions** $\beta(s, t) = \mathbf{B}_s \Theta \mathbf{B}_t^T$ are drawn randomly for each setting. The associated coefficients are drawn as

$$\text{vec}(\Theta) \sim N\left(\mathbf{0}, (0.1\mathbf{I} + \mathbf{P}(\lambda_s, \lambda_t))^{-1}\right),$$

where $\mathbf{P}(\lambda_s, \lambda_t)$ is a first order difference tensor penalty matrix.

- The marginal B-spline bases B_s and B_t have either 4 or 8 basis functions for each direction.
- $\lambda_s = \lambda_t$ are either 0.1 or 1.

This generates coefficient surfaces of varying complexity and roughness. We do not fit models where the basis used to generate $\beta(s, t)$ is larger than that used to estimate $\beta(s, t)$ since that could introduce a distracting approximation error not relevant to the issues at hand.

In order to make results comparable across the different settings, we use the relative integrated mean squared errors $\text{rIMSE}_\beta = \frac{\int (\beta(s, t) - \hat{\beta}(s, t))^2 ds dt}{\int (\beta(s, t))^2 ds dt}$ and $\text{rIMSE}_Y = \frac{1}{N} \sum_{i=1}^N \frac{\int (\hat{Y}_i(t) - E(Y_i(t)))^2 dt}{\int (Y_i(t) - \bar{Y}_i(t))^2 dt}$, where $\bar{Y}_i(t)$ is the mean of $Y_i(t)$ over t .

4.2. Results

Identifiability The estimation accuracy for $\hat{Y}(t)$ (not shown) is excellent across the board even for the very noisy settings, with 90% of relative integrated

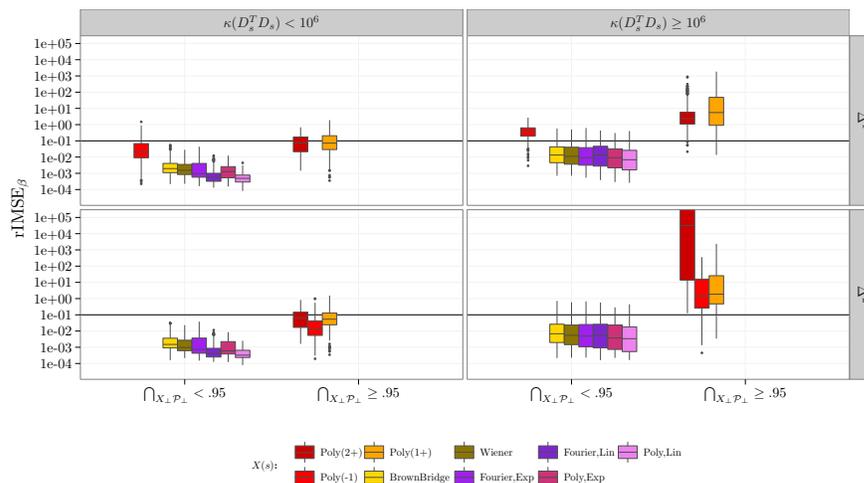


FIG 2. Boxplots for relative integrated mean square error rIMSE_{β} for all 18000 results for $\text{SNR}_{\epsilon} = 10$ with conventional difference penalties. Columns show results for full rank \mathbf{D}_s versus numerically rank deficient \mathbf{D}_s . Rows show results for the first and second order difference penalties. Boxplots grouped by overlap between $\text{ke}(K^X)$ and $\mathcal{P}_{s_{\perp}}$ as computed by $\cap_{X_{\perp} \mathcal{P}_{s_{\perp}}}$, color-coded for the different processes the $X(s)$ are sampled from. Vertical axis on \log_{10} -scale, extreme errors $> 10^5$ for Poly(2+) are cut off.

mean square errors below 0.01 and a median of 0.00051. Estimation accuracy for $\hat{\beta}(s, t)$, however, varies wildly over a range of 18 magnitudes between 1.3×10^{-8} and 2.2×10^{10} . Further analysis shows that the simulation study design succeeds in creating the identifiability issues described by the results in Section 3.2. To quantify the severity of identifiability issues, we compute rank correlations between rIMSE_{β} and rIMSE_{γ} over the 20 replicates of each simulation setting. As expected, we observe low or even negative correlations mostly for settings in which \mathbf{D}_s is rank-deficient. This effect increases both for lower signal-to-noise ratios and for more complex true shapes of $\beta(s, t)$. For intuition, consider that the “best” solution for (8) for any given error will be the smoothest surface (i.e., the one with the smallest penalty term) in the set of surfaces that can be generated by adding functions from $\text{ke}(K^X)$ to any initial $\beta(s, t)$ with the given error. This may be quite close or quite far from the true $\beta(s, t)$, depending on the specific setting, with more noisy data and more complex true shapes more likely to result in fits that are quite far from the truth and still producing good model fit.

Estimation performance for $\beta(s, t)$ Figure 2 shows the estimation errors for coefficient surfaces generated with $\text{SNR}_{\epsilon} = 10$. The right column shows results for numerically rank deficient \mathbf{D}_s (i.e., condition number $\kappa(\mathbf{D}_s^{\top} \mathbf{D}_s) \geq 10^6$), the left column for designs with $\kappa(\mathbf{D}_s^{\top} \mathbf{D}_s) < 10^6$. The top row shows results for first differences penalty, bottom row for second differences penalty. Boxplots are grouped by the amount of overlap between $\text{ke}(K^X)$ and $\mathcal{P}_{s_{\perp}}$ as computed by

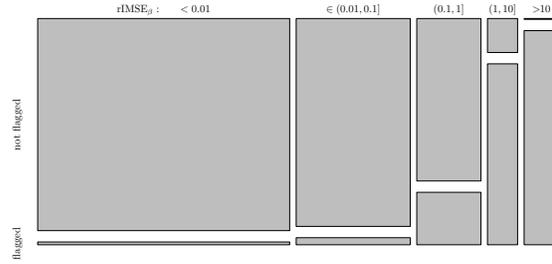


FIG 3. Mosaic plot for contingency table of “flagged” replicates and categorized $rIMSE_\beta$.

$\bigcap_{X_\perp \mathcal{P}_{s_\perp}}$ (see (9)), color-coded for the different processes the $X(s)$ -trajectories are sampled from. Results for $SNR_\epsilon = 2$ and 10^3 were qualitatively very similar – errors obviously become larger for noisier data but the pattern shown in Figure 2 remains the same. Note that relative estimation errors below ≈ 0.01 correspond to estimates that are visually indistinguishable from the true surfaces, and that errors below ≈ 0.1 (thick black horizontal line) usually preserve most essential features of the true $\beta(s, t)$ well. Results with $rIMSE_\beta > 1$ bear little resemblance to the “true” function. Also recall that all of these fits, with $rIMSE_\beta$ values varying by more than 14 magnitudes, resulted in a comparatively small range of $rIMSE_Y$ between 10^{-5} and 10^{-3} . Closer inspection of results shows that the extremely large errors for $Poly(1+)$, $Poly(-1)$ and $Poly(2+)$ are caused by the expected behavior: estimates are shifted by functions from the overlap of $ke(K^X)$ and \mathcal{P}_{s_\perp} . The top row of Figure 4 shows an example of this behavior: the estimate for first order difference penalty is shifted by a large constant, while the estimate for the second order difference penalty is shifted by both a constant and a huge linear trend in s -direction. The fitted values of all models shown in Figure 4 are practically identical.

These results mostly corroborate the results derived in Section 3 – we see that:

- Serious errors $rIMSE_\beta > 0.1$ are rare for both full-rank and rank-deficient \mathbf{D}_s if the generating process for $X(s)$ is not antagonistic in the sense that $ke(K^X) \cap \mathcal{P}_{s_\perp} = \{0\}$, i.e. for the *Poly*, *Fourier*, *Wiener* and *BrownBridge* processes.
- As long as $ke(K^X) \cap \mathcal{P}_{s_\perp} = \{0\}$ (approximated numerically by the criterion that $\bigcap_{X_\perp \mathcal{P}_{s_\perp}} < 0.95$), regularization of the estimated coefficient surface allows us to achieve good estimates even if the unpenalized regression model *per se* would not be identifiable due to rank deficiency of \mathbf{D}_s .
- The larger \mathcal{P}_{s_\perp} (top to bottom), and the closer $ke(K^X)$ is to \mathcal{P}_{s_\perp} (left to right in each group of boxplots), the larger the likelihood of estimates far from the truth and the larger the average $rIMSE_\beta$.

Diagnostics As in Figure 2, we use $\bigcap_{X_\perp \mathcal{P}_{s_\perp}}$ (see (9)) as a measure of the overlap between $ke(K^X)$ and \mathcal{P}_{s_\perp} . We consider a replicate to be “flagged” as problematic if both $\bigcap_{X_\perp \mathcal{P}_{s_\perp}} \geq 0.95$ and $\kappa(\mathbf{D}_s^T, \mathbf{D}_s) > 10^6$.

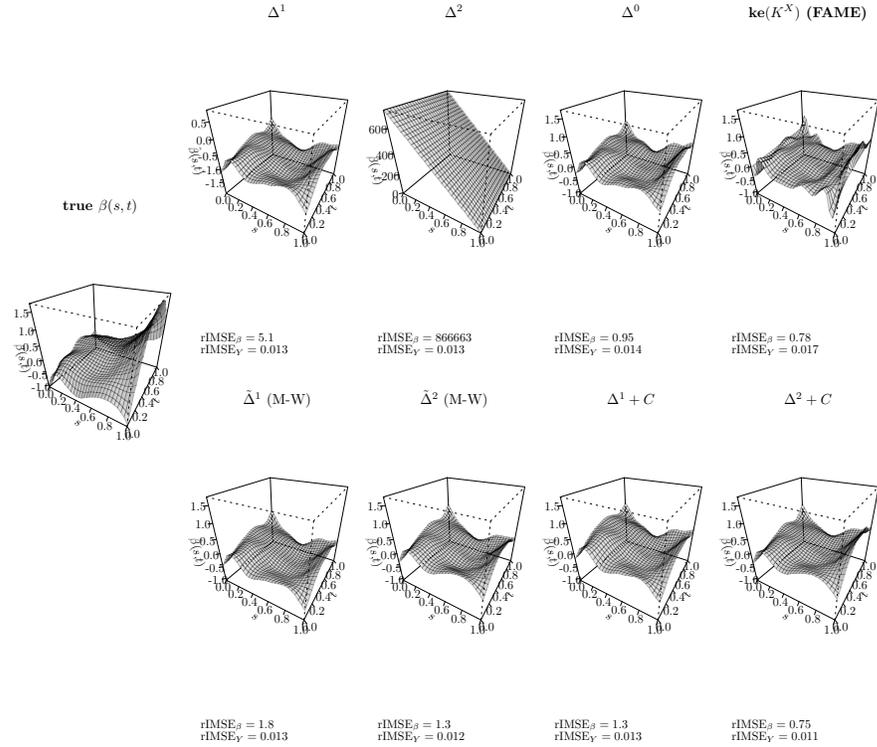


FIG 4. Example estimates for different penalties in a very difficult setting with $X(s)$ from Poly(2+), $M = 5$, $K_s = 12$, $\text{SNR}_\varepsilon = 2$. Center left panel: true coefficient surface. Top row, left to right: first order difference penalty, second order difference penalty, ridge penalty, FAME penalty. Bottom row, left to right: full-rank first and second order difference penalties, first and second order difference penalties with constraints. Note the different z-axis scales in the two left panels of the top row. Subtitles give $r\text{IMSE}_\beta$ and $r\text{IMSE}_\gamma$ for each fit.

Figure 3 shows a mosaic plot of the contingency table of “flagged” replicates and categorized $r\text{IMSE}_\beta$. While the sensitivity for identifying replicates with $r\text{IMSE}_\beta > 1$ is 0.92, the specificity is only 0.28. The sensitivity for identifying replicates with $r\text{IMSE}_\beta > 0.1$ is 0.58, the specificity is 0.09. These fairly low specificities indicate that the penalized approach to function-on-function regression discussed here can outperform theoretical expectations and frequently finds good solutions even in very difficult settings. Total accuracy for identifying settings with $r\text{IMSE}_\beta > 0.1$ is 0.88 and 0.94 for $r\text{IMSE}_\beta > 1$. The positive predictive value (precision) of the criterion for $r\text{IMSE}_\beta > 1$ is 0.72, while it is 0.91 for $r\text{IMSE}_\beta > 0.1$. The negative predictive value of the criterion for $r\text{IMSE}_\beta > 1$ is 0.99, while it is 0.87 for $r\text{IMSE}_\beta > 0.1$. Also note that certainly not all errors in the (0.1, 1]-range are due to identifiability issues, so not all of the (rare) non-detections are failures of the criterion.

Performance of modified penalties We broadened the scope of our simulation study by additionally comparing the performance of the non-standard penalties and constraints introduced in Section 3.3:

- a full-rank ridge penalty (“ Δ^0 ”),
- the modified full-rank roughness penalties as suggested by Marra and Wood (2011); in our case we used both full-rank first order differences penalties (“ $\tilde{\Delta}^1$ ”) and full-rank second order differences penalties (“ $\tilde{\Delta}^2$ ”),
- the FPC penalty of James and Silverman (2005) (“ $\text{ke}(K^X)$ (FAME)”). We replace $\hat{\nu}_m$ by $\max(\hat{\nu}_m, 10^{-10}\hat{\nu}_1)$ in order to remove any (numerically) zero or negative eigenvalues,
- and conventional first and second order differences penalties with additional constraints (“ $\Delta^1 + C$ “, “ $\Delta^2 + C$ “) that enforce orthogonality of the estimated coefficient surface to functions in $\text{ke}(K^X) \cap \mathcal{P}_{s\perp}$ if the data are flagged as problematic by our diagnostic criterion ($\kappa(D_s^T D_s) \geq 10^6$ and $\bigcap_{X\perp \mathcal{P}_{s\perp}} > .95$). Note that $\Delta^1 + C$ and $\Delta^2 + C$ reduce to Δ^1 and Δ^2 , respectively, if $\text{ke}(K^X) \cap \mathcal{P}_{s\perp} = \emptyset$ and need not be compared to the other techniques in that case.

Figure 5 shows the rIMSE_β for $\text{SNR}_\epsilon = 10$ for the different $X(s)$ -processes and penalties. Note that, in contrast to Figure 2, colors now represent the different penalties, not the different $X(s)$ -processes. Boxplots for Δ^1 and Δ^2 contain the same results as those shown in Figure 2.

The full-rank difference penalties $\tilde{\Delta}^1$ and $\tilde{\Delta}^2$ (cyan and turquoise boxes) seem to drastically reduce the size and likelihood of severe estimation errors in the difficult settings, especially compared to Δ^2 (dark blue). We occasionally incur slightly worse estimates for the easier settings, but these differences are small and hardly relevant here. However, we have found that neither $\tilde{\Delta}^1$ nor $\tilde{\Delta}^2$ with $\epsilon = .1$ are effective in removing non-identifiability artifacts if the estimated smoothing parameter is very small, i.e., if the effect of the penalty on the fit is weak. This explains the large outliers observed for some replicates for the antagonistic X -processes (top row) for $\tilde{\Delta}^1$ and $\tilde{\Delta}^2$. Figure 8 (bottom row, leftmost two panels) shows this for the DTI data set, where the estimated smoothing parameters are small. Note that no such outliers occur for the ridge (Δ^0) and FAME penalties (green boxes). However, both penalties are not competitive for numerically rank deficient D_s for the non-pathological $X(s)$ -processes in the bottom two rows and perform worse than the (modified) difference penalties in all other settings, and so cannot be recommended for general use as well. Furthermore, in the simulated settings we use, the true $\beta(s, t)$ are centered around 0, which reduces the negative effect of these penalties’ bias towards small absolute values of $\beta(s, t)$.

Figure 6 shows results only for simulated data sets flagged as problematic by our diagnostic criterion for Δ^1 or Δ^2 . In such settings, it makes sense to use difference penalties combined with “orthogonal-to-null-space-overlap” constraints on the fitted surface, these results are denoted by $\Delta^1 + C$ or $\Delta^2 + C$ (violet and purple), respectively.

Figure 4 shows illustrative exemplary fits for the modified penalties and penalties with constraints in a data setting with $\kappa(D_s^T D_s) \geq 10^6$. in the two

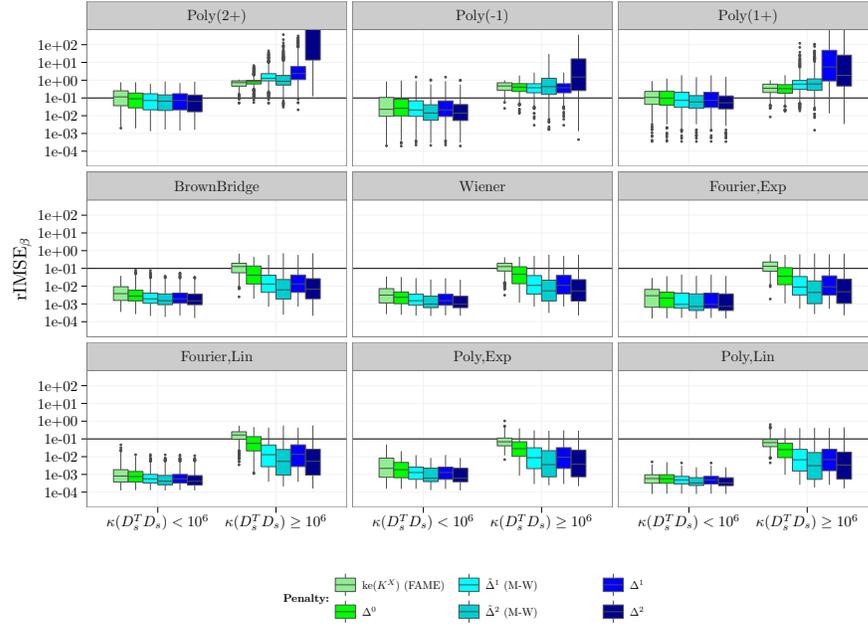


FIG 5. $rIMSE_\beta$ for $SNR_\epsilon = 10$ for the 8 $X(s)$ -processes (panels) and the 8 different penalties (color). Separate boxplots in each panel for settings with numerically rank deficient D_s ($\kappa(D_s^T D_s) \geq 10^6$) versus settings with full rank D_s .

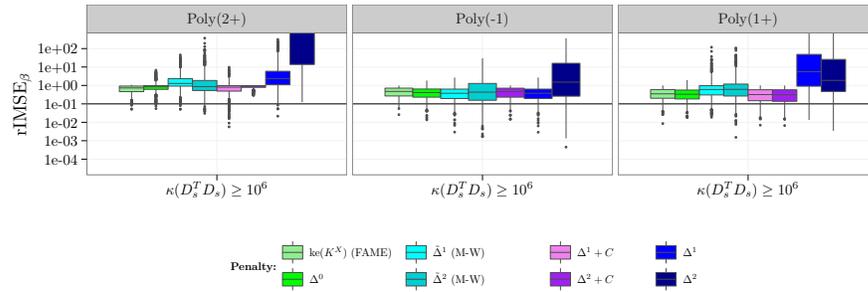


FIG 6. $rIMSE_\beta$ for $SNR_\epsilon = 10$ only for “pathological” data-sets. $X(s)$ -processes in panels, penalties coded by color.

rightmost panel of the top row and the panels in the bottom row. While $rIMSE_Y$ is very similar for all 8 penalties except FAME (top right), the shapes of the corresponding coefficient surface estimates differ from each other in the expected fashion: The ridge penalty’s bias towards small $|\beta(s, t)|$ and tendency to under-smooth the estimated surface is visible, as is the typical wiggleness of FAME fits which are also not roughness penalized in the conventional sense. Compared to

the result for Δ^1 , the fit for $\tilde{\Delta}^1$ is much closer to the level of the true $\beta(s, t)$ and adds a much smaller (spurious) constant to the fit. Similar remarks apply for the comparison between Δ^2 and $\tilde{\Delta}^2$: The huge spurious linear trend in s is almost completely removed. The overlap criterion $\bigcap_{X \perp \mathcal{P}_{s \perp}}$ is ≈ 1 for Δ^1 , ≈ 2 for Δ^2 and 0 by construction for all others. Both constrained fits ($\Delta^1 + C$, $\Delta^2 + C$, two bottom right panels) completely suppress the spurious constant (and trend) present in the unconstrained fits (two top right panels), and $\Delta^2 + C$ yields a slightly smoother fit than $\Delta^1 + C$ as expected.

Note that the true $\beta(s, t)$ are centered around 0 in our simulation. As all modified penalties and penalties with additional constraints result in estimated surfaces around 0, results will be shifted up- or downwards compared to the true surfaces if this is not the case in a real application. Such a shift reflects the fact that the average level of the surface cannot be inferred from the data in settings with corresponding persistent non-identifiability. Thus, in cases where our diagnostics indicate persistent non-identifiability, estimated coefficient surface values can only be interpreted relative to one another, but no interpretation of their sign is possible due to the estimated absolute level being essentially arbitrary. A corresponding implication regarding linear trends applies to the case of second order difference penalties.

Summary of simulation results We can draw the following conclusions based on the entirety of simulation results:

- the potential for extreme estimation errors is large for $X(s)$ -processes with low effective rank whose $\text{ke}(K^X)$ is not disjunct from $\mathcal{P}_{s \perp}$.
- there is no strong positive correlation between accuracy of fitted values (rIMSE_Y) and accuracy of the estimated coefficient surface (rIMSE_β) for rank deficient \mathbf{D}_s . Even extremely wrong estimates of $\beta(s, t)$ can yield good model fits.
- calculating $\bigcap_{X \perp \mathcal{P}_{s \perp}}$ and $\kappa(\mathbf{D}_s^T \mathbf{D}_s)$ yields a suitable criterion that can diagnose persistent non-identifiability reliably, albeit with a substantial rate of “false alarms” in which conventional difference penalties work well despite an antagonistic data situation.
- the full-rank roughness penalties often stabilize estimates in persistent non-identifiability settings without diminishing accuracy in other settings by a relevant amount, but not reliably so.
- for problematic data settings flagged by the criterion developed here, the combination of difference penalties with suitable constraints on the coefficient surface (“ $\Delta + C$ ”), ridge penalties (“ Δ^0 ”) and penalties based on the spectrum of the functional covariate (“FAME”) all perform similarly and outperform both full-rank and conventional difference penalties.

5. Case study

To illustrate the practical relevance of the theoretical development given in Section 3 and the performance of the countermeasures developed therein, we de-

scribe a deliberately constructed but realistic setting in which non-identifiability persists despite penalization. We use a subset of 100 patients from the DTI-data described in the introduction and fit a realistic model in which investigators try to separate the effect of mean CCA-FA levels on RCST-FA from that of the shape of CCA-FA along the tract. Such a model can be defined as

$$Y_i(t) = \beta_0(t) + \bar{x}_i \beta_1(t) + \int_{\mathcal{S}} X_i^c(s) \beta(s, t) ds + \epsilon_{it}; \quad \epsilon_{it} \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2),$$

where $Y_i(t)$ is RCST-FA for patient i , $\beta_0(t)$ is a global functional intercept, $\bar{x}_i = \overline{X_i(s)}$ are the mean CCA-FA levels with associated effect $\beta_1(t)$, and $X_i^c(s) \approx X_i(s) - \bar{x}_i$ are the denoised, curve-wise centered CCA-FA curves. In this example, we use a reconstruction $X_i^c(s)$ based on the six largest empirical FPCs (ca. 90 % of total variance explained) to denoise CCA-FA and interpolate missing values. Note that curve-specific mean centering causes all constant functions to be in the kernel of the covariance of $X_i^c(t)$, i.e. $\bigcap_{X \perp \mathcal{P}_{s \perp}} \geq .95$ for first and second order difference penalties. All fits shown below use $K_s = K_t = 12$ marginal basis functions, so by construction we expect $\kappa(\mathbf{D}_s^T \mathbf{D}_s)$ to be large since the rank of the functional covariate is at most $M = 6$. More than 6 marginal basis functions are required, however, so that the basis for $\beta(s, t)$ is flexible enough for approximating the rather complex shape we observe here. Note that curve-wise centering is also a necessary pre-processing step for many spectroscopic data analyses where absorption spectra are used as functional covariates and different mean intensity levels of such spectra are often pure laboratory artifacts (Fuchs et al., 2015), and more generally in settings such as this one where practitioners try to separate effects of the mean level of a functional covariate from those of its shape.

From left to right, the top row of Figure 7 shows the responses and fitted values for this model, the uncentered functional covariates used for the models shown in Figure 1, and the centered functional covariates used for the models in this section. For uncentered covariates, this setting provides an example of *simple* non-identifiability, while centered covariates induce *persistent* non-identifiability for difference penalties. Figure 8 shows the estimated coefficient surfaces $\hat{\beta}(s, t)$ for this model for various penalties. Despite the divergent shapes of the different coefficient surface estimates, all 8 fits lead to practically identical fitted values on the training data (100 patients) and practically identical predictions for the test set (155 patients), with integrated MSE ≈ 0.0031 and predictive integrated MSE ≈ 0.0046 for each model. Both $\hat{\beta}_0(t)$ and $\hat{\beta}_1(t)$ are also practically identical for all 8 models, they are shown in the bottom row of Figure 7 along with the estimated contributions $\int_{\mathcal{S}} X_i^c(s) \hat{\beta}(s, t) ds$ of the functional covariate to the additive predictor.

The model specification is appropriately flagged as causing *persistent* non-identifiability (i.e., $\bigcap_{X \perp \mathcal{P}_{s \perp}} = .98$ [1.15] and $\kappa(\mathbf{D}_s^T \mathbf{D}_s) \rightarrow \infty$) for first [second] order difference penalties Δ^1 [Δ^2]. It is instructive to compare the resulting estimates for $\beta(s, t)$ for different penalty specifications: the first difference penalty fit (top, left) includes a constant offset from 0, while the second difference penalty fit (top, second from left) includes different offsets from 0 for

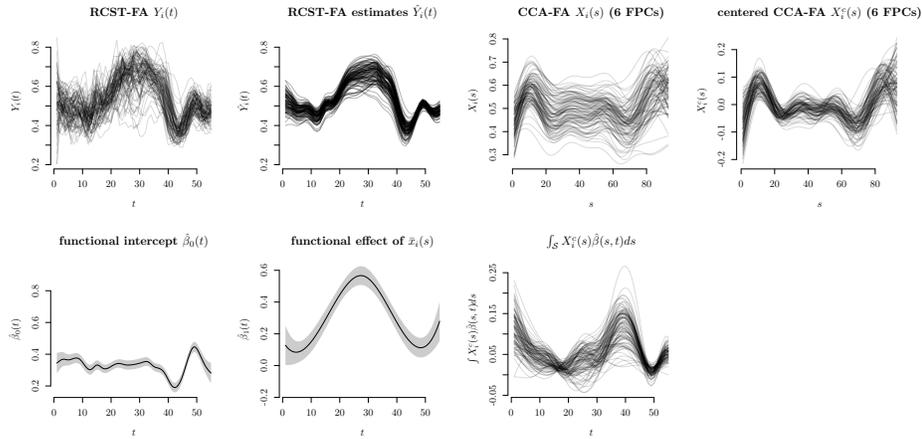


FIG 7. Top row, left to right: Observed RCST-FA $Y_i(t)$; estimated RCST-FA $\hat{Y}_i(t)$ for the model described in this Section; presmoothed CCA-FA based on 6 largest FPCs as used in the example shown in Figure 1; presmoothed, curve-wise centered CCA-FA based on 6 largest FPCs as used in the example below. Bottom row, left to right: Estimated functional intercept $\hat{\beta}_0(t)$ and estimated functional effect of mean CCA-FA $\hat{\beta}_1(t)$, with approximate point-wise 95% confidence intervals; estimated contributions of the functional effect $\int_{\mathcal{S}} X_i^c(s) \hat{\beta}(s, t) ds$ to the additive predictor.

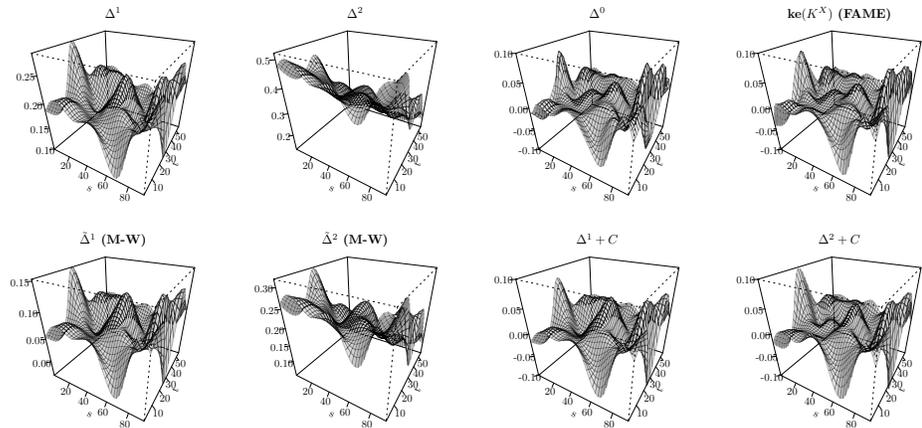


FIG 8. Estimated coefficient surfaces $\hat{\beta}(s, t)$ for regressing RCST-FA on mean CCA-FA and centered CCA-FA truncated to its first 6 empirical FPCs. All 8 fits lead to practically identical fitted values. Top row from left: First order (Δ^1), second order difference penalty (Δ^2), ridge penalty (Δ^0), FPC-kernel-based penalty (FAME). Bottom row from left: Full rank first order ($\hat{\Delta}^1$ (M-W)), second order difference penalty ($\hat{\Delta}^2$ (M-W)), first order with constraint on kernel overlap ($\Delta^1 + C$), second order with constraint on kernel overlap ($\Delta^2 + C$). All fits performed with `pffr()`. Note that z-axis limits for the 4 panels on the left differ from each other and from those of the 4 panels on the right.

each t , in a linear decrease that is unpenalized by the marginal second order difference penalty in t direction, as well as a linear trend in s . The presence of this spurious linear trend is surprising since we did not explicitly remove linear components from the functional covariate and it is not present in the fits for the uncentered data (c.f. Figure 1). Further investigation revealed that the centered functional covariate’s null-space contains a fairly strong linear component ($\langle \bigcap_{LV}((\mathbf{X}^c)^\perp), \mathbf{w} \cdot \mathbf{s} \rangle = 0.77$, where $(\mathbf{X}^c)^\perp$ denotes the orthogonal complement of the observed $X^c(s)$), which is overlap enough to produce spurious linear shifts in this example. Note that the linear component is much stronger in the non-centered covariates ($\langle \bigcap_{LV}((\mathbf{X})^\perp), \mathbf{w} \cdot \mathbf{s} \rangle = 0.11$) and its lack is caused by the curve-wise centering in this example. It is straightforward to show that this can occur if $\int X_i(s) s ds \approx \bar{x}_i \int s ds \forall i = 1, \dots, n$, as is the case here.

Both offset and trends are reduced, but not entirely removed, by instead using the full-rank difference penalties with $\epsilon = 0.1$ (bottom row, two leftmost panels). Due to the wiggleness of the coefficient surface, the estimated smoothing parameters are quite small, so there is only little penalization going on. Consequently, the additional small penalty on \mathcal{P}_{s^\perp} is not strong enough to eliminate non-identifiability artifacts from the fit in this case. Estimated coefficient surfaces with full-rank difference penalties with $\epsilon = 1$ (not shown), however, are very similar to the those in the four right-most panels of Figure 8. That results for the full-rank difference penalties depend so strongly on this tuning parameter is a considerable disadvantage.

The four right-most panels show results based on the ridge penalty (Δ^0 , second from right, top row), the penalty based on the FPCs of $\text{ke}(K^X)$ (FAME, top right), and the difference penalties with additional “orthogonal-to-kernel-overlap” constraints ($\Delta + C$, bottom row). It is reassuring to see that all four of these penalties lead to very similar results in this setting despite their different motivations and mathematical properties. Admittedly, understanding the estimated coefficient surfaces in terms of the supposed data generating process remains difficult. First, because of their complex shape and secondly (and more pertinently to the main points of this paper), because the restriction to surfaces centered around 0 that is enforced for $\Delta + C$ and implied by the Δ^0 and $\text{ke}(K^X)$ penalties precludes facile interpretation of, e.g., the sign of features of $\beta(s, t)$. Since the “true” offset of the coefficient surface is not estimable from the data, negativity or positivity of certain peaks or troughs of $\hat{\beta}(s, t)$ is not directly interpretable and interpretation can thus only rely on relative heights across the surface. Higher dimensional overlaps between penalty and covariate null-spaces than the one encountered here will compound these expositional difficulties.

This example demonstrates how reasonable, but unfortunate combinations of data pre-processing and model specifications can lead to simply as well as persistently non-identifiable models despite penalization. The diagnostics and countermeasures developed in Section 3, however, seem to be suitable for detecting and remedying such problems in a real data setting.

6. Conclusion and discussion

Coefficient surface estimates in spline-based function-on-function-regression (1) can suffer from persistent identifiability problems if the span of the marginal basis for the coefficient surface over a functional covariate's domain overlaps the kernel of its covariance operator. A rank deficient design matrix can occur in particular if the functional covariate's covariance operator is of effective rank smaller than the number of marginal basis functions – either because the number of eigenfunctions with non-zero eigenvalues is truly below the number of marginal basis functions, or if eigenvalues of the covariance operator decrease too rapidly compared to the noise level of the data.

In practice, spline based approaches are typically fitted with a regularization penalty corresponding to a smoothness assumption on the coefficient surface. We have shown that identifiability problems persist if, and only if, in addition to a rank deficiency of the design matrix, the kernel of the functional predictor's covariance $\text{ke}(K^X)$ overlaps the function space $\mathcal{P}_{s\perp}$ spanned by parameter vectors in the null-space of the spline's roughness penalty. In the case of no overlap, there is a unique smoothest representative on any hyperplane of parameter vectors leading to the same additive predictor and the penalized estimation problem finds a unique smoothest solution. Similar results hold for the simpler case of penalized scalar-on-function regression models (Happ, 2013). They are also expected to hold for more general loss functions than the quadratic loss analysed here since generalized additive models are typically estimated by the penalized iteratively re-weighted least squares (P-IWLS) method, where the system of equations solved in each step is identical to that of (8) except for the introduction of a vector of IWLS weights and the substitution of \mathbf{y} by IWLS working responses that are element-wise linear transformations of \mathbf{y} .

A lack of identifiability also implies a lack of correlation between accuracy of the coefficient estimates and goodness of fit for the responses. As this extends to prediction errors for out-of-sample data from the same process, it is usually not possible to detect identifiability issues for a given data set based on subsampling or cross-validation schemes. Instead, based on theoretical considerations and simulation results, we have identified two easily computable diagnostic criteria in order to detect non-identifiable model specifications before estimation. The criteria combine the condition number of a partial design matrix with a measure of the amount of overlap between the kernel of the functional predictor's covariance and the null-space of the penalty. Non-identifiability may in particular be an issue if both criteria are indicative of a problematic setting, or if the partial design matrix is numerically rank deficient and the penalty smoothing parameter is estimated to be close to zero. If a non-identifiable model specification is discovered, we recommend that practitioners choose modified full-rank roughness penalties to safeguard against spurious estimates or estimate coefficient surfaces under constraints that force these spurious components to zero. The `pffr`-function in the `refund` package uses a first order differences penalty by default, incorporates the diagnostic checks developed and evaluated in this work, and issues corresponding user warnings and en-

forces suitable constraints if a persistently non-identifiable model specification is detected.

Another practical consequence of our results is that pre-processing methods for functional covariates should be avoided if they reduce their effective rank (such as pre-smoothing with low-dimensional bases) or if they increase the amount of overlap between $\text{ke}(K^X)$ and $\mathcal{P}_{s\perp}$ (such as curve-wise centering).

Jointly, these provisions seem to be sufficient to diagnose and safeguard against most serious artifacts of non-identifiability in practice. Our results indicate that in many cases, penalization allows the reasonable estimation of coefficient surfaces that are not identifiable in the theoretical model under an additional smoothness assumption, avoiding instead the common assumption that the estimated coefficient surface lies in the span of the first few eigenfunctions of the covariance operator of the covariate.

This work drives home the point that we cannot hope to reliably estimate arbitrarily complex effect shapes from functional covariates with low information content. In that sense, assuming smoothness of the coefficient surface and constraining its non-identifiable components to be zero is simply following a principle of parsimony. At the same time, substantial interpretation of coefficient surface estimates derived from rank-deficient designs is difficult and has the potential to be very misleading. Functional principal component regression approaches do not suffer from the potential identifiability issues discussed here, but they do so at the price of restricting the estimated coefficient surface to the span of the estimated functional principal components. These are significant challenges for the maturing field of functional regression methods, at least in applications where these methods are used not only for prediction, but also for inferring and understanding the underlying data generating processes. It is our hope that the theoretical development and practical examples presented here can serve as a starting point for critical reflection on this important issue.

Appendix A: Proofs

A.1. Proof of Proposition 3.2

Assume that B_t is of full rank K_t . Then, the design matrix $D = B_t \otimes (XWB_s)$ in model (7) is rank-deficient if and only if

- a) $M < K_s$ or
- b) if $M \geq K_s$, but $\text{rank}(\Phi WB_s) < K_s$.

Proof. As $\text{rank}(X) = \text{rank}(\Xi\Phi) = \text{rank}(\Phi) = M$,

$$\begin{aligned} \text{rank}(\Phi WB_s) &\geq \text{rank}(\Xi\Phi WB_s) = \text{rank}(XWB_s) \\ &\geq \text{rank}(\Xi\Phi) + \text{rank}(\Phi WB_s) - \text{rank}(\Phi) = \text{rank}(\Phi WB_s) \end{aligned}$$

using Harville (1997, Th. 17.5.1). Thus, $\text{rank}(XWB_s) = \text{rank}(\Phi WB_s)$. The rank will be less than full if $\text{rank}(\Phi WB_s) < K_s$, including if $M < K_s$, as $\text{rank}(\Phi WB_s) \leq M$ by construction. \square

A.2. Proof of Proposition 3.3

Let $\mathbf{P} = \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s})$, with \mathbf{P}_s and \mathbf{P}_t positive semi-definite matrices. Assume that \mathbf{B}_t is of full rank K_t , that $\text{rank}(\mathbf{P}_t) < K_t$ and that $\lambda_s > 0, \lambda_t \geq 0$. Then, for any $\mathbf{f} \in \text{im}(\mathbf{D})$ there is a unique minimum $\min_{\boldsymbol{\theta} \in \mathcal{H}_f} \{\boldsymbol{\theta}^\top \mathbf{P} \boldsymbol{\theta}\}$ if and only if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$.

Proof. We have

$$\begin{aligned} & \min \{\boldsymbol{\theta}^\top \mathbf{P} \boldsymbol{\theta}\} && \text{s.t. } \boldsymbol{\theta} \in \mathcal{H}_f \\ & = \min \{\boldsymbol{\theta}^\top \mathbf{U} \mathbf{U}^\top \mathbf{P} \mathbf{U} \mathbf{U}^\top \boldsymbol{\theta}\} && \text{s.t. } \mathbf{D} \boldsymbol{\theta} = \mathbf{f} \\ & = \min \left\{ (\boldsymbol{\theta}^\top \mathbf{U}_+ | \boldsymbol{\theta}^\top \mathbf{U}_0) \begin{pmatrix} \mathbf{U}_+^\top \mathbf{P} \mathbf{U}_+ & \mathbf{U}_+^\top \mathbf{P} \mathbf{U}_0 \\ \mathbf{U}_0^\top \mathbf{P} \mathbf{U}_+ & \mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0 \end{pmatrix} \begin{pmatrix} \mathbf{U}_+^\top \boldsymbol{\theta} \\ \mathbf{U}_0^\top \boldsymbol{\theta} \end{pmatrix} \right\} && \text{s.t. } \mathbf{U}_+^\top \boldsymbol{\theta} = \boldsymbol{\Sigma}_+^{-1} \mathbf{V}_+^\top \mathbf{f}. \end{aligned}$$

Denote $\mathbf{v}_+ = \mathbf{U}_+^\top \boldsymbol{\theta}$ and $\mathbf{v}_0 = \mathbf{U}_0^\top \boldsymbol{\theta}$, with $(\mathbf{v}_+^\top, \mathbf{v}_0^\top)^\top = \mathbf{U}^\top \boldsymbol{\theta}$ a bijective re-parametrization of $\boldsymbol{\theta}$. Note that \mathbf{v}_+ is fixed while \mathbf{v}_0 is free to vary within the hyperplane. Setting the derivative with respect to \mathbf{v}_0 equal to zero yields

$$\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0 \mathbf{v}_0 = -\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_+ \mathbf{v}_+. \quad (10)$$

Now, if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) \neq \{\mathbf{0}\}$, choose $\mathbf{u}_s \neq \mathbf{0}$ with $\mathbf{U}_{s0} \mathbf{u}_s \in \text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s)$ and $\mathbf{u}_t \neq \mathbf{0}$ with $\mathbf{U}_t \mathbf{u}_t \in \text{ke}(\mathbf{P}_t)$. As $\mathbf{U}_0 = \mathbf{U}_t \otimes \mathbf{U}_{s0}$, we have

$$\begin{aligned} (\mathbf{u}_t^\top \otimes \mathbf{u}_s^\top) \mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0 (\mathbf{u}_t \otimes \mathbf{u}_s) &= \\ \lambda_s (\mathbf{u}_t^\top \mathbf{u}_t) \mathbf{u}_s^\top \mathbf{U}_{s0}^\top \mathbf{P}_s \mathbf{U}_{s0} \mathbf{u}_s + \lambda_t \mathbf{u}_t^\top \mathbf{U}_t^\top \mathbf{P}_t \mathbf{U}_t \mathbf{u}_t (\mathbf{u}_s^\top \mathbf{u}_s) &= 0. \end{aligned}$$

Thus, $\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0$ is not of full rank, no unique solution \mathbf{v}_0 of (10) exists, so there is no unique minimum $\min_{\boldsymbol{\theta} \in \mathcal{H}_f} \{\boldsymbol{\theta}^\top \mathbf{P} \boldsymbol{\theta}\}$.

On the other hand, if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$, for any \mathbf{x} with $\mathbf{x}^\top \mathbf{U}_{s0}^\top \mathbf{P}_s \mathbf{U}_{s0} \mathbf{x} = 0$ we have $\mathbf{U}_{s0} \mathbf{x} \in \text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$ and thus $\mathbf{x} = \mathbf{U}_{s0}^\top \mathbf{U}_{s0} \mathbf{x} = \mathbf{0}$. As this means that $\mathbf{U}_{s0}^\top \mathbf{P}_s \mathbf{U}_{s0}$ is positive definite, we also have that $\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0 = \lambda_s \mathbf{I}_{K_t} \otimes (\mathbf{U}_{s0}^\top \mathbf{P}_s \mathbf{U}_{s0}) + \lambda_t (\mathbf{U}_t^\top \mathbf{P}_t \mathbf{U}_t) \otimes \mathbf{I}_{(K_s - d/K_t)}$ is positive definite and thus invertible. Therefore, there is a unique minimum $\mathbf{v}_0 = -(\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0)^{-1} \mathbf{U}_0^\top \mathbf{P} \mathbf{U}_+ \mathbf{v}_+$ and a unique smoothest point

$$\boldsymbol{\theta}_f = \mathbf{U} (\mathbf{v}_+^\top, -[(\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0)^{-1} \mathbf{U}_0^\top \mathbf{P} \mathbf{U}_+ \mathbf{v}_+]^\top)^\top = \mathbf{H} \mathbf{U}_+ \mathbf{v}_+ = \mathbf{H} \boldsymbol{\theta}, \quad (11)$$

with $\mathbf{H} = (\mathbf{I}_{K_s K_t} - \mathbf{U}_0 (\mathbf{U}_0^\top \mathbf{P} \mathbf{U}_0)^{-1} \mathbf{U}_0^\top \mathbf{P})$ and $\boldsymbol{\theta}_f^\top \mathbf{P} \boldsymbol{\theta}_f = \min_{\boldsymbol{\theta} \in \mathcal{H}_f} \boldsymbol{\theta}^\top \mathbf{P} \boldsymbol{\theta}$. \square

A.3. Proof of Proposition 3.4

Assume that \mathbf{B}_t is of full rank K_t , that $\text{rank}(\mathbf{P}_t) < K_t$ and that $\lambda_s > 0, \lambda_t \geq 0$. Then, there is a unique penalized least squares solution for (8) if and only if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$.

Proof. Problem (8) has a unique solution iff $(\mathbf{D}^\top \mathbf{D} + \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s})) \geq 0$ is invertible, i.e., positive definite. Now, suppose that $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$. For any $\mathbf{x} \in \mathbb{R}^{K_t K_s}$ with

$$\mathbf{x}^\top (\mathbf{D}^\top \mathbf{D} + \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s})) \mathbf{x} = 0,$$

we have, with $\mathbf{b} = \mathbf{U}^\top \mathbf{x} = (b_{kj})_{kj \in \{11, 12, \dots, K_t K_s\}}$,

$$\begin{aligned} \mathbf{x}^\top \mathbf{D}^\top \mathbf{D} \mathbf{x} &= \mathbf{b}^\top \left(\boldsymbol{\Sigma}_t^2 \otimes \begin{pmatrix} \boldsymbol{\Sigma}_{s+}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \right) \mathbf{b} = 0 \\ \Rightarrow b_{kj} &= 0 \forall 1 \leq j \leq d/K_t; 1 \leq k \leq K_t, \end{aligned} \quad (12)$$

and also

$$\mathbf{b}^\top [\mathbf{I}_{K_t} \otimes ((\mathbf{U}_{s+} | \mathbf{U}_{s0})^\top \mathbf{P}_s (\mathbf{U}_{s+} | \mathbf{U}_{s0}))] \mathbf{b} \stackrel{(12)}{=} \tilde{\mathbf{b}}^\top [\mathbf{I}_{K_t} \otimes (\mathbf{U}_{s0}^\top \mathbf{P}_s \mathbf{U}_{s0})] \tilde{\mathbf{b}} = 0,$$

where $\tilde{\mathbf{b}}$ is obtained by removing the zero entries given by (12) which correspond to \mathbf{U}_{s+} from \mathbf{b} . Thus, for all $1 \leq k \leq K_t$, and letting $\tilde{\mathbf{b}}_k = (b_{k(d/K_t+1)}, \dots, b_{kK_s})$, we have $\mathbf{U}_{s0} \tilde{\mathbf{b}}_k \in \text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) = \{\mathbf{0}\}$. Thus, $\tilde{\mathbf{b}} = \mathbf{0}$, $\mathbf{b} = \mathbf{0}$, $\mathbf{x} = \mathbf{0}$ and $(\mathbf{D}^\top \mathbf{D} + \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s}))$ is of full rank.

On the other hand, if $\text{ke}(\mathbf{D}_s^\top \mathbf{D}_s) \cap \text{ke}(\mathbf{P}_s) \neq \{\mathbf{0}\}$, there is a $\mathbf{u}_s \neq \mathbf{0}$ with $\mathbf{D}_s \mathbf{u}_s = \mathbf{0}$ and $\mathbf{P}_s \mathbf{u}_s = \mathbf{0}$. Choose $\mathbf{0} \neq \mathbf{u}_t \in \text{ke} \mathbf{P}_t$. Then,

$$\begin{aligned} (\mathbf{u}_t^\top \otimes \mathbf{u}_s^\top) (\mathbf{D}^\top \mathbf{D} + \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s})) (\mathbf{u}_t \otimes \mathbf{u}_s) \\ = \mathbf{u}_t^\top \mathbf{B}_t^\top \mathbf{B}_t \mathbf{u}_t \cdot 0 + \lambda_s \mathbf{u}_t^\top \mathbf{u}_t \cdot 0 + \lambda_t \cdot 0 \cdot \mathbf{u}_s^\top \mathbf{u}_s = 0. \end{aligned}$$

Thus, $(\mathbf{D}^\top \mathbf{D} + \lambda_s(\mathbf{I}_{K_t} \otimes \mathbf{P}_s) + \lambda_t(\mathbf{P}_t \otimes \mathbf{I}_{K_s}))$ is singular and not invertible. \square

Acknowledgements

The authors wish to thank Ciprian Crainiceanu, Ludwig Fahrmeir, Jeff Goldsmith and Clara Happ for helpful discussions and encouragement. Comments by an anonymous reviewer spurred substantial improvements of our methodology and presentation. The DTI data were collected at Johns Hopkins University and the Kennedy-Krieger Institute. Financial support by the German Research Foundation through the Emmy Noether Programme, grant GR 3793/1-1 is gratefully acknowledged.

References

- H. CARDOT, F. FERRATY, and P. SARDA. Functional linear model. *Statistics and Probability Letters*, 45:11–22, 1999. [MR1718346](#)
- H. CARDOT, F. FERRATY, and P. SARDA. Spline estimators for the functional linear model. *Statistica Sinica*, 13(3):571–592, 2003. [MR1997162](#)
- J. M. CHIOU, H. G. MÜLLER, and J. L. WANG. Functional response models. *Statistica Sinica*, 14(3):675–694, 2004. [MR2087968](#)

- C. M. CRAINICEANU, A.-M. STAICU, and C.-Z. DI. Generalized multi-level functional regression. *Journal of the American Statistical Association*, 104:1550–1561, 2009. [MR2750578](#)
- M. FEBRERO-BANDE and M. OVIEDO DE LA FUENTE. Statistical computing in functional data analysis: The R package `fda.usc`. *Journal of Statistical Software*, 51(4):1–28, 2012.
- K. FUCHS, F. SCHEIPL, and S. GREVEN. Penalized scalar-on-functions regression with interaction term. *Computational Statistics & Data Analysis*, 81:38–51, 2015. [MR3257399](#)
- J. GOLDSMITH, J. BOBB, C.M. CRAINICEANU, B. CAFFO, and D. REICH. Penalized functional regression. *Journal of Computational and Graphical Statistics*, 20(4):830–851, 2011. [MR2878950](#)
- C. HAPP. Identifiability in scalar-on-functions regression. Master’s thesis, LMU München, 2013.
- D. A. HARVILLE. *Matrix Algebra from a Statistician’s Perspective*. Springer, 1997. [MR1467237](#)
- G. HE, H. G. MÜLLER, and J. L. WANG. Extending correlation and regression from multivariate to functional data. In M.L. PURI, editor, *Asymptotics in Statistics and Probability*, pages 301–315. VSP International Science Publishers, 2000.
- R. HERRICK. *WFMM*. The University of Texas M.D. Anderson Cancer Center, version 3.0 edition, 2013. URL https://biostatistics.mdanderson.org/SoftwareDownload/SingleSoftware.aspx?Software_Id=70
- L. HUANG, F. SCHEIPL, J. GOLDSMITH, J. GELLAR, J. HAREZLAK, M. W. MCLEAN, B. SWIHART, L. XIAO, C. M. CRAINICEANU, and P. T. REISS. *Refund: Regression with Functional Data*, 2015. URL <http://CRAN.R-project.org/package=refund>. R package version 0.1-12.
- A. E. IVANESCU, A.-M. STAICU, F. SCHEIPL, and S. GREVEN. Penalized function-on-function regression. *Computational Statistics*, 30(2):539–568, 2015. [MR3357075](#)
- G. M. JAMES. Generalized linear models with functional predictors. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):411–432, 2002. [MR1924298](#)
- G. M. JAMES and B. W. SILVERMAN. Functional adaptive model estimation. *Journal of the American Statistical Association*, 100(470):565–576, 2005. [MR2160560](#)
- R. LARSSON and M. VILLANI. A distance measure between cointegration spaces. *Economics Letters*, 70(1):21–27, 2001. [MR1798930](#)
- G. MARRA and S. N. WOOD. Practical variable selection for generalized additive models. *Computational Statistics & Data Analysis*, 55(7):2372–2387, 2011. [MR2786996](#)
- L. PRCHAL and P. SARDA. Spline estimator for functional linear regression with functional response. Unpublished, 2007. URL http://www.math.univ-toulouse.fr/staph/PAPERS/flm_prchal_sarda.pdf.
- J. O. RAMSAY and B. W. SILVERMAN. *Functional Data Analysis*. Springer, 2005. [MR2168993](#)

- J. O. RAMSAY, G. HOOKER, D. CAMPBELL, and J. CAO. Parameter estimation for differential equations: A generalized smoothing approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(5):741–796, 2007. [MR2368570](#)
- J. O. RAMSAY, H. WICKHAM, S. GRAVES, and G. HOOKER. *fda: Functional Data Analysis*, 2014. URL <http://CRAN.R-project.org/package=fda>. R package version 2.4.4.
- P. T. REISS and R. T. OGDEN. Functional principal component regression and functional partial least squares. *Journal of the American Statistical Association*, 102(479):984–996, 2007. [MR2411660](#)
- F. SCHEIPL, A.-M. STAICU, and S. GREVEN. Functional additive mixed models. *Journal of Computational and Graphical Statistics*, 24(2):477–501, 2015. [MR3357391](#)
- S. N. WOOD. Modelling and smoothing parameter estimation with multiple quadratic penalties. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(2):413–428, 2000. [MR1749600](#)
- S. N. WOOD. Low rank scale invariant tensor product smooths for generalized additive mixed models. *Biometrics*, 62, 2006. [MR2297673](#)
- Y. WU, J. FAN, and H. G. MÜLLER. Varying-coefficient functional linear regression. *Bernoulli*, 16(3):730–758, 2010. [MR2730646](#)
- N. T. YANG, H.-G. MÜLLER, and J.-L. WANG. *PACE: Principal Analysis by Conditional Expectation*, 2012. URL <http://www.stat.ucdavis.edu/PACE/>. MATLAB package version 2.16.
- F. YAO and H. G. MÜLLER. Functional quadratic regression. *Biometrika*, 97(1):49–64, 2010. [MR2594416](#)
- F. YAO, H. G. MÜLLER, and J. L. WANG. Functional linear regression analysis for longitudinal data. *The Annals of Statistics*, 33(6):2873–2903, 2005. [MR2253106](#)