



Chapter 15

Goodness-of-fit Tests for Functional Linear Models Based on Integrated Projections

Eduardo García-Portugués, Javier Álvarez-Liébana,
Gonzalo Álvarez-Pérez and Wenceslao González-Manteiga

Abstract Functional linear models are one of the most fundamental tools to assess the relation between two random variables of a functional or scalar nature. This contribution proposes a goodness-of-fit test for the functional linear model with functional response that neatly adapts to functional/scalar responses/predictors. In particular, the new goodness-of-fit test extends a previous proposal for scalar response. The test statistic is based on a convenient regularized estimator, is easy to compute, and is calibrated through an efficient bootstrap resampling. A graphical diagnostic tool, useful to visualize the deviations from the model, is introduced and illustrated with a novel data application. The R package `goffda` implements the proposed methods and allows for the reproducibility of the data application.

15.1 Functional Linear Models

15.1.1 Formulation

Given two separable Hilbert spaces \mathbb{H}_1 and \mathbb{H}_2 , we consider the regression setting with centered \mathbb{H}_2 -valued response \mathcal{Y} and centered \mathbb{H}_1 -valued predictor \mathcal{X} :

Eduardo García-Portugués (✉)

Department of Statistics and UC3M-Santander Big Data Institute, Carlos III University of Madrid, Avda. Universidad 30, 28911 Leganés, Spain, e-mail: edgarcia@est-econ.uc3m.es

Javier Álvarez-Liébana

Department of Statistics and Operations Research and Mathematics Didactics, University of Oviedo, C/ Federico García Lorca, 18, 33007 Oviedo, Spain, e-mail: alvarezljavier@uniovi.es

Gonzalo Álvarez-Pérez

Department of Physics, University of Oviedo, C/ Federico García Lorca, 18, 33007 Oviedo, Spain, e-mail: gonzaloalvarez@uniovi.es

Wenceslao González-Manteiga

Department of Statistics, Mathematical Analysis and Optimization, University of Santiago de Compostela, Rúa Lope Gómez de Marzoa s/n, 15782 Santiago de Compostela, Spain, e-mail: wenceslao.gonzalez@usc.es

$$\mathcal{Y} = m(\mathcal{X}) + \mathcal{E}, \quad (15.1)$$

where $m : \mathcal{X} \in \mathbb{H}_1 \mapsto \mathbb{E}[\mathcal{Y}|\mathcal{X} = \mathcal{X}] \in \mathbb{H}_2$ is the regression operator and the \mathbb{H}_2 -valued error \mathcal{E} is such that $\mathbb{E}[\mathcal{E}|\mathcal{X}] = 0$. When $\mathbb{H}_1 = L^2([a, b])$ and $\mathbb{H}_2 = L^2([c, d])$, the Functional Linear Model with Functional Response (FLMFR; see, e.g., [15, Chapter 16]) is the most well-known parametric instance of (15.1). If the regression operator is assumed to be Hilbert–Schmidt, m is parametrizable as

$$m_\beta(\mathcal{X}) = \int_a^b \beta(s, \cdot) \mathcal{X}(s) ds =: \langle\langle \beta, \mathcal{X} \rangle\rangle, \quad (15.2)$$

for $\beta \in \mathbb{H}_1 \otimes \mathbb{H}_2 = L^2([a, b] \times [c, d])$ a square-integrable kernel. The present work considers this framework and is concerned with the goodness-of-fit of the family of \mathbb{H}_2 -valued and \mathbb{H}_1 -conditioned linear models

$$\mathcal{L} := \{ \langle\langle \beta, \cdot \rangle\rangle : \beta \in \mathbb{H}_1 \otimes \mathbb{H}_2 \}. \quad (15.3)$$

Any $\mathcal{X} \in \mathbb{H}_1$ and $\mathcal{Y}, \mathcal{E} \in \mathbb{H}_2$ can be represented in terms of orthonormal bases $\{\Psi_j\}_{j=1}^\infty$ and $\{\Phi_k\}_{k=1}^\infty$ as $\mathcal{X} = \sum_{j=1}^\infty x_j \Psi_j$, $\mathcal{Y} = \sum_{k=1}^\infty y_k \Phi_k$, and $\mathcal{E} = \sum_{k=1}^\infty e_k \Phi_k$, where $x_j = \langle \mathcal{X}, \Psi_j \rangle_{\mathbb{H}_1}$, $y_k = \langle \mathcal{Y}, \Phi_k \rangle_{\mathbb{H}_2}$, and $e_k = \langle \mathcal{E}, \Phi_k \rangle_{\mathbb{H}_2}$, $\forall j, k \geq 1$. Also, $\beta \in \mathbb{H}_1 \otimes \mathbb{H}_2$ can be expressed as

$$\beta = \sum_{j=1}^\infty \sum_{k=1}^\infty b_{jk} (\Psi_j \otimes \Phi_k), \quad b_{jk} = \langle \beta, \Psi_j \otimes \Phi_k \rangle_{\mathbb{H}_1 \otimes \mathbb{H}_2}, \quad \forall j, k \geq 1.$$

Therefore, the population version of the FLMFR based on (15.2) can be expressed as

$$y_k = \sum_{j=1}^\infty b_{jk} x_j + e_k, \quad k \geq 1. \quad (15.4)$$

15.1.2 Model Estimation

The projection of (15.4) into the truncated bases $\{\Psi_j\}_{j=1}^p$ and $\{\Phi_k\}_{k=1}^q$ opens the way for the estimation of β given a centered sample $\{(\mathcal{X}_i, \mathcal{Y}_i)\}_{i=1}^n$. Indeed, the truncated sample version of (15.4) is expressed as

$$\mathbf{Y}_q = \mathbf{X}_p \mathbf{B}_{p,q} + \mathbf{E}_q, \quad (15.5)$$

where \mathbf{Y}_q and \mathbf{E}_q are $n \times q$ matrices with the respective coefficients of $\{\mathcal{Y}_i\}_{i=1}^n$ and $\{\mathcal{E}_i\}_{i=1}^n$ on $\{\Phi_k\}_{k=1}^q$, \mathbf{X}_p is the $n \times p$ matrix of coefficients of $\{\mathcal{X}_i\}_{i=1}^n$ on $\{\Psi_j\}_{j=1}^p$, and $\mathbf{B}_{p,q}$ is the $p \times q$ matrix of coefficients of β on $\{\Psi_j \otimes \Phi_k\}_{j,k=1}^{p,q}$.

Several estimators for β have been proposed; see, e.g., [16, 13, 5, 1, 14]. A popular estimation paradigm is Functional Principal Components Regression (FPCR; [15]), which considers the (empirical) Functional Principal Components (FPC) $\{\hat{\Psi}_j\}_{j=1}^p$

and $\{\hat{\Phi}_k\}_{k=1}^q$ as a plug-in for $\{\Psi_j\}_{j=1}^p$ and $\{\Phi_k\}_{k=1}^q$ underneath (15.5). Estimation by FPCR yields $\hat{\mathbf{B}}_{p,q} = \arg \min_{\mathbf{B}_{p,q}} \|\mathbf{Y}_q - \mathbf{X}_p \mathbf{B}_{p,q}\|^2 = (\mathbf{X}'_p \mathbf{X}_p)^{-1} \mathbf{X}'_p \mathbf{Y}_q$, with $j = 1, \dots, p$ and $k = 1, \dots, q$. The estimator $\hat{\mathbf{B}}_{p,q}$ depends on (p, q) and an automatic data-driven selection of (p, q) is of most practical interest. However, cross-validators are computationally expensive, especially since two tuning parameters must be optimized. A simple alternative for selecting q is to guarantee a certain proportion of explained variance (say, 0.99) for $\{\mathcal{Y}_i\}_{i=1}^n$. The more critical selection of p can be done by first ensuring a certain proportion of explained variance (say, 0.99) and then performing a LASSO-regularized FPCR regression (FPCR-L1 henceforth):

$$\hat{\mathbf{B}}_{p,q}^{(\lambda)} = \arg \min_{\mathbf{B}_{p,q}} \left\{ \frac{1}{2n} \sum_{i=1}^n \|(\mathbf{Y}_q)_i - (\mathbf{X}_p \mathbf{B}_{p,q})_i\|^2 + \lambda \sum_{j=1}^p \|(\mathbf{B}_{p,q})_j\| \right\},$$

where the notation $(\mathbf{A})_i$ stands for the i -th row of the matrix \mathbf{A} . This regularization applies a row-wise penalty that enables variable selection for a given λ , which can be efficiently selected by cross-validation and its *one standard error* variant [9].

However, FPCR-L1 lacks an explicit expression for the hat matrix (in contrast with FPCR), an important handicap for the bootstrap algorithm outlined in Section 15.2.3. To combine the flexible variable selection of FPCR-L1 with the analytical form of FPCR, we propose the FPCR-L1S estimator, which firstly implements FPCR-L1 for variable selection and then performs FPCR on the selected predictors. It returns the hat matrix $\mathbf{H}_C^{(\lambda)} = \tilde{\mathbf{X}}_{\tilde{p}} (\tilde{\mathbf{X}}'_{\tilde{p}} \tilde{\mathbf{X}}_{\tilde{p}})^{-1} \tilde{\mathbf{X}}'_{\tilde{p}}$, where $\tilde{\mathbf{X}}_{\tilde{p}}$ is the matrix of the coefficients of the \tilde{p} LASSO-selected predictors (not necessarily sorted).

Simulations [11, Section 2.4] report that FPCR-L1S outperforms FPCR.

15.2 Proposed Goodness-of-fit Tests

15.2.1 Test Statistic Genesis

Our aim is to test whether the regression operator belongs to the class of linear operators described in (15.3), that is, to test

$$\mathcal{H}_0 : m \in \mathcal{L} \quad \text{vs.} \quad \mathcal{H}_1 : m \notin \mathcal{L}.$$

To do so, we use the following lemma to characterize \mathcal{H}_0 in terms of the one-dimensional projections of \mathcal{Y} and \mathcal{X} . The lemma requires from analogues of the Euclidean $(p - 1)$ -sphere $\mathbb{S}^{p-1} := \{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{x}\| = 1\}$: the $(p - 1)$ -sphere of \mathbb{H}_1 for $\{\Psi_j\}_{j=1}^\infty$, $\mathbb{S}_{\mathbb{H}_1, \{\Psi_j\}_{j=1}^\infty}^{p-1} := \{\sum_{j=1}^p x_j \Psi_j \in \mathbb{H}_1 : \|\mathbf{x}\| = 1\}$ and, analogously, $\mathbb{S}_{\mathbb{H}_2, \{\Phi_k\}_{k=1}^\infty}^{q-1}$.

Lemma 1 (\mathcal{H}_0 characterization on finite-dimensional directions; [11]) *Let \mathcal{X} and \mathcal{Y} be \mathbb{H}_1 - and \mathbb{H}_2 -valued random variables, respectively, $\beta \in \mathbb{H}_1 \otimes \mathbb{H}_2$, and let $\{\Psi_j\}_{j=1}^\infty$ and $\{\Phi_k\}_{k=1}^\infty$ be bases of \mathbb{H}_1 and \mathbb{H}_2 , respectively. Then, the next statements are equivalent:*

- i. \mathcal{H}_0 holds, that is, $m(x) = \langle\langle x, \beta \rangle\rangle, \forall x \in \mathbb{H}_1$.
- ii. $\mathbb{E} \left[\langle\langle \mathcal{Y} - \langle\langle X, \beta \rangle\rangle, \gamma_{\mathcal{Y}}^{(q)} \rangle_{\mathbb{H}_2} \mathbb{1} \left\{ \langle\langle X, \gamma_X^{(p)} \rangle_{\mathbb{H}_1} \leq u \right\} \right] = 0$, for almost every $u \in \mathbb{R}, \forall \gamma_X^{(p)} \in \mathbb{S}_{\mathbb{H}_1, \{\Psi_j\}_{j=1}^{\infty}}^{p-1}, \forall \gamma_{\mathcal{Y}}^{(q)} \in \mathbb{S}_{\mathbb{H}_2, \{\Phi_k\}_{k=1}^{\infty}}^{q-1}$, and for all $p, q \geq 1$.

The reader is referred to [11] for the proof of the lemma.

We use the above characterization to detect deviations from \mathcal{H}_0 . We do so by means of the (p, q) -truncated empirical version of the doubly-projected integrated regression function in statement ii, that is, the residual marked empirical process

$$R_{n,p,q}(u, \gamma_X^{(p)}, \gamma_{\mathcal{Y}}^{(q)}) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \langle \hat{\mathcal{E}}_i^{(q)}, \gamma_{\mathcal{Y}}^{(q)} \rangle_{\mathbb{H}_2} \mathbb{1} \left\{ \langle X_i^{(p)}, \gamma_X^{(p)} \rangle_{\mathbb{H}_1} \leq u \right\}, \quad u \in \mathbb{R}, \quad (15.6)$$

with residual marks $\langle \hat{\mathcal{E}}_i^{(q)}, \gamma_{\mathcal{Y}}^{(q)} \rangle_{\mathbb{H}_2} = \hat{\mathbf{e}}'_{i,q} \mathbf{h}_q$ and jumps $\langle X_i^{(p)}, \gamma_X^{(p)} \rangle_{\mathbb{H}_1} = \mathbf{x}'_{i,p} \mathbf{g}_p$, where $\hat{\mathbf{e}}'_{i,q}$ represents the i -th row of the $n \times q$ matrix of residual coefficients $\hat{\mathbf{E}}_q$ on $\{\Phi_k\}_{k=1}^q$, $\mathbf{x}_{i,p}$ are the first p coefficients of X_i on $\{\Psi_j\}_{j=1}^p$, and $\mathbf{g}_p \in \mathbb{S}^{p-1}$ and $\mathbf{h}_q \in \mathbb{S}^{q-1}$ are the coefficients of $\gamma_X^{(p)}$ and $\gamma_{\mathcal{Y}}^{(q)}$, respectively.

To measure the proximity of (15.6) to zero (and hence to \mathcal{H}_0), and following the ideas of [7] and [12], we consider a Cramér-von Mises norm on $\Pi^{(p,q)} = \mathbb{S}_{\mathbb{H}_2, \{\Phi_k\}_{k=1}^{\infty}}^{q-1} \times \mathbb{S}_{\mathbb{H}_1, \{\Psi_j\}_{j=1}^{\infty}}^{p-1} \times \mathbb{R}$, yielding the so-called Projected Cramér-von Mises (PCvM) statistic:

$$\text{PCvM}_{n,p,q} = \int_{\mathbb{S}^{q-1} \times \mathbb{S}^{p-1} \times \mathbb{R}} [R_{n,p,q}(u, \mathbf{g}_p, \mathbf{h}_q)]^2 F_{n,\mathbf{g}_p}(du) d\mathbf{g}_p d\mathbf{h}_q,$$

where F_{n,\mathbf{g}_p} is the empirical cumulative distribution function of $\{\mathbf{x}'_{i,p} \mathbf{g}_p\}_{i=1}^n$.

From the developments in [11], we get an easily computable form of the statistic:

$$\text{PCvM}_{n,p,q} = \frac{1}{n^2} \frac{2\pi^{p/2+q/2-1}}{q\Gamma(p/2)\Gamma(q/2)} \text{Tr} \left[\hat{\mathbf{E}}_q' \mathbf{A}_{\bullet} \hat{\mathbf{E}}_q \right], \quad (15.7)$$

where $\text{Tr}(\cdot)$ denotes the trace operator and \mathbf{A}_{\bullet} is a certain $n \times n$ symmetric matrix that only depends on $\{\mathbf{x}_{i,p}\}_{i=1}^n$.

15.2.2 Statistic Interpretation and Particular Cases

The statistic (15.7) can be regarded as a weighted quadratic norm:

$$\text{PCvM}_{n,p,q} = \frac{1}{n^2} \frac{2\pi^{p/2+q/2-1}}{q\Gamma(p/2)\Gamma(q/2)} \sum_{k=1}^q \|(\hat{e}_{1,k}, \dots, \hat{e}_{n,k})\|_{\mathbf{A}_{\bullet}},$$

where $\hat{\mathcal{E}}_i^{(q)} = \sum_{k=1}^q \hat{e}_{i,k} \Phi_k$, $i = 1, \dots, n$, and $\|\mathbf{v}\|_{\mathbf{A}_{\bullet}} := (\mathbf{v}' \mathbf{A}_{\bullet} \mathbf{v})^{1/2}$ is a norm in \mathbb{R}^n induced by \mathbf{A}_{\bullet} . Therefore, the statistic aggregates across the dimensions of the

truncated response the \mathbf{A}_\bullet -weighted norms of the coefficients of the functional errors on $\{\Phi_k\}_{k=1}^q$. The basis of such interpretation is the next lemma (proof given in [11]).

Lemma 2 ([11]) *Assume that the functional sample $\{X_i\}_{i=1}^n$ has pairwise distinct coefficients $\{\mathbf{x}_{i,p}\}_{i=1}^n$ on an arbitrary p -truncated basis $\{\Psi_j\}_{j=1}^p$ of \mathbb{H}_1 . Then, for any sample size $n \geq 1$, the $n \times n$ matrix \mathbf{A}_\bullet is positive definite.*

The general framework of the FLMFR seamlessly adapts to scalar response or predictor. So do the estimation methods discussed in Section 15.1.2 and the statistic (15.7). Indeed, in the case of scalar response (see, e.g., [2] and [4]), $\mathbb{H}_2 = \mathbb{R}$ is identifiable with the subspace of $L^2([c, d])$ of constant functions with basis $\{(d - c)^{-1/2}\}$ and $\beta(\cdot, \star) \equiv \beta(\cdot) \in L^2([a, b])$ is a univariate function. The statistic $\text{PCvM}_{n,p,1}$ precisely corresponds to the PCvM statistic for the functional linear model with scalar response given in [12]. In the case of scalar predictor (see [3]), $\beta(\cdot, \star) \equiv \beta(\star) \in L^2([c, d])$ and $\text{PCvM}_{n,1,q}$ results in a test statistic specific for such model.

15.2.3 Bootstrap Calibration and Graphical Tool

The calibration of the statistic (15.7) is done through a wild bootstrap on the residuals. We sketch next the main steps of such resampling, referring to Algorithm 1 in [11] for the specifics and its adaptation to the β -specified case.

1. Compute the statistic $\text{PCvM}_{n,\tilde{p},q}$ from the residuals $\hat{\mathbf{e}}_{i,q} = \mathbf{Y}_{i,q} - \mathbf{X}_{i,\tilde{p}} \hat{\mathbf{B}}_{\tilde{p},q}^{(\lambda),C}$, $i = 1, \dots, n$, associated to the FPCR-L1S estimate $\hat{\mathbf{B}}_{\tilde{p},q}^{(\lambda),C}$ (which selects \tilde{p}).
2. For $b = 1, \dots, B$:
 - a. Perturb the residuals as $\mathbf{e}_{i,q}^{*b} := V_i^{*b} \hat{\mathbf{e}}_{i,q}$, $i = 1, \dots, n$, where $\{V_i^{*b}\}_{i=1}^n$ are independent zero-mean and unit-variance random variables.
 - b. Using $\{\mathbf{e}_{i,q}^{*b}\}_{i=1}^n$, simulate $\{\mathbf{Y}_{i,q}^{*b}\}_{i=1}^n$ from the multivariate linear model.
 - c. Fit the multivariate model from $\{(\mathbf{X}_{i,\tilde{p}}, \mathbf{Y}_{i,q}^{*b})\}_{i=1}^n$ and obtain $\hat{\mathbf{B}}_{\tilde{p},q}^{*b}$.
 - d. Compute the bootstrapped statistic $\text{PCvM}_{n,\tilde{p},q}^{*b}$ from the bootstrap residuals $\hat{\mathbf{e}}_{i,q}^{*b} := \mathbf{Y}_{i,q}^{*b} - \mathbf{X}_{i,\tilde{p}} \hat{\mathbf{B}}_{\tilde{p},q}^{*b}$, $i = 1, \dots, n$.
3. Estimate the p -value by Monte Carlo as $\#\{\text{PCvM}_{n,\tilde{p},q} \leq \text{PCvM}_{n,\tilde{p},q}^{*b}\} / B$.

The bootstrap procedure yields as a by-product a graphical diagnostic tool of the goodness-of-fit of the FLMFR that helps visualizing the possible deviations from \mathcal{H}_0 . The tool compares the empirical process on which the PCvM statistic is applied,

$$R_{n,p,q}(u, \mathbf{g}_p, \mathbf{h}_q) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \hat{\mathbf{e}}_{i,q}' \mathbf{h}_q \mathbb{1}_{\{\mathbf{x}'_{i,p} \mathbf{g}_p \leq u\}},$$

with G samples of its bootstrapped version:

$$R_{n,p,q}^{*b}(u, \mathbf{g}_p, \mathbf{h}_q) = \frac{1}{\sqrt{n}} \sum_{i=1}^n (\hat{\mathbf{e}}_{i,q}^{*b})' \mathbf{h}_q \mathbb{1}_{\{\mathbf{x}'_{i,p} \mathbf{g}_p \leq u\}}, \quad b = 1, \dots, G.$$

The graphical tool employs the FPC bases $\{\hat{\Psi}_j\}_{j=1}^p$ and $\{\hat{\Phi}_k\}_{k=1}^q$ and considers \mathbf{g}_p and \mathbf{h}_q as the canonical vectors in \mathbb{R}^p and \mathbb{R}^q , respectively. This allows to visualize the deviations from \mathcal{H}_0 when “it is projected” in the first FPC of $\{\mathcal{X}_i\}_{i=1}^n$ and the first FPC of $\{\mathcal{Y}_i\}_{i=1}^n$ (or any other combination thereof). [Figure 15.2](#) shows and explains two outputs of this diagnostic tool, for the situations in which \mathcal{H}_0 is and is not rejected.

15.3 Application: AEMET Temperatures Dataset

The `aemet_temp` dataset in the `goffda` [10] package contains daily temperatures of $n = 73$ weather stations from the Meteorological State Agency of Spain (AEMET) during the time span 1974–2013. The dataset is split in two 20-year periods, 1974–1993 and 1994–2013, and the daily temperatures on each weather station are averaged for both periods. This results in two functional samples for the average temperatures across Spain on 1974–1993 (predictor \mathcal{X}) and 1994–2013 (response \mathcal{Y}). Both samples were smoothed with local linear estimators using cross-validated bandwidths to ease visualization. [Figure 15.1](#) (left) shows the samples of \mathcal{X} and \mathcal{Y} .

The PCvM test based on $\tilde{p} = 4$ (selected by FPCR-LIS with λ chosen by one standard error cross-validation) and $q = 3$ (selected such that the proportion of explained variance is 0.99) yielded a p -value equal to 0.4155 using $B = 10^4$ bootstrap replicates. Therefore, the FLMFR is not rejected. The estimated β , shown in [Figure 15.1](#) (right), reveals a temperature increment on the latter period with respect to the former, a conclusion supported by the predominance of positive values on the $\hat{\beta}$ surface and the positiveness of almost all the temperature curves. The diagnostic tool in [Figure 15.2](#) (left) shows no remarkable deviations of the residual marked empirical process from \mathcal{H}_0 . The PCvM test rejects emphatically the simple hypotheses $\mathcal{H}_0 : \beta = 0$ and $\mathcal{H}_0 : \beta(s, t) = \mathbb{1}_{\{s=t\}}$ (stationary-temperature hypothesis; right panel in [Figure 15.2](#)), thus corroborating a significant change in the temperatures between both periods. The diagnostic tool for the latter hypothesis reveals that the non-stationarity is due to the relations between the second FPC of $\{\mathcal{X}_i\}_{i=1}^n$ and $\{\mathcal{Y}_i\}_{i=1}^n$, both related with the variation shape of the temperature curves along the year.

15.4 Software: goffda R Package

The R package `goffda` [10] implements all the methods described and allows for replication of the data application. The implementation of the critical parts of the goodness-of-fit tests, such as the computation of the \mathbf{A}_\bullet matrix and the computation of the PCvM statistic, are implemented in C++ (through Rcpp [6]) for the sake of efficiency. The `goffda` package relies on the `fdata` class from the `fda.usc` [8] package, so it is fully compatible with the latter.

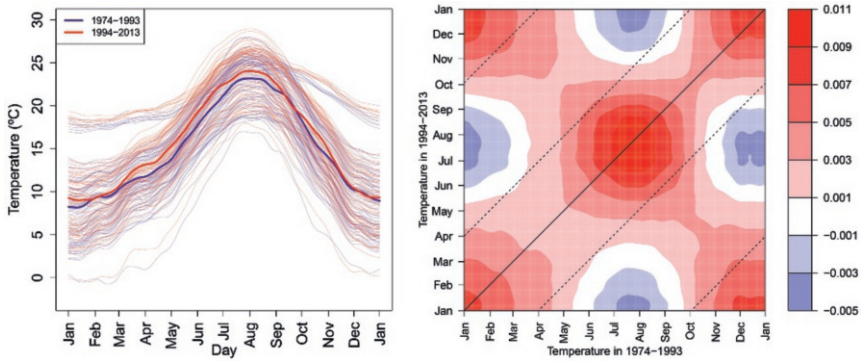


Fig. 15.1 Left: Temperatures of 73 AEMET weather stations for the periods 1974–1983 (X) and 1994–2013 (Y), along with their means. Right: FPCR-L1S estimator $\hat{\beta}$ for the FLMFR.

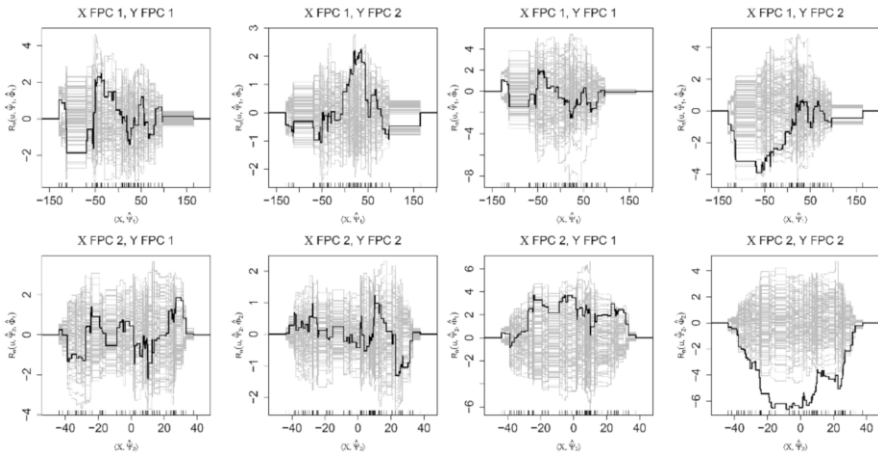


Fig. 15.2 Graphical tool of the PCvM test. The black curve represents the observed process $R_{n,p,q}(u, \mathbf{e}_j, \mathbf{e}_k)$ for its projections on the j -th FPC of $\{X_i\}_{i=1}^n$ and the k -th FPC of $\{Y_i\}_{i=1}^n$, $j, k = 1, 2$. The grey curves stand for the bootstrapped processes under \mathcal{H}_0 , i.e., $R_{n,p,q}^{*b}(u, \mathbf{e}_j, \mathbf{e}_k)$, $b = 1, \dots, 100$. The left 2×2 panel shows the diagnostic output for $\mathcal{H}_0 : m \in \mathcal{L}$ in the AEMET temperatures dataset. The non-rejection of \mathcal{H}_0 is manifested in the centrality of the observed process within the bootstrapped ones. The right 2×2 panel shows the diagnostic for $\mathcal{H}_0 : \beta(s, t) = \mathbb{1}_{\{s=t\}}$, with rejection of \mathcal{H}_0 evidenced by the outlyingness of $R_{n,p,q}(u, \mathbf{e}_2, \mathbf{e}_2)$.

The main functions of `goffda` are: `flm_est` (several estimation methods for the FLMFR); `Adot` (efficient implementation of the \mathbf{A}_\bullet matrix); `flm_stat` (computation of (15.7)); `flm_test` (implementation of the test with its bootstrap resampling). `flm_est` and `flm_test` deal seamlessly with either functional/scalar responses/predictors.

Acknowledgements The authors gratefully thank Prof. Manuel Febrero-Bande for discussions and for providing access to the dataset of raw AEMET temperatures. The first author acknowledges support from grants PGC2018-097284-B-I00 and IJCI-2017-32005 from the Spanish Ministry of Economy and Competitiveness (co-funded with FEDER funds). The second author acknowledges support from grant PGC2018-099549-B-I00 from the same agency. The first and fourth authors acknowledge support from grant MTM2016-76969-P also from the same agency. The authors gratefully acknowledge the computing resources of the Supercomputing Center of Galicia (CESGA).

References

- [1] Benatia, D., Carrasco, M., Florens, J.P.: Functional linear regression with functional response. *J. Econometrics* **201**(2), 269–291 (2017)
- [2] Cardot, H., Ferraty, F., Sarda, P.: Functional linear model. *Statist. Prob. Lett.* **45**(1), 11–22 (1999)
- [3] Chiou, J.M., Müller, H.G., Wang, J.L., Carey, J.R.: A functional multiplicative effects model for longitudinal data, with application to reproductive histories of female medflies. *Statist. Sinica* **13**(4), 1119–1133 (2003)
- [4] Crambes, C., Kneip, A., Sarda, P.: Smoothing splines estimators for functional linear regression. *Ann. Statist.* **37**(1), 35–72 (2009)
- [5] Crambes, C., Mas, A.: Asymptotics of prediction in functional linear regression with functional outputs. *Bernoulli* **19**(5B), 2627–2651 (2013)
- [6] Eddelbuettel, D., François, R.: Rcpp: Seamless R and C++ integration. *J. Stat. Softw.* **40**(8), 1–18 (2011)
- [7] Escanciano, J.C.: A consistent diagnostic test for regression models using projections. *Econometric Theory* **22**(6), 1030–1051 (2006)
- [8] Febrero-Bande, M., Oviedo de la Fuente, M.: Statistical computing in functional data analysis: The R package *fda.usc*. *J. Stat. Softw.* **51**(4), 1–28 (2012)
- [9] Friedman, J., Hastie, T., Tibshirani, R.: Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**(1), 1–22 (2010)
- [10] García-Portugués, E., Álvarez-Liébaná, J.: *goffda*: Goodness-of-fit tests for functional data. R package version 0.0.6 (2019). <https://CRAN.R-project.org/package=goffda>.
- [11] García-Portugués, E., Álvarez-Liébaná, J., Álvarez-Pérez, G., González-Manteiga, W.: A goodness-of-fit test for the functional linear model with functional response. *arXiv:1909.07686* (2019)
- [12] García-Portugués, E., González-Manteiga, W., Febrero-Bande, M.: A goodness-of-fit test for the functional linear model with scalar response. *J. Comp. Graph. Stat.* **23**(3), 761–778 (2014)
- [13] He, G., Müller, H.G., Wang, J.L., Yang, W.: Functional linear regression via canonical analysis. *Bernoulli* **16**(3), 705–729 (2010)
- [14] Imaizumi, M., Kato, K.: PCA-based estimation for functional linear regression with functional responses. *J. Multivariate Anal.* **163**, 15–36 (2018)
- [15] Ramsay, J.O., Silverman, B.W.: *Functional Data Analysis*. Springer Series in Statistics. Springer, New York (2005)
- [16] Yao, F., Müller, H.G., Wang, J.L.: Functional linear regression analysis for longitudinal data. *Ann. Statist.* **33**(6), 2873–2903 (2005)