



scikit-fda: A Python Package for Functional Data Analysis

Carlos Ramos-Carreño 

Universidad Autónoma
de Madrid

José Luis Torrecilla 

Universidad Autónoma
de Madrid

Miguel Carbajo-Berrocal

Universidad Autónoma
de Madrid

Pablo Marcos

Universidad Autónoma
de Madrid

Alberto Suárez 

Universidad Autónoma
de Madrid

Abstract

The library **scikit-fda** is a Python package for functional data analysis (FDA). It provides a comprehensive set of tools for representation, preprocessing, and exploratory analysis of functional data. The library is built upon and integrated in Python's scientific ecosystem. In particular, it conforms to the **scikit-learn** application programming interface so as to take advantage of the functionality for machine learning provided by this package: Pipelines, model selection, and hyperparameter tuning, among others. The **scikit-fda** package has been released as free and open-source software under a 3-clause BSD license and is open to contributions from the FDA community. The library's extensive documentation includes step-by-step tutorials and detailed examples of use.

Keywords: functional data analysis, computational statistics, interactive data visualization, Python, **scikit-learn**.

1. Introduction

Functional data analysis (FDA) is the branch of statistics that deals with observations varying over a continuous parameter, such as curves, surfaces, and other types of functions (Cuevas 2014; Wang, Chiou, and Müller 2016). These types of data appear in many different fields, such as biology (Cremona, Xu, Makova, Reimherr, Chiaromonte, and Madrigal 2019), demographics (Hyndman and Shahid Ullah 2007), economics (Frois Caldeira, Gupta, Suleman, and

Torrent 2020), energy security (Gong, Wang, and Lin 2021), genomics (Leng and Müller 2006; Chen, Cremona, Qi, Mitra, Chiaromonte, and Makova 2020), medicine (Sørensen, Goldsmith, and Sangalli 2013; Ferrando, Ventura-Campos, and Epifanio 2020; Horsley, Ramsay, Ditto, and Da Costa 2021), meteorology (Beyaztas and Yaseen 2019), oceanography (Assunção *et al.* 2020), traffic control (Wagner-Muns, Guardiola, Samaranayake, and Kayani 2018; Hu, Yuan, Zhu, Yang, and Xie 2019), and other areas of application (Ullah and Finch 2013). The functional nature of these data and, in particular, their continuous structure, entails important differences with respect to the classical multivariate statistics. These characteristics require the development of specific statistical and computational tools for their analysis.

Due to the growing interest in functional data, several specialized software tools for FDA have emerged in recent years (Scheipl, Arnone, Hooker, Tucker, and Wrobel 2022). One of the main references in the field is the **fda** package (Ramsay, Hooker, and Graves 2024), which is available in R (R Core Team 2024) and MATLAB (The MathWorks Inc. 2021). This general-purpose library provides an implementation of the methods described in Ramsay and Silverman (2005) and Ramsay, Hooker, and Graves (2009). It utilizes a basis expansion representation of the functional observations. Another important reference in the FDA community is the **fda.usc** package (Febrero-Bande and Oviedo de la Fuente 2012), in which the non-parametric approach developed in Ferraty and Vieu (2006) is adopted. One of the contributions of this package is the introduction of a novel structure for the representation of functional data in discrete form, as a collection of measurements at a grid of points. This R package provides an extensive range of FDA tools, including methods for regression and classification.

A more recent general-purpose R package is **tidyfun** (Scheipl, Goldsmith, and Wrobel 2024). In this package, a novel data structure (vectors of class ‘**tf**’) is introduced to represent functional observations. These ‘**tf**’ vectors can be included as columns in an R ‘**data.frame**’ alongside with other variables. Furthermore, they can be manipulated using the tools of the **tidyverse** ecosystem (Wickham *et al.* 2019). Another recent contribution is **funData** (Happ-Kurz 2020). This package provides a representation for discretized univariate and multivariate functional data of arbitrary dimensions based on S4 classes in R.

Finally, a variety of computational tools have been developed to address specific problems in FDA. Some relevant examples are the packages **refund** (Goldsmith *et al.* 2024), which, among others, provides tools for functional regression and principal component analysis (FPCA), **FDboost** (Brockhaus, Rügamer, and Greven 2020), focused on regression problems, **funFEM** (Bouveyron, Côme, and Jacques 2015) and **funHDDC** (Schmutz, Jacques, Bouveyron, Chèze, and Martin 2020) for functional clustering, **fpca** (Peng and Paul 2011) and **fdapace** (Zhou *et al.* 2022), which are mainly devoted to functional principal component analysis (FPCA). The **fdapace** package available in R, and its MATLAB counterpart **PACE** (Yao, Müller, and Wang 2015), provide methods for both sparsely or densely sampled random trajectories based on FPCA, via the principal analysis by conditional estimation (PACE) algorithm. The package **fdasrvf** (Tucker 2024b) contains tools for alignment, elastic registration, PCA, and regression with functional data based on the square-root velocity framework (SRVF) described in Srivastava and Klassen (2016). This package is available under this name both in R and in MATLAB, as **fdasrsf** (Tucker 2024a) in Python (Van Rossum *et al.* 2011), and as **ElasticFDA** (Tucker 2021) in Julia (Bezanson, Edelman, Karpinski, and Shah 2017). The package **roahd** (Ieva, Paganoni, Romo, and Tarabelloni 2019) includes a collection of methods for robust analysis of functional data. Outlier detection tools are provided also in **fdaoutlier** (Ojo, Lillo, and Fernandez Anta 2023). Visualization tools, including interactive ones, are provided also

in the packages **rainbow** (Hyndman and Shang 2010) and **refund.shiny** (Wrobel, Park, Staicu, and Goldsmith 2016). A recent contribution is the R package **mlrFDA** described in Pfisterer, Beggel, Sun, Scheipl, and Bischl (2021), which gives access and extends the machine learning framework **mlr** (Bischl *et al.* 2016) for the analysis of functional data.

In recent years, the **Python** language has become more relevant in statistics, data science, and machine learning. However, in contrast to the wide variety of alternatives available for FDA in R, the options in **Python** are much more limited both in number and functionality. Some **Python** libraries devoted to FDA are **fdasrsf**, which has been described earlier, and the recently released **FDAPy** (Golovkine 2021), that provides methods for principal component analysis and clustering.

In this context, we present **scikit-fda** (Ramos-Carreño *et al.* 2022), a general-purpose library that makes functional data processing and analysis accessible to the **Python** community. This package offers data structures for the representation of functional data both in discretized form and as a basis expansion, and an extensive set of tools for preprocessing (smoothing, registration and dimensionality reduction), and statistical analysis, including interactive visualization and outlier detection tools. In addition, it provides infrastructure to facilitate the application of the machine learning tools of **scikit-learn** to functional data. Comprehensive documentation is supplied that includes installation instructions, tutorials, API reference, and illustrative examples.

In the short time since its inception, several libraries have been developed using **scikit-fda** as their foundation (Bernard, Lejeune, Ferrieres, and Teste 2021; Consagra, Venkataraman, and Qiu 2022). In addition, it has been employed in a number of recent investigations (Fermanian 2022; Torrecilla, Ramos-Carreño, Sánchez-Montañés, and Suárez 2020; Tan, Bergmeir, Petitjean, and Webb 2020, 2021; Pegoraro and Beraha 2021).

In what follows, an overview of the functionality provided by the **scikit-fda** package is given: In Section 2, the discretized and the basis expansion representations of the functional observations are introduced. Interpolation, extrapolation, and derivation tools are presented also in this section. In Section 3, the functionality of the package is described. It includes subsections devoted to preprocessing (smoothing, registration, dimensionality reduction, and variable selection), exploratory analysis (descriptive statistics, depth measures, interactive visualization, and outlier detection), and the integration with the machine learning tools of **scikit-learn**. Examples of code that illustrate **scikit-fda**'s functionality are provided throughout the article. The tools and processes implemented to ensure the quality of the code are described Section 4. This section includes also an overview of the project's extensive documentation.

2. Representation of functional data in **scikit-fda**

The **scikit-fda** package provides tools for the representation of functional observations of the form $x : \mathcal{T} \subseteq \mathbb{R}^p \rightarrow \mathbb{R}^q$, with $p \geq 1$, and $q \geq 1$. The parameter p is the dimension of the domain of the functions ($p = 1$ for curves, $p = 2$ for surfaces, etc.). The parameter q is the dimension of the co-domain; that is, the number of output coordinates for vector-valued functions. For instance, a gray-scale two-dimensional image can be treated as a functional datum with $p = 2$, for the location of the pixels in the image, and $q = 1$, for the intensity at each pixel. A color image consisting of three channels (e.g., red, green, and blue) would have $p = 2$ and $q = 3$. The values of p and q are the same for all the observations in a functional

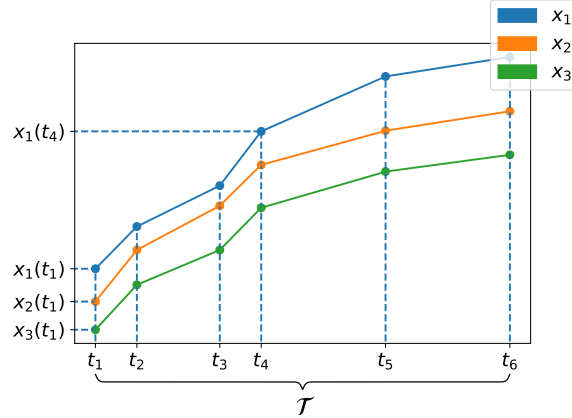


Figure 1: Functional observations in discretized form. The quantity $x_n(t_j)$ represents the value of the n -th trajectory at t_j .

dataset. For the sake of clarity, we focus on the case of real-valued univariate functions (that is, $p = q = 1$) for which most of the functionality described in this article is implemented. We will further assume that the functions are defined on a compact interval in the real line. In higher dimensions, the domain is assumed to be rectangular. Typically, the continuous parameter on which the functions depend, $t \in \mathcal{T}$, is assumed to be time.

A functional dataset consists of a collection of N observations $\{x_i(t), t \in \mathcal{T}\}_{i=1}^N$, where $x_i(t)$ is the i -th observation in the sample. Each observation can be represented either in discretized form, or as a basis expansion. In the former representation, a functional observation consists of a set of measurements at a grid of points $\mathbf{t} = (t_1, \dots, t_M) \in \mathcal{T}^M$, which is common for all observations. This grid need not be regularly spaced. The i -th observation in the sample is represented by the vector $\{x_i(\mathbf{t})\}_{i=1}^N$, where $x_i(\mathbf{t}) = (x_i(t_1), \dots, x_i(t_M))$. The discretization grid is assumed to be sufficiently fine so that the functional character of the data is apparent (Ramsay and Silverman 2005; Ferraty and Vieu 2006; Hsing and Eubank 2015). As an illustration of this representation, three example trajectories measured at a grid of irregularly-spaced points are displayed in Figure 1.

Alternatively, a functional observation can be represented as an expansion in a functional basis $\{\phi_i(t), t \in \mathcal{T}\}_{i \geq 1}$

$$x(t) = \sum_{i \geq 1} c_i \phi_i(t),$$

where $\{c_i\}_{i \geq 1}$ are the coefficients of the expansion.

The package **scikit-fda** provides data structures for both types of representation: The class ‘**FDataGrid**’ for discretized data, and the class ‘**FDataBasis**’ for expansions in a functional basis. They are derived from the abstract class ‘**FData**’, which provides common properties and methods. In what follows, these classes are described in detail.

2.1. The class ‘FData’

In the class ‘**FData**’, the attributes and methods shared by the discretized and the basis expansion representations of the functional dataset are collected. Thus, it provides a common

interface for both ‘FDataGrid’ and ‘FDataBasis’ objects. Specifically, objects of the ‘FData’ class have the following attributes:

- **dataset_name**: Name of the functional dataset.
- **n_samples**: Size (number of functional observations) of the dataset.
- **dim_domain**: Dimension of the domain in which the functions are defined ($p \geq 1$).
- **argument_names**: Names of each of the p arguments of the function (domain dimensions).
- **domain_range**: Limits of the intervals for each of the domain arguments. They are used as the default ranges for plotting and numerical integration.
- **dim_codomain**: Dimension of the codomain q (output). For scalar functions, $q = 1$. Values $q > 1$ correspond to vector-valued functions.
- **coordinate_names**: Names of the q codomain coordinates.
- **extrapolation**: Default extrapolation strategy; for instance, constant or periodic. See Section 2.4 for details.

As an illustration of this data structure, consider the case of observations that are bidimensional RGB images. For this functional dataset, **dim_domain** would be 2, corresponding to the two dimensions of the image. The names of these dimensions (e.g., “x” and “y”, or “horizontal” and “vertical”) would be stored in the attribute **argument_names**. The attribute **dim_codomain** would be 3, corresponding to the three color channels. The names of these channels (“R”, “G”, and “B”) would be stored in the attribute **coordinate_names**.

The class ‘FData’ provides also methods that are common to both representations; for instance, methods for evaluation, addition, multiplication by a scalar, and plotting. Since ‘FData’ is abstract, it is not possible to directly instantiate an object of this class. Instead, objects of one of its subclasses, ‘FDataGrid’ or ‘FDataBasis’, need to be created.

2.2. Discretized representation: The class ‘FDataGrid’

Functional data are often the result of monitoring a continuous process at a discrete set of points. For the general case, $x : \mathbb{R}^p \rightarrow \mathbb{R}^q$, with $p, q \geq 1$, we assume that the discretization grid in the j -th dimension is $\mathbf{t}_j = (t_{j1}, \dots, t_{jM_j})$, with $j = 1, \dots, p$. In the **scikit-fda** library, the points in the grid need not be regularly spaced. The grid has to be the same for all observations in the functional dataset. The dataset $\{x_i : \mathbb{R}^p \rightarrow \mathbb{R}^q\}_{i=1}^N$ is represented as an object of the class ‘FDataGrid’. The values of the functional observations are stored in the tensor $\{x_i(\mathbf{t}) = x_i(\mathbf{t}_1 \times \dots \times \mathbf{t}_p)\}_{i=1}^N$ of dimension $N \times M_1 \times \dots \times M_p \times q$. Here $\mathbf{t}_1 \times \dots \times \mathbf{t}_p$ is the grid of points obtained as the Cartesian product of $\mathbf{t}_1, \dots, \mathbf{t}_p$. In the case $p = q = 1$, the discretized sample is simply an $N \times M$ matrix $\{x_i(\mathbf{t}) = (x_i(t_1), \dots, x_i(t_M))\}_{i=1}^N$.

In addition to those inherited from ‘FData’, objects of this class have the following attributes:

- **grid_points**: Sequence of discretization grids, one for each of the domain dimensions ($\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_p$). The values of the functions are specified at the Cartesian product of the 1-D arrays in the sequence.

- **data_matrix**: **NumPy** array of dimensions $N \times M_1 \times \cdots \times M_p \times q$ in which the values of the N functional observations are stored.
- **interpolation**: Default interpolation strategy for locations within the rectangular discretization grid. See Section 2.4 for details.

Since the attribute **data_matrix** is a **NumPy** array, it is possible to carry out pointwise operations, such as powers, exponentials, logarithms, and trigonometric functions by directly applying the corresponding **NumPy** functions (Harris *et al.* 2020).

As an illustration, in the following code a ‘**FDataGrid**’ object is created with three functional observations measured at grid points $\mathbf{t} = (0, 0.1, 0.3, 0.4, 0.7, 1)$. These data are depicted in Figure 1.

```
>>> import skfda
>>> grid_points = [0.0, 0.1, 0.3, 0.4, 0.7, 1.0]
>>> data_matrix = [
...     [109.5, 115.8, 121.9, 130.0, 138.2, 141.1],
...     [104.6, 112.3, 118.9, 125.0, 130.1, 133.0],
...     [100.4, 107.1, 112.3, 118.6, 124.0, 126.5],
... ]
>>> fd = skfda.FDataGrid(data_matrix=data_matrix, grid_points=grid_points)
```

In the previous example, the discretization points and the values of the functions are specified manually. The values of **grid_points** and **data_matrix** can be imported from data files in standard formats, such as CSV, XLSX, ARFF, MATLAB files, with the help of the corresponding functions from **NumPy** (Harris *et al.* 2020), **SciPy** (Virtanen *et al.* 2020), **pandas** (pandas Development Team 2020), and similar packages. An example of how to import data from a CSV file is provided in https://fda.readthedocs.io/importing_data.

2.3. Basis expansion representation: The class ‘**FDataBasis**’

Assume that the functional observations in the dataset considered belong to \mathcal{F} , a separable Hilbert space; for instance, L^2 . Under such assumption there exists a countable basis $\{\phi_i(t)\}_{i \geq 1}$ that is complete, so that any $x(t) \in \mathcal{F}$ can be expressed as

$$x(t) = \sum_{i \geq 1} c_i \phi_i(t),$$

where $\{c_i\}_{i \geq 1}$, are the coefficients of basis expansion. The **scikit-fda** library provides support for such a representation in the constant, monomial, B-spline, and Fourier bases. In addition, other types of bases can be implemented by the user. An example showing how to define new bases is available in https://fda.readthedocs.io/create_new_bases. The choice of basis should be made taking into consideration the characteristics of the data at hand. For instance, the Fourier basis is well-suited to representing periodic functions. For non-periodic data, a representation in the B-splines basis is probably more appropriate. The monomial basis is useful to represent polynomials. Monomials are also the building blocks of the Maclaurin series. Therefore, the monomial basis can be used to represent local approximations of analytic functions. The first five elements of the different bases provided in **scikit-fda** are shown in Figure 2.

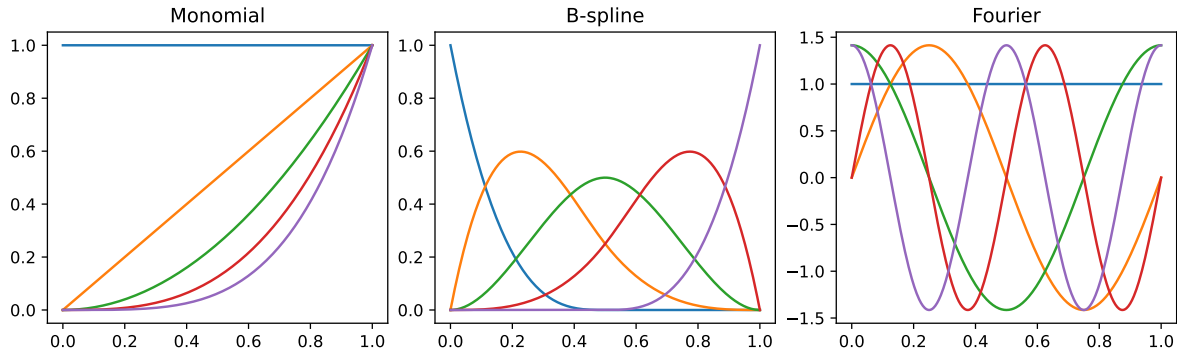


Figure 2: First five elements of the bases available in **scikit-fda**: Monomial (left), B-splines (center), and Fourier (right).

In general, these types of representation are infinite-dimensional. In practice, the expansion is truncated to the first K terms

$$x(t) \approx \sum_{i=1}^K c_i \phi_i(t). \quad (1)$$

Truncation often results in the smoothing of the original functional observations. This smoothing effect can be beneficial for the representation of noisy data (see Section 3.3.1).

In **scikit-fda**, the class ‘`FDataBasis`’ is used to represent functional data as a finite basis expansion. In addition to those inherited from ‘`FData`’, objects of this class have the following attributes:

- **basis**: The basis for the representation of the functional observations. It is an object of the class ‘`Basis`’, one of whose attributes, `n_basis` is the number of elements of the basis considered. The bases available are ‘`ConstantBasis`’, ‘`MonomialBasis`’, ‘`BSplineBasis`’ and ‘`FourierBasis`’.
- **coefficients**: The $N \times K$ matrix that contains the coefficients of the basis expansion.

The following code is used to illustrate **scikit-fda**’s support for this type of representation.

```
>>> import skfda
>>> from skfda.representation.basis import FourierBasis, BSplineBasis
>>> X, y = skfda.datasets.fetch_phoneme(return_X_y=True)
>>> X.to_basis(BSplineBasis(n_basis=5))
>>> X.to_basis(FourierBasis(n_basis=5))
```

In this code, the *Phoneme* dataset (Hastie, Tibshirani, and Friedman 2009) is used. The curves in this dataset correspond to log-periodograms of the time series of utterances of five different phonemes by different speakers. The functional observations are transformed from their original discretized representation into different basis expansions. A sample of the transformed trajectories, together with the original ones, is displayed in Figure 3. In these plots, the smoothing effect of the transformation to the basis representations is apparent.

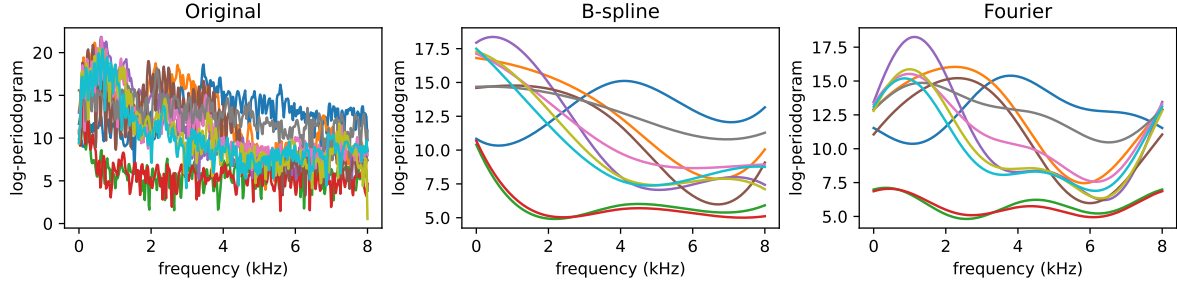


Figure 3: Different representations of the first ten trajectories of the *Phoneme* dataset. From left to right: original trajectories, B-spline, and Fourier basis representation. In both cases, 5 basis functions are considered

2.4. Interpolation and extrapolation

The **scikit-fda** package provides a variety of interpolation methods for functional data in discretized form. By default, linear interpolation is performed. Other types of interpolation can be specified in the attribute `interpolation` of the `FDataGrid` object. Specifically, support for spline interpolation is provided by the class `SplineInterpolation`. It is also possible to employ other interpolation and extrapolation strategies defined by the user. An example showing how to define such custom strategies is available at https://fda.readthedocs.io/create_new_interpolation.

Different extrapolation strategies for `FDataGrid` and `FDataBasis` objects are available in **scikit-fda**. In particular, it is possible to specify a constant value (for instance, the value of the function at one of the limits of the domain), or to assume a periodic structure. In the case of `FDataBasis` objects, one can also directly evaluate the basis expansion outside the domain. Finally, as in interpolation, user-defined extrapolation strategies can be utilized.

2.5. Derivatives

The computation of derivatives is of particular importance in functional data analysis. For instance, derivatives can reveal significant information that is not apparent in the original curves. Furthermore, the norm of a derivative is a natural measure of the function's roughness (Ramsay and Silverman 2005). For this reason, they are often employed to define penalties for regularization. In the **scikit-fda** package, the method `derivative()` can be used to perform this operation for both `FDataGrid` and `FDataBasis` objects. In the case of `FDataGrid` objects, derivatives are approximated using finite differences. For `FDataBasis` objects, they are computed exactly in terms of the derivatives of the basis functions. Therefore, if a new type of basis is designed, it is necessary to implement the derivatives of the basis functions in the corresponding class.

2.6. Regularization

Regularization methods consist in favoring simpler models to improve the quality and robustness of the solutions of an optimization problem. In FDA, regularization is used to obtain smooth functional approximations to noisy discrete data, for registration, and for principal component analysis, among others. For the purpose of regularization, the complexity

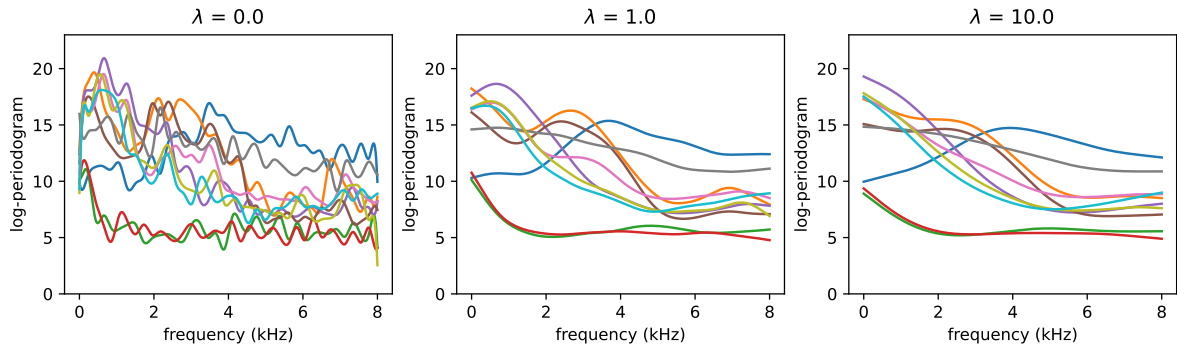


Figure 4: Smoothed representation of the first ten trajectories of the *Phoneme* dataset in a B-spline basis with 40 basis functions for different values for the regularization parameter $\lambda \geq 0$; from left to right: $\lambda = 0$ (no regularization), $\lambda = 1$, and $\lambda = 10$.

of a function can be quantified in terms of its norm, or of a linear transformation thereof (e.g., the function derivatives). A penalty term proportional to this measure of complexity is then added to the cost function to be minimized. The package **scikit-fda** provides the necessary infrastructure to implement regularization based on the L^2 -norm of the function in class ‘**L2Regularization**’. Alternatively, a linear operator can be passed as a parameter to the constructor of ‘**L2Regularization**’ objects. Some common linear operators are readily available in **scikit-fda**’s operators module for that purpose. The following code illustrates this type of regularization to obtain smooth representations of a set of functions in the basis of B-splines. In this example, the L^2 norm of second derivatives of the trajectories is used to penalize their curvature.

```
>>> import skfda
>>> X, y = skfda.datasets.fetch_phoneme(return_X_y=True)
>>> X = X.coordinates[0]
>>> basis = skfda.representation.basis.BSplineBasis(
...     domain_range=X.domain_range, n_basis=40)
>>> regularization = skfda.misc.regularization.L2Regularization(
...     skfda.misc.operators.LinearDifferentialOperator(2),
...     regularization_parameter=1)
>>> smoother = skfda.preprocessing.smoothing.BasisSmoother(
...     basis=basis, regularization=regularization, return_basis=True)
>>> X_basis = smoother.fit_transform(X)
```

The effect of this type of smoothing is illustrated in Figure 4. In this figure, 10 trajectories of the *Phoneme* dataset are represented in a B-spline basis composed of 40 basis functions for different values for the regularization parameter λ ; from left to right: $\lambda = 0$ (no regularization), $\lambda = 1$, and $\lambda = 10$. As shown in the left plot of this figure, when the number of basis functions is large, the non-regularized trajectories represented in this basis still exhibit significant fluctuations. Progressively smoother basis representations are obtained as the *regularization parameter* $\lambda > 0$ increases.

3. Functionality of scikit-fda

In this section, an overview of the utilities for functional data analysis provided by the **scikit-fda** package is given. The first step in the analysis is to generate functional datasets or to retrieve them from external sources. In the **scikit-fda** package it is possible to generate synthetic data, to simulate random trajectories from stochastic processes, and to load data from files in standard formats and from repositories of real-world datasets. The library provides also an extensive set of tools for the analysis of functional data both in discretized and basis expansion forms. In particular, it offers methods for exploratory analysis, smoothing, registration, dimensionality reduction, the computation of functional depths, outlier detection, and interactive visualization, among others. An important feature of **scikit-fda** is the integration with the extensive collection of **scikit-learn**'s tools for machine learning, including data preprocessing, training, testing, and hyperparameter selection. Specifically, the methods provided are designed so that they can be utilized in **scikit-learn** *pipelines*. In what follows these functionalities will be described in detail.

3.1. Generation of synthetic data

A variety of methods for the generation of functional data, either from some simple models or sampled from stochastic processes, are available in **scikit-fda**. In particular, the function `make_multimodal_samples()` can be used to generate functions with several maxima. This is used in the synthetic registration example illustrated in Figure 8. The function `make_gaussian_process()` can be used to simulate trajectories sampled from a Gaussian process with a specified mean and covariance function. Several commonly employed covariance functions, such as Brownian, exponential, radial basis function (RBF), Matérn, and polynomial kernels are supplied in the package. Additional types of covariance functions can be defined by the user.

The following code illustrates the generation of 50 trajectories from standard Brownian Motion, an Ornstein-Uhlenbeck process (a Gaussian process with an exponential covariance function), and a Gaussian process with an RBF covariance function. A regular grid of 100 equally spaced points in $\mathcal{T} = [0, 1]$ is employed for the discretized representation. The simulated trajectories are displayed in Figure 5.

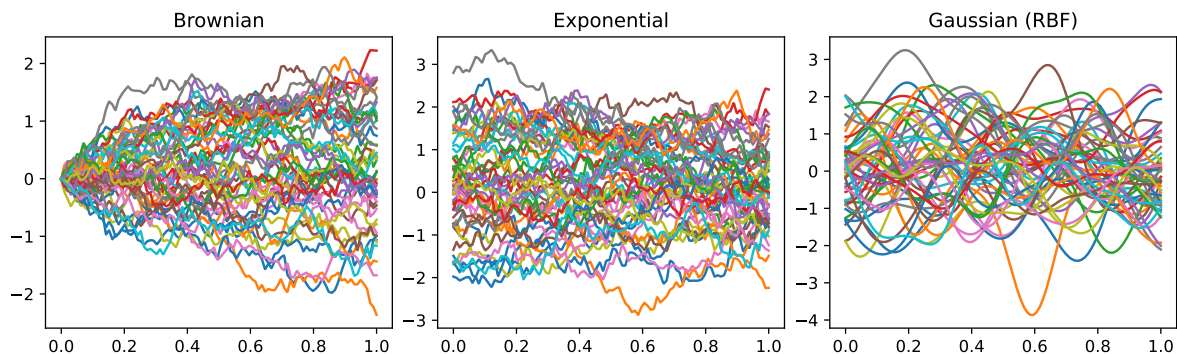


Figure 5: Trajectories sampled from Gaussian processes with different covariance functions. From left to right, standard Brownian, Exponential ($l = 1$) and RBF ($l = 0.1$).

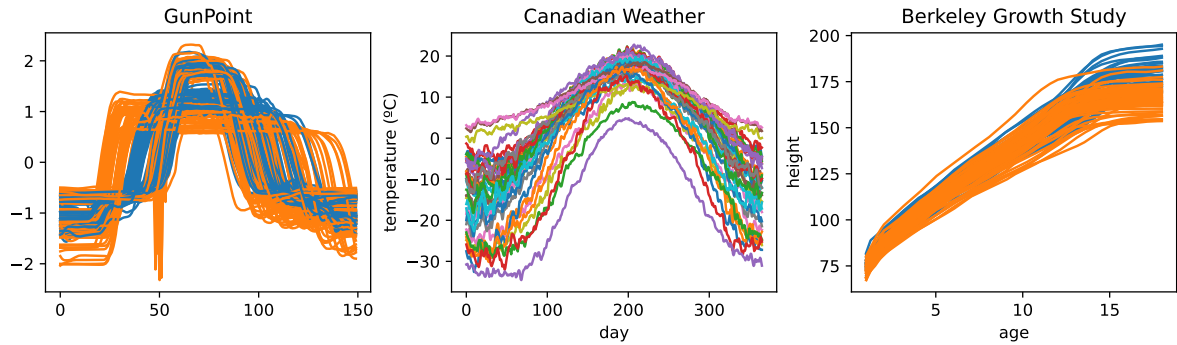


Figure 6: Real datasets fetched with **scikit-fda**. From left to right: *GunPoint*, *Canadian Weather*, and the *Berkeley Growth Study*.

```
>>> import skfda
>>> from skfda.misc.covariances import Brownian, Gaussian, Exponential
>>> cov_dict = {
...     "Brownian": Brownian(variance=1),
...     "Exponential": Exponential(variance=1, length_scale=1),
...     "Gaussian (RBF)": Gaussian(length_scale=0.1)}
>>> for i, (name, cov) in enumerate(cov_dict.items()):
...     fd = skfda.datasets.make_gaussian_process(
...         n_samples=50, n_features=100, mean=0, cov=cov, random_state=0)
...     fd.plot()
```

3.2. Real-world data

The package **scikit-fda** provides tools to retrieve functional datasets from other libraries and public access repositories. The datasets themselves are retrieved using the package **scikit-datasets** (Díaz-Vico and Ramos-Carreño 2022). The data are downloaded only once and cached on disk, so as to reduce network traffic and make it possible to work with the data offline. They are then converted to the ‘FData’ format for further processing and analysis. For instance, the function `fetch_cran()` can be used to retrieve datasets from R packages in the CRAN repository. Datasets from the *UCR & UEA Time Series Classification Archive* (Dau et al. 2019; Bagnall et al. 2018; Bagnall, Lines, Vickers, and Keogh 2021) can be accessed making use of the function `fetch_ucr()`. Moreover, there exist specific functions to import some widely-used datasets such as `fetch_growth()` for the *Berkeley Growth Study* dataset (Tuddenham and Snyder 1954) and `fetch_weather()` for the *Canadian Weather* dataset (Ramsay and Silverman 2005).

The following code illustrates the functionality described for three well-known datasets: *GunPoint* from the UCR repository, *Canadian Weather*, and the *Berkeley Growth Study*. **scikit-fda**’s functions are used to load the data and plot some of the datasets’ trajectories. Note that the *Canadian Weather* dataset has two codomain dimensions ($q = 2$): temperature and precipitation. In this example, the first is selected using the `coordinates` property of an ‘FData’ object.

```
>>> import skfda
```

```

>>> dataset = skfda.datasets.fetch_ucr("GunPoint")
>>> dataset["data"].plot(group=dataset["target"])
>>> X, _ = skfda.datasets.fetch_weather(return_X_y=True)
>>> X.coordinates[0].plot()
>>> X, y = skfda.datasets.fetch_growth(return_X_y=True)
>>> X.plot(group=y)

```

3.3. Preprocessing

Functional observations often need to be subject to some form of processing to facilitate ulterior manipulation. To this end, the **scikit-fda** package provides utilities for smoothing, registration, and dimensionality reduction. In what follows, these utilities are described in detail.

Smoothing

Smoothing consists in replacing the original functional observation $x(t)$ with a smoothed version $\hat{x}(t)$. This replacement yields a more manageable, possibly more faithful, representation of the underlying process. In particular, smoothing can be used to recover the signal component from noisy measurements. Furthermore, smoothed approximations with continuous derivatives can be used for regularization (Wang *et al.* 2016). The methods of **scikit-fda**'s classes `'BasisSmoother'` and `'KernelSmoother'` can be employed to this end. The package provides also utilities to determine an appropriate degree of smoothing using some form of statistical validation.

As discussed in Section 2, the approximation of a function by a truncated basis expansion, as in (1), is a form of smoothing. The coefficients of the finite basis expansion can be estimated by least squares (Ramsay and Silverman 2005). This kind of smoothing is performed in **scikit-fda** by the class `'BasisSmoother'`. Further smoothing can be achieved by a regularization approach based on roughness penalties, as described in Section 2.6 (Green and Silverman 1993; Ramsay and Silverman 2005).

Smoothing can be achieved also by performing a linear transformation of the original functional observations

$$\hat{x}(t) = \int_{\mathcal{T}} s_t(\tau) x(\tau) d\tau. \quad (2)$$

The weighting function $s_t(\tau)$ quantifies the contribution of the value of the function at τ to the smoothed value at t . This weighting function should be localized, so that the values of the function at points close to t contribute more to the average.

For functional data in discretized form, (2) can be expressed as a matrix transformation

$$\hat{\mathbf{x}} = \mathbf{S}\mathbf{x},$$

where $\mathbf{x} = x(\mathbf{t})$ is the vector of values of the function at the discretization points $\mathbf{t} = (t_1, \dots, t_M)$, the vector $\hat{\mathbf{x}}$ consists of the smoothed function values at a grid of points, which can be different from the original ones, and \mathbf{S} is the smoothing matrix. By default, the grid at which the smoothed values are computed is \mathbf{t} , the set of sampling points. The smoothing matrix \mathbf{S} is sometimes referred to as the “hat” matrix, a name borrowed from regression analysis because it transforms the dependent variable vector \mathbf{x} into its fitted version $\hat{\mathbf{x}}$ (Ramsay and Silverman 2005).

Smoothing with local weights can be implemented using kernels (Wasserman 2006). **scikit-fda** provides three smoothers of this type: Nadaraya-Watson, local linear regression, and k -nearest neighbors. As an illustration, for the the Nadaraya-Watson smoother, the hat matrix is

$$S_{ij}(h) = \frac{K\left(\frac{t_i - t_j}{h}\right)}{\sum_{m=1}^M K\left(\frac{t_i - t_m}{h}\right)}, \quad 1 \leq i, j \leq M,$$

where K is the kernel function and h is the parameter that controls the degree of smoothing. Commonly used kernel functions, such as Gaussian, uniform, and Epanechnikov, are available in **scikit-fda**. It is also possible to employ user-defined kernels for this type of smoothing.

In these methods, the value of the smoothing parameter needs to be carefully adjusted to avoid under- or oversmoothing. In **scikit-fda** this value can be determined by statistical validation with the help of the class ‘`SmoothingParameterSearch`’. Particular scoring criteria, such as the leave-one-out cross-validation (‘`LinearSmootherLeaveOneOutScorer`’) or, alternatively, generalized cross-validation (‘`LinearSmootherGeneralizedCVScorer`’) are provided to guide the search. The criterion that is maximized in leave-one-out cross validation is

$$CV_{loo}(h) = \frac{1}{M} \sum_{m=1}^M \left(\frac{x(t_m) - \hat{x}(t_m; h)}{1 - S_{mm}(h)} \right)^2,$$

where h is the smoothing parameter and $\hat{x}(t_m; h)$ is the smoothed value.

The generalized cross validation (GCV) criterion is

$$GCV(h) = \Xi(\mathbf{S}(h)) \frac{1}{M} \sum_{m=1}^M (x(t_m) - \hat{x}(t_m; h))^2,$$

where Ξ is a penalty function. By default, the penalty is

$$\Xi(\mathbf{S}(h)) = \frac{1}{(1 - \text{trace}(\mathbf{S}(h))/M)^2}.$$

Additional penalty functions, such as Akaike’s information criterion (AIC), implemented as `akaike()`, or Shibata’s model selector, implemented as `shibata()`, are provided as well.

In the code that follows, the smoothing functionality provided by **scikit-fda** is illustrated using the functions of the *Phoneme* dataset (Hastie *et al.* 2009), which are rather noisy.

```
>>> import skfda
>>> from skfda.preprocessing.smoothing import KernelSmoother, validation
>>> from skfda.misc.hat_matrix import KNeighborsHatMatrix
>>> X, y = skfda.datasets.fetch_phoneme(return_X_y=True)
>>> grid = validation.SmoothingParameterSearch(
...     KernelSmoother(KNeighborsHatMatrix()), [2, 3, 4, 5],
...     param_name="kernel_estimator__n_neighbors",
...     scoring=validation.LinearSmootherGeneralizedCVScorer(
...         validation.shibata))
>>> grid.fit(X)
>>> X_smooth = grid.transform(X)
```

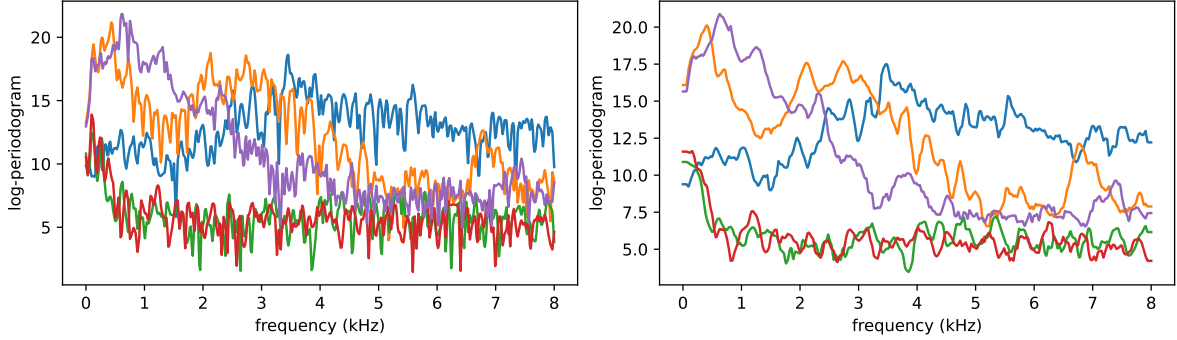


Figure 7: Five trajectories of the *Phoneme* dataset before (left) and after (right) nearest neighbors smoothing.

In this example, k-nearest neighbors is used to smooth the data. The optimal number of neighbors is selected between the values 2, 3, 4 and 5. Model selection is carried out using ‘`LinearSmootherGeneralizedCVScorer`’ with the `shibata()` penalty function. The first five original curves and the smoothed ones are displayed in Figure 7.

Registration

Another type of preprocessing, which is especially relevant in FDA, is registration. Registration consists in applying transformations to the raw data so that the functional observations are properly aligned. There is a variety of reasons why misalignment can occur. In some cases, it is the result of errors in the measurement process. In others, the domain has to be warped because the functions depend on an internal parameter, which is different from the one observed. For periodic functions, such as the signal of a heartbeat, the starting time for the different measurements could be different. A number of strategies can be used for registration. For instance, maxima, minima, zeros, and other landmarks can be used as reference points for alignment. Alternatively, some measure of dispersion between the observations can be minimized. It is also possible to register a set of functional observations to a reference function. After registration, it may be necessary to evaluate the functional observations at points in the domain that are different from the ones in the original grid. This can be made utilizing the interpolation and extrapolation techniques described in Section 2.4. To carry out such an alignment, the package **scikit-fda** offers support for shift registration, and for elastic registration.

Shift registration consists in aligning the functional observations by a translation

$$\tilde{x}_i(t) = x_i(t + \delta_i), \quad i = 1, \dots, N,$$

where δ_i is the time shift applied to $x_i(t)$, and $\tilde{x}_i(t)$ is the registered function (Ramsay and Silverman 2005). Shifting modifies the lower and upper bounds of the interval on which the function observations are defined. The values of the shifted functions that lie outside the original interval are discarded. For the subinterval in which the function values are not available, they are estimated by extrapolation. The method for extrapolation can be provided as an input.

The shifting constants $\{\delta_i\}_{i=1}^N$ can be determined using different procedures. If a single landmark, such as a maximum, a minimum, or a zero crossing, is present in every curve, and

their locations, $\{\tau_i\}_{i=1}^N$, are known, then the i -th curve can be shifted by $\delta_i = \tau_i - \tau^*$. After registration, the location of the landmark is τ^* for every curve,

$$\tilde{x}_i(\tau^*) = x_i(\tau_i), \quad i = 1, \dots, N.$$

In **scikit-fda**, the function `landmark_shift_registration()` can be used to carry out this transformation. The values of the δ_i can be retrieved using the `landmark_shift_deltas()` function.

Alternatively, the values δ_i can be computed by minimizing a least squares criterion ([Ramsay and Silverman 2005](#))

$$REGSSE = \sum_{i=1}^N \int_{\mathcal{T}} [\tilde{x}_i(t) - \hat{\mu}(t)]^2 dt,$$

where $\hat{\mu}(t)$ is the sample mean of the registered data $\{\tilde{x}_i(t)\}_{i=1}^N$. This type of shift registration can be performed with methods of class `LeastSquaresShiftRegistration`. Instead of the sample mean, which is the default value, a user-defined template function can be employed. In this case, the values for the δ_i are stored as the attribute `deltas_` after the registration.

Another type of registration available in the package **scikit-fda** is elastic registration. In elastic registration, one attempts to align the data by applying a warping transformation

$$\tilde{x}_i(t) = x_i(\gamma_i(t)), \quad i = 1, \dots, N.$$

The warping γ_i is a monotonically increasing function defined in $\mathcal{T} = [a, b]$. Assuming that the values of the function at the endpoints of this interval are fixed, it obeys the constraints $\gamma_i(a) = a$ and $\gamma_i(b) = b$. If the locations of some landmarks are known, the warping function for elastic registration can be approximated by monotonically increasing splines ([Ramsay and Silverman 2005](#)). Besides the boundary constraints specified earlier, the spline interpolator for the i -th functional observation has to satisfy

$$\tilde{x}_i(\tau_l^*) = x_i(\tau_{il}), \quad l = 1, \dots, L,$$

where $\{\tau_{il}\}_{l=1}^L$ are the landmark locations in the i -th observation, and $\tau_l^* = \gamma_i^{-1}(\tau_{il})$ is the location of the l -th landmark in the registered curves. In **scikit-fda** this type of registration can be carried out using the function `landmark_elastic_registration()`. The warpings can be retrieved with the function `landmark_elastic_registration_warping()`. A drawback of this approach is that the landmarks and their locations need to be identified beforehand.

An alternative type of elastic registration is to align the observations to a reference template. This has the advantage that no information on landmarks is needed. In the elastic registration method described in [Srivastava, Wu, Kurtek, Klassen, and Marron \(2011\)](#), the template is defined in terms of the Karcher mean under the Fisher-Rao metric ([Srivastava and Klassen 2016](#)). Then, an energy function depending on the Fisher-Rao distance between each curve and the template is minimized. To efficiently compute this distance, the square root velocity function (SRVF) transform is introduced ([Joshi, Klassen, Srivastava, and Jermyn 2007](#)). The main reason for introducing this transform is that the Fisher-Rao distance between two functions is given by the L^2 distance between their SRVF representations. The **scikit-fda**'s class `FisherRaoElasticRegistration` includes methods for this particular type of registration. The warping functions used are stored in the `warping_` attribute of this class. The implementation makes use of the dynamic programming routines for alignment to a

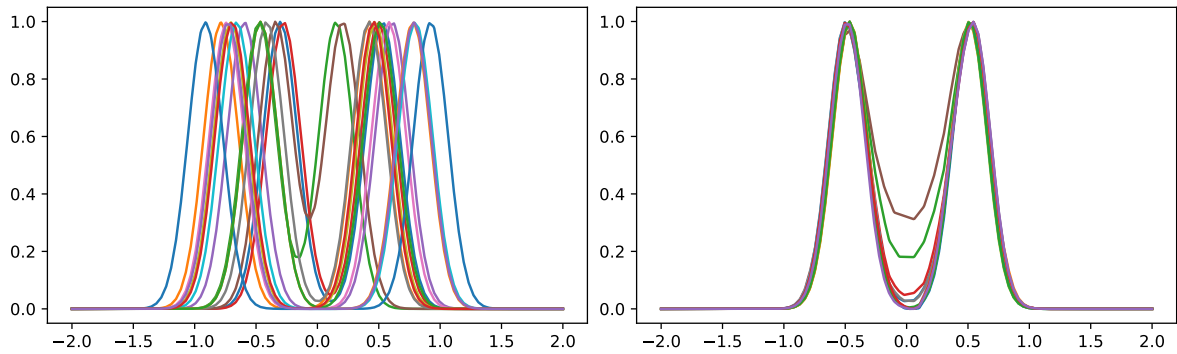


Figure 8: Elastic Fisher-Rao registration with synthetic data: The original curves are shown in the left panel. The registered curves are displayed in the right panel.

template from the Python package **fdasrsf** (Tucker 2024a). This type of elastic registration is illustrated in Figure 8 with synthetic data. The original trajectories, which are displayed on the left plot, exhibit two local maxima whose relative locations are different in each of the curves. In consequence, alignment cannot be achieved via a simple shift. Note that after this type of elastic registration curves are well aligned even without previous information about the location of the landmarks.

The results of applying shift registration by least squares and elastic Fisher-Rao registration to the *Berkeley Growth Study* data (Tuddenham and Snyder 1954) are compared in Figure 9. This figure has been generated using the following code:

```
>>> import skfda
>>> from skfda.preprocessing.registration import (
...     FisherRaoElasticRegistration,
...     LeastSquaresShiftRegistration)
>>> X, y = skfda.datasets.fetch_growth(return_X_y=True)
>>> X_aligned_elastic = FisherRaoElasticRegistration().fit_transform(X)
>>> X_aligned_shift = LeastSquaresShiftRegistration().fit_transform(X)
>>> X.plot()
>>> X_aligned_shift.plot()
>>> X_aligned_elastic.plot()
```

The curves displayed in the left panel of Figure 9 trace the evolution of the heights of 54 girls and 39 boys since their birth until their 18th birthday. The nominal ages at which the measurements are made coincide for all individuals. However, each child has a different growth profile. In particular, landmark features manifest themselves at different ages. For instance, even though most curves exhibit a growth spurt at puberty, the precise ages at which this occurs are different for each individual. Therefore, an elastic deformation of the actual age axis may uncover an internal age, which is more meaningful from a biological perspective. The effects of shift and elastic registration based on the Fisher-Rao distance are displayed in the middle and right panels of Figure 9 respectively.

Dimensionality reduction

Functional data are infinite-dimensional objects. Even in the case that they are represented by

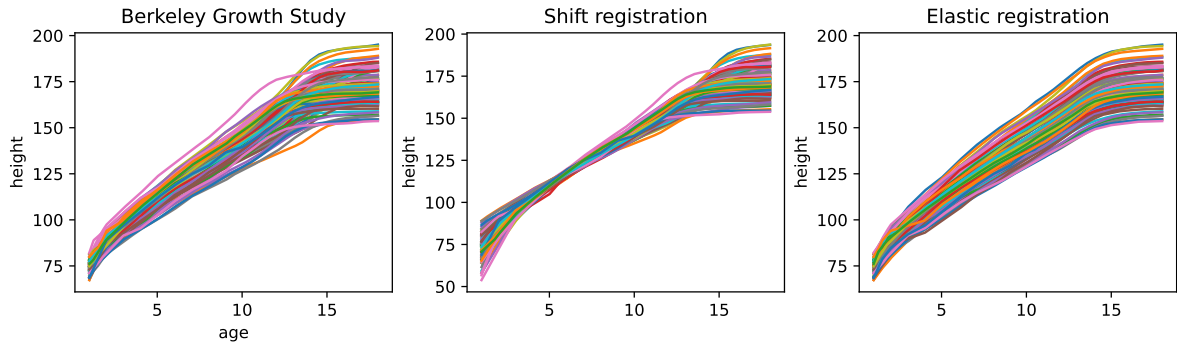


Figure 9: Registration of the *Berkeley Growth Study* data with different methods. From left to right: the original curves, shift registration by least squares and elastic Fisher-Rao registration.

a set of discrete measurements, their dimensionality is typically very high. In addition, nearby observations exhibit a large degree of dependence. Due to these characteristics, technical and computational difficulties arise in the analysis of these types of data. To alleviate such difficulties one can represent the functional data in a lower-dimensional space while preserving as much information as possible. The use of dimensionality reduction methods leads to gains in efficiency and, in some cases, improvements in interpretability and predictive capacity. Furthermore, in this lower-dimensional representation, the methods of multivariate statistics can be employed (Vieu 2018).

A simple dimensionality-reduction method is to select a set of impact or design points that are relevant for the task at hand; for instance, the most informative points for clustering, classification, or regression (Delaigle, Hall, and Bathia 2012; Ferraty, Hall, and Vieu 2010; Kneip, Poss, and Sarda 2016). Specifically, `scikit-fda`'s class `EvaluationTransformer` can be used to evaluate the functions as a set of points in the domain of the function. Alternatively, a truncated basis representation can be used (Biau, Bunea, and Wegkamp 2005; Poskitt and Sengarapillai 2013). The coefficients of a functional data object represented as a basis expansion can be extracted using the methods of the class `CoefficientsTransformer`. Besides these, the `scikit-fda` package provides methods for functional principal components analysis (FPCA) and variable selection methods. These types of methods are described in what follows.

Functional principal components analysis. Functional principal component analysis (FPCA) is a widely used dimensionality reduction method in FDA. In this method, the individual functions are represented in the orthonormal basis of eigenfunctions of the stochastic process' covariance operator. Dimensionality reduction is achieved by retaining the projections of the original functions onto the subspace of L^2 spanned by the set of eigenfunctions that correspond to the largest eigenvalues. This representation is the one, among those of the same dimension, that explains the most of the data's variance.

A random function $X \in L^2$ can be represented as

$$X(t) = \mu(t) + \sum_{j=1}^{\infty} \xi_j \phi_j(t), \quad (3)$$

where $\mu(t) = \mathbb{E}[X(t)]$, $\phi_j(t)$ is the j -th principal component, and $\xi_j = \int_{\mathcal{T}} (X(t) - \mu(t))\phi_j(t)dt$, denotes the projection (score) along the j -th principal component. By the Karhunen-Loève theorem, the scores $\{\xi_j\}_{j \geq 1}$ are uncorrelated random variables (Wang *et al.* 2016). Smoothed versions of the principal components can also be computed by applying the regularization penalties described in Section 2.6, using the procedure described in Section 9.4.2 of Ramsay and Silverman (2005). The smooth principal components are obtained by optimizing a function that takes into account not only the sample variance but also a term that penalizes the roughness of the principal components. A reduction of the dimension of the data can be achieved by truncating the basis expansion in (3), so that only the first K components are included

$$X(t) = \mu(t) + \sum_{j=1}^K \xi_j \phi_j(t).$$

In **scikit-fda**, functional principal component analysis can be carried out using the methods of class ‘FPCA’. The following code illustrates this functionality for the *Berkeley Growth Study* data:

```
>>> import skfda
>>> import matplotlib.pyplot as plt
>>> X, y = skfda.datasets.fetch_growth(return_X_y=True)
>>> fpca = skfda.preprocessing.dim_reduction.FPCA(n_components=2)
>>> fpca.fit(X)
>>> skfda.exploratory.visualization.FPCAPlot(X.mean(),
...     fpca.components_, factor=30).plot()
>>> scores = fpca.transform(X)
>>> scores_class_0 = scores[y == 0]
>>> scores_class_1 = scores[y == 1]
>>> plt.figure()
>>> plt.scatter(scores_class_0[:, 0], scores_class_0[:, 1])
>>> plt.scatter(scores_class_1[:, 0], scores_class_1[:, 1])
```

In this example, the first two principal components are computed. Then, the functional observations are projected onto the two-dimensional subspace spanned by these components. A numerical quadrature is used to compute the corresponding inner products. The class ‘FPCAPlot’ is used to display the curves $\{\mu(t) \pm \phi_j(t); j = 1, 2\}$, which are the result of adding and subtracting the j -th eigenfunction to the sample mean. In the case of the Berkely growth study, the first component captures overall deviations (either positive or negative) with respect to the mean. The second one reveals patterns associated to differences of growth speed. In particular, it exhibits a maximum followed by a sign change at around puberty. The resulting plots are shown in the left and middle panels in Figure 10. Finally, the projection of the curves onto the first two principal components is obtained using the `transform` method of the class ‘FPCA’. The scores of these two components are displayed as points in the right panel of Figure 10. Note that, with some exceptions, boys (blue) and girls (orange) appear grouped in two separate clusters in this plot.

Variable selection. The package **scikit-fda** provides a variety of tools to carry out variable selection. A simple approach is to apply a multivariate variable selection method to the

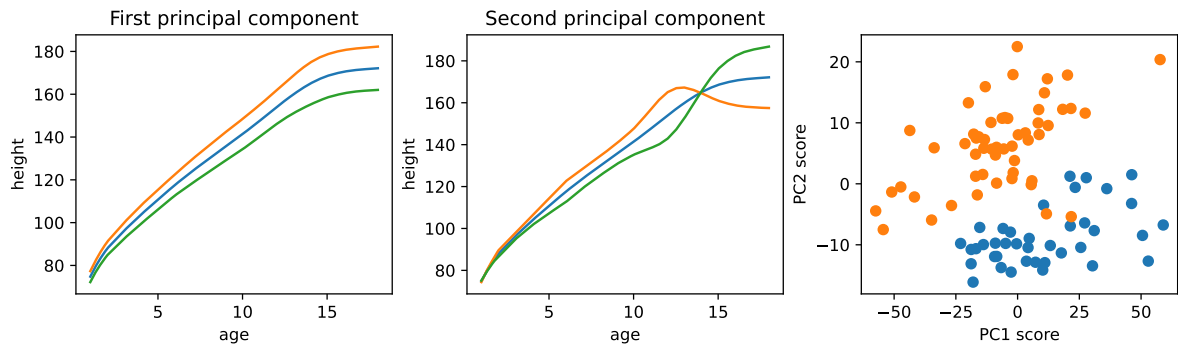


Figure 10: Principal components analysis for the *Berkeley Growth Study* data. The first and second principal components are plotted in the leftmost panel as perturbations around the mean (in blue). In the right panel, the scores of individual functional observations for the first two components are plotted. The orange and blue points correspond to girls' and boys' growth curves, respectively.

discretized representation of the functional observations (Berrendero, Cuevas, and Torrecilla 2016a; Jiménez-Cordero and Maldonado 2021). The `scikit-fda` class 'EvaluationTransformer' can be used to transform 'FData' objects into `NumPy` arrays. Then, any Python library for multivariate variable selection can be used. If this approach does not take into account the functional nature of the data, there can be difficulties in the analysis (Aneiros and Vieu 2016). A multivariate method that takes into account the redundancy that arises from the continuity of functional data is minimum-redundancy-maximum-relevance (mRMR) (Ding and Peng 2005; Peng, Long, and Ding 2005; Berrendero *et al.* 2016a). This method is implemented in `scikit-fda` in the class 'MinimumRedundancyMaximumRelevance'. The dependence measures that quantify the relevance and the redundancy can be specified by the user.

In addition, the `scikit-fda` package includes a collection of variable selection methods that specifically take into account the functional nature of the data: a method based on the theory of Reproducing Kernel Hilbert Spaces (RKHS-VS), maxima hunting (MH), and recursive maxima hunting (RMH).

The RKHS-VS method, implemented in the class 'RKHSVariableSelection', was introduced for binary classification problems (Berrendero, Cuevas, and Torrecilla 2018). For a specified value of d , the goal is to identify the set of points $\mathbf{t} = (t_1, \dots, t_d)^\top \in \mathcal{T}^d$ and select the corresponding function values, $X(t_1), \dots, X(t_d)$, that maximize the Mahalanobis distance between groups

$$(\mu_1(\mathbf{t}) - \mu_0(\mathbf{t}))^\top K(\mathbf{t}, \mathbf{t})^{-1} (\mu_1(\mathbf{t}) - \mu_0(\mathbf{t})),$$

where $\mu_0(\mathbf{t})$, $\mu_1(\mathbf{t})$, and $K(\mathbf{t}, \mathbf{t})$ are the mean functions of each class and the covariance function evaluated at $X(t_1), \dots, X(t_d)$, respectively. In homoscedastic binary classification problems, with a fixed dimension d this selection is optimal in terms of classification error. In practice, the exploration of all possible combinations is often infeasible. To reduce the computational costs, a greedy search is implemented.

In MH (Berrendero, Cuevas, and Torrecilla 2016b; Ordóñez, Oviedo de la Fuente, Roca-Pardiñas, and Rodríguez-Pérez 2018) one selects the values of $t \in \mathcal{T}$ that correspond to local maxima of a non-negative dependence measure between $X(t)$ and the class label. The selected

variables are thus the most relevant in a region. Furthermore, the values of $X(t)$ that are close to these local maxima, which generally provide redundant information, are automatically discarded. In Berrendero *et al.* (2016b), the distance correlation (Székely, Rizzo, and Bakirov 2007) is used as the dependence measure. This variable selection method is implemented in the class ‘MaximaHunting’, using the implementation of distance correlation function provided by the **dcor** package (Ramos-Carreño 2020). MH is an interpretable, fully functional method with optimal performance in an important class of functional classification problems. In spite of its simplicity and good performance, MH has some limitations. Specifically, there can be numerical difficulties to identify the local maxima of the dependence measure. Furthermore, MH takes into account only the marginal relevance of a single variable. Variables that are only relevant when selected in combination with other cannot be identified by these procedures.

RMH (Torrecilla and Suárez 2016) addresses these limitations by assuming a particular form of the stochastic process from which the trajectories are sampled. The algorithm proceeds as follows: First the value of t that is the global maximum of the dependence between the variable $X(t)$ and the class label is selected. Let t^* be such optimum and, therefore, $X(t^*)$ the variable selected. The information conveyed by $X(t^*)$ is removed by subtracting from the trajectories the conditional expectation of the process given the value of the selected variable. Then, the global maximum of the resulting process is identified. The algorithm proceeds in this iterative manner until a pre-specified number of variables have been selected, or until a convergence criterion is fulfilled. In **scikit-fda**, the class ‘RecursiveMaximaHunting’ provides an enhanced, very customizable implementation of this method.

3.4. Exploratory analysis

Exploratory analysis methods are used to identify salient features, visualize, and describe the data from a statistical point of view. Specifically, the **scikit-fda** package provides tools for the computation of summary statistics, including robust ones, interactive tools for visual analysis, and outlier detection.

Summary statistics

Common summary statistics, such as the sample mean function and the sample covariance function can be estimated using the tools provided by **scikit-fda**. Consider a set of functional observations $\{x_i(t)\}_{i=1}^N$. The sample mean,

$$\hat{\mu}(t) = \frac{1}{N} \sum_{i=1}^N x_i(t),$$

can be computed by applying the function `mean()` to the ‘FData’ object in which the data are stored. The functional observations can be either in discrete form or in a basis representation. The resulting mean function is a ‘FData’ object of the same type as the input (i.e., discretized or in a basis representation).

The sample covariance function \hat{k} ,

$$\hat{k}(t, s) = \frac{1}{N-1} \sum_{i=1}^N (x_i(t) - \hat{\mu}(t))(x_i(s) - \hat{\mu}(s)),$$

can be computed applying the function `cov()` to the corresponding ‘FData’ object. Similarly, the function `var()` can be used to compute the sample variance, $\hat{k}(t, t)$. Irrespective of the

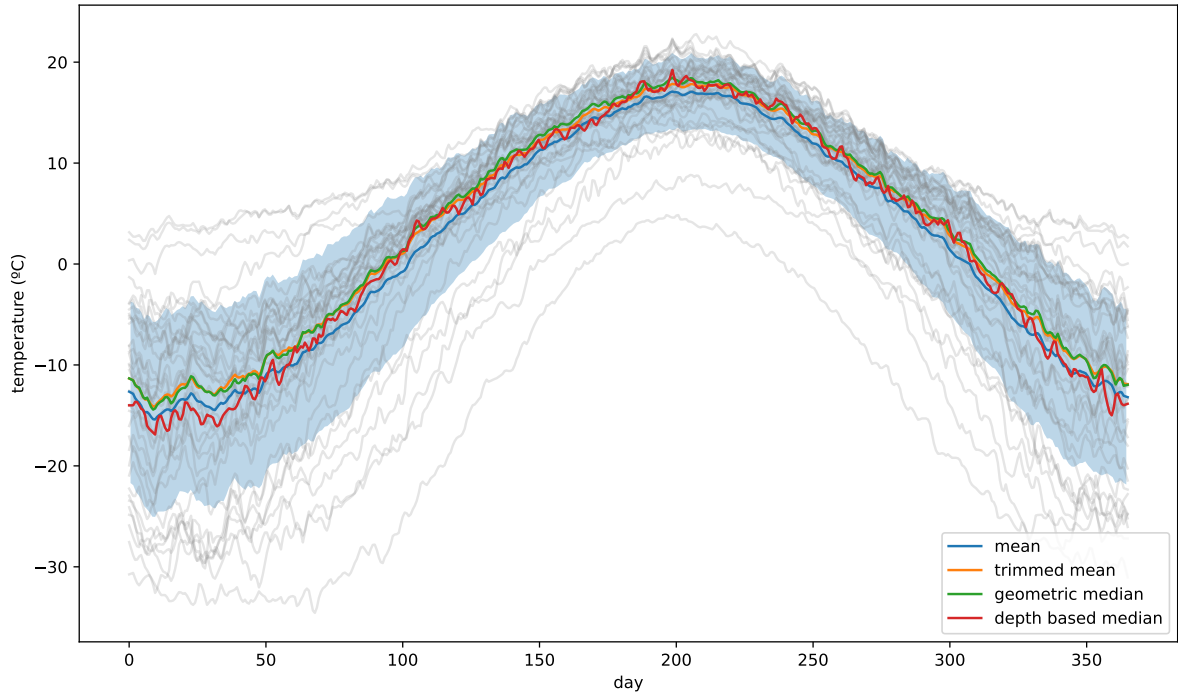


Figure 11: Centrality statistics of the *Canadian Weather* dataset. The shaded band corresponds to one standard deviation around the mean.

representation of the functional observations, the sample variance and covariance are returned in discretized form. Instead of the functions `mean()`, `cov()` and `var()`, the ‘FData’ methods of the same name can be used to compute these summary statistics.

Figure 11 presents an illustration of this functionality for the *Canadian Weather* dataset. The sample estimate of the mean temperature curves is shown as a blue curve. The shaded area corresponds to one standard deviation around the estimated mean. Other measures of centrality, such as the trimmed mean, the geometric mean, and the median, are displayed in this figure as well. These robust statistics will be described in some detail in Section 3.4.3.

Depth measures

Depth measures quantify the centrality of a function in relation to a set of functions. These measures are used for exploratory analysis, to compute robust statistics, detect outliers, and for data visualization (e.g., the functional box-plot). In contrast to the univariate case, a variety of definitions of functional depth can be given that yield different orderings of the functional observations in the sample. Each of these functional depths lead to different definitions of robust statistics and of degrees of outlyingness.

Some common functional depth measures are implemented in the package **scikit-fda**. In particular, the methods of class ‘`IntegratedDepth`’ can be used to compute integrated depth measures (Fraiman and Muniz 2001), which are averages of univariate depths. Specifically, the integrated depth of the function x is

$$\text{ID}(x) = \int_{\mathcal{T}} D(x(t)) dt.$$

where D is an univariate depth function. This function can be selected by the user. The default is the measure proposed by [Fraiman and Muniz \(2001\)](#):

$$D(x(t)) = 1 - \left| \frac{1}{2} - F_{X(t)}(x(t)) \right|,$$

where $F_{X(t)}$ denotes the distribution function of the marginal.

An alternative definition is the band depth (BD), introduced by [López-Pintado and Romo \(2009\)](#). To compute this functional depth, one needs to identify the bands that are delimited by all possible pairs of functional observations in the sample. The BD value is the fraction of bands that completely encompass the curve. In **scikit-fda**, this quantity can be computed using methods of the class `'BandDepth'`. A related, less restrictive measure, is the modified band depth (MBD). This measure takes into account not only the number of bands that contain x , but also the time that x lies within each band. The MBD has better statistical properties than the original BD, in part because it is an integrated depth measure ([Nagy, Gijbels, Omelka, and Hlubinka 2016](#)). In **scikit-fda**, MBD is implemented in the class `'ModifiedBandDepth'`.

Robust statistics

The package **scikit-fda** provides support for the computation of robust statistics. Robust statistics may provide a better characterization of the data than non-robust ones (e.g., the mean or the covariance functions), especially in the presence of outliers. One of the most important robust statistics is the geometric median ([Lardin-Puech, Cardot, and Goga 2014](#))

$$\text{median} = \arg \min_{z \in \mathcal{F}} \sum_{i=1}^N \|x_i - z\|.$$

It can be computed with the function `geometric_median()`. Alternatively, the median can be defined as the deepest point in the sample. Different depth measures yield different definitions of the median. These types of medians can be computed with the function `depth_based_median()`.

Functional depth measures can be used also to define the degree of outlyingness of a function in a sample: The larger the depth value the more central the functional observation is. Finally, functional depth measures can be used to define trimmed means ([Fraiman and Muniz 2001](#)). A trimmed mean is a robust version of the standard mean in which the most outlying functional observations (the ones with the lowest depth values) are discarded. In **scikit-fda**, the trimmed mean is implemented in function `trim_mean()`.

The geometric median, the MBD-based median, and the MBD-based trimmed mean in which 10% of the data are discarded, of the *Canadian Weather* dataset are shown in Figure 11.

Interactive visualization tools and outlier detection

Visualization tools can be used to gain insight into the data. In particular, trends, salient features, and other patterns in the data can be identified simply by inspection. Visualization tools can be utilized also to single out functional observations that are markedly different from the other observations in the sample (outliers). Outlier detection is useful to identify rare events, novel patterns, anomalies, or erroneous measurements. The package **scikit-fda** provides a number of interactive tools for data visualization and outlier detection. Their implementation utilizes the functionality provided by **matplotlib** ([Hunter 2007](#)).

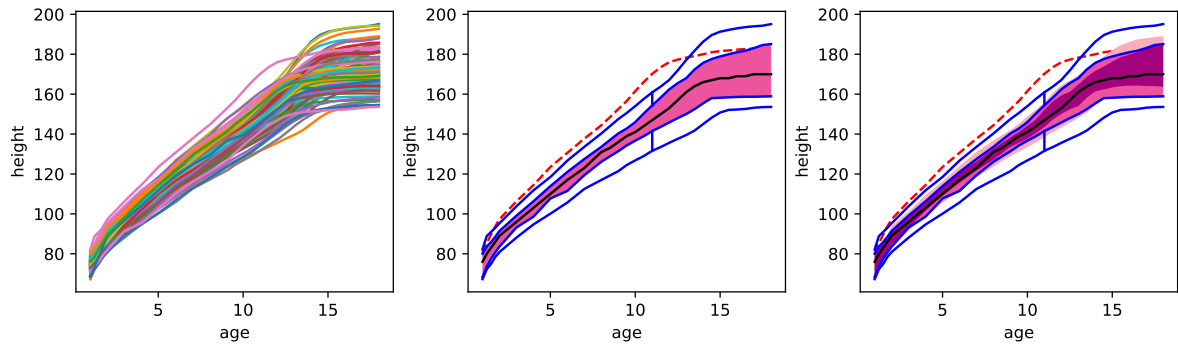


Figure 12: Functional boxplots of the *Berkeley Growth Study* dataset. The original curves are depicted in the left panel. The standard functional boxplot is shown in the central panel. In this panel, The black line stands for the functional median. The central envelope is displayed as a pink band around the median. The blue whiskers and their fences mark the maximum non-outlying envelope. Outliers are shown as a red dashed lines. In the right panel, different shades of pink are used for the deepest 25%, 50%, and 75% of the data.

Functional data objects have a `plot()` method that can be used to graph the curves. Some customization options, such as group colors or labels, are available for this method. An illustration of its use with the *Berkeley Growth Study* dataset is shown in the left panel of Figure 12. For ‘`FDataGrid`’ objects, the `scatter()` method can be used to display the values of the function as individual points in a graph. This method was used to generate Figure 1.

Another tool for visual exploration provided by **scikit-fda** is the functional boxplot (Sun and Genton 2011). This is an generalization of the univariate boxplot for functional data. The functional boxplot consists of a graph of the functional median (i.e., the deepest curve in the sample) surrounded by a central envelope, which encompasses the deepest 50% of the observations, and a maximum non-outlying envelope. The width of this outer envelope is determined by scaling the central one by a constant factor. This constant factor can be selected by the user. Its default value is 1.5. In **scikit-fda**, the class ‘`Boxplot`’ can be used to generate and customize functional boxplots. In this plot, a trajectory is marked as an outlier if it lies beyond the maximum non-outlying envelope for some interval. The class ‘`BoxplotOutlierDetector`’ can be used for outlier detection based on this criterion. Some customizable elements of ‘`Boxplot`’ objects are the depth measure, and the definition of centered bands that encompasses a user-specified fraction of the deepest observations. The following code provides an illustration of these functionalities with the *Berkeley Growth Study* dataset. The plots that result from the execution of this code are displayed in Figure 12.

```
>>> import skfda
>>> X, _ = skfda.datasets.fetch_growth(return_X_y=True)
>>> X.plot()
>>> boxplot = skfda.exploratory.visualization.Boxplot(X,
...   depth_method=skfda.exploratory.depth.ModifiedBandDepth())
>>> boxplot.plot()
>>> boxplot = skfda.exploratory.visualization.Boxplot(X,
...   depth_method=skfda.exploratory.depth.ModifiedBandDepth(),
...   prob=[0.75, 0.5, 0.25])
```

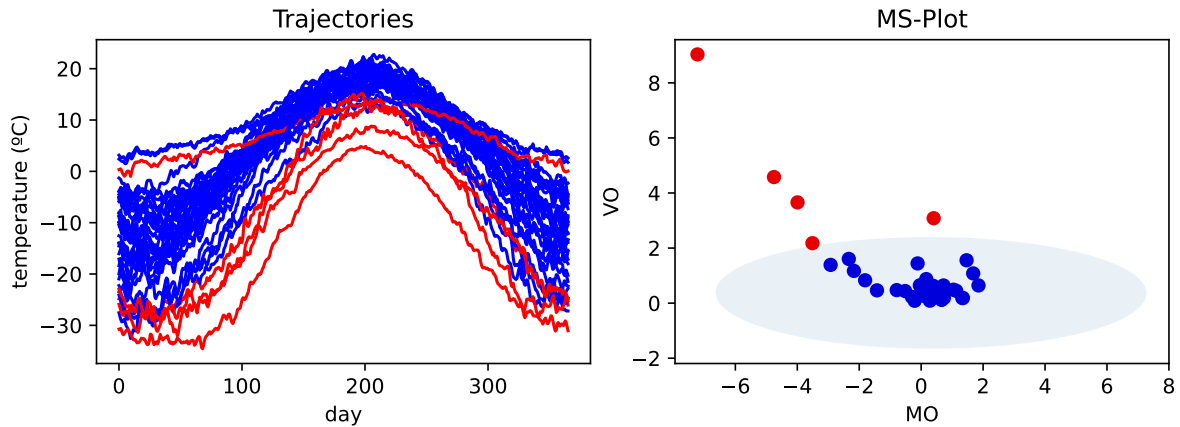


Figure 13: MS-plot and outliers of the *Canadian Weather* dataset. The original yearly temperature curves are displayed in the left panel. The corresponding MS-plot is shown in the right panel. The observations identified as outliers by the MS-plot criterion (those outside the blue ellipse) are marked in red.

```
>>> boxplot.plot()
```

Another tool for functional data visualization and outlier detection is the magnitude-shape plot (MS-plot) (Dai and Genton 2018, 2019). In this method, the degree of outlyingness of a functional observation is characterized in terms of two quantities: The magnitude outlyingness (MO) and the shape outlyingness (VO). The MS-plot is the scatter plot of the values MO and VO for each functional observation. This two-dimensional representation of the data can be used, for instance, to identify clusters of functions, or detect potential outliers, either in shape or in magnitude.

The following code can be used to display the MS-plot for the temperature curves of the *Canadian Weather* dataset together with the original trajectories. Additionally, outliers are identified according to the MS-plot criterion and marked in red. The class ‘MagnitudeShapePlot’ generates the MS-plot and uses internally the methods of the class ‘MSPlotOutlierDetector’ for outlier detection. The resulting plots are shown in Figure 13.

```
>>> import skfda
>>> X, y = skfda.datasets.fetch_weather(return_X_y=True)
>>> X = X.coordinates[0]
>>> ms_plot = skfda.exploratory.visualization.MagnitudeShapePlot(X)
>>> ms_plot.plot()
>>> fig = X.plot(group=ms_plot.outliers, group_colors=["blue", "red"])
```

The class ‘Outliergram’ provides an additional method for data visualization and detection of shape outliers (Arribas-Gil and Romo 2014). The graph is defined in terms of two related quantities: The modified epigraph index (MEI) and the MBD. The MEI of a trajectory is the average over time of the fraction of curves in the sample that lie above it. Each curve is a point (MEI, MBD) in the scatter plot. The outliergram takes advantage of the fact that points corresponding to typical functional observations lie on a parabola, whose analytical form is known. This parabola is used as a reference for the identification of shape outliers.

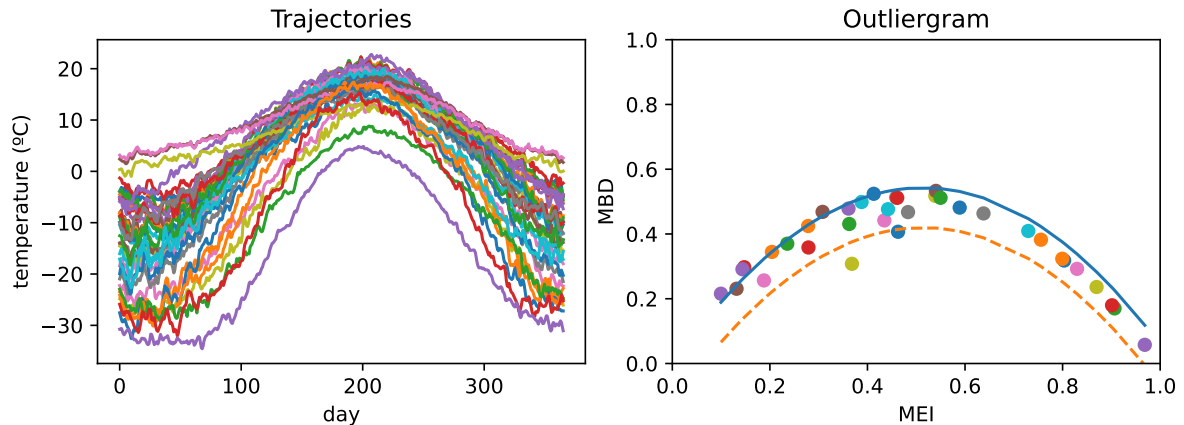


Figure 14: Outliergram of the yearly temperature curves for the *Canadian Weather* dataset. The original trajectories are in the left panel and the corresponding outliergram is shown in the right panel. The blue line corresponds to the reference parabola. The orange dashed line separates the typical curves (above) from the outliers (below).

Specifically, the degree of outlyingness of a curve is quantified in terms of its vertical distance to the parabola. The **scikit-fda**'s classes `'Outliergram'` and `'OutliergramOutlierDetection'` can be used to generate the outliergram and to detect outliers by using this criterion, respectively. The following code illustrates this functionality with the temperatures of the *Canadian Weather* dataset. The original trajectories and the corresponding outliergram are shown in Figure 14.

```
>>> import skfda
>>> import matplotlib.pyplot as plt
>>> X, y = skfda.datasets.fetch_weather(return_X_y=True)
>>> X = X.coordinates[0]
>>> fig = X.plot()
>>> fig = skfda.exploratory.visualization.Outliergram(X).plot()
```

In addition to standard plotting capabilities, most graphs generated with **scikit-fda** incorporate some interactive features. For example, the cursor can be placed at a point in the graph to display the actual coordinate values and the label of the observation. In addition, if different plots are used for visual exploration of some functional dataset, selecting a particular curve in one plot highlights the corresponding curve in the other active plots. Finally, widgets such as sliders can be used to select curves by some property, such as the label of the observation, or their depth in the sample.

An illustration of this interactive functionality is presented in Figure 15. In this figure, three different kinds of plots are displayed for the temperature curves of the *Canadian Weather* dataset: A graph of the sample trajectories, the MS-plot, and the outliergram. In the lower right side a slider has been created that displays the MBD value of the selected curve. A functional observation can be selected either by choosing a value in the slider widget or by clicking on the corresponding point in the MS-plot or in the outliergram. The datum selected is then highlighted in all plots. Finally, the cursor has been placed at a point in the MS-

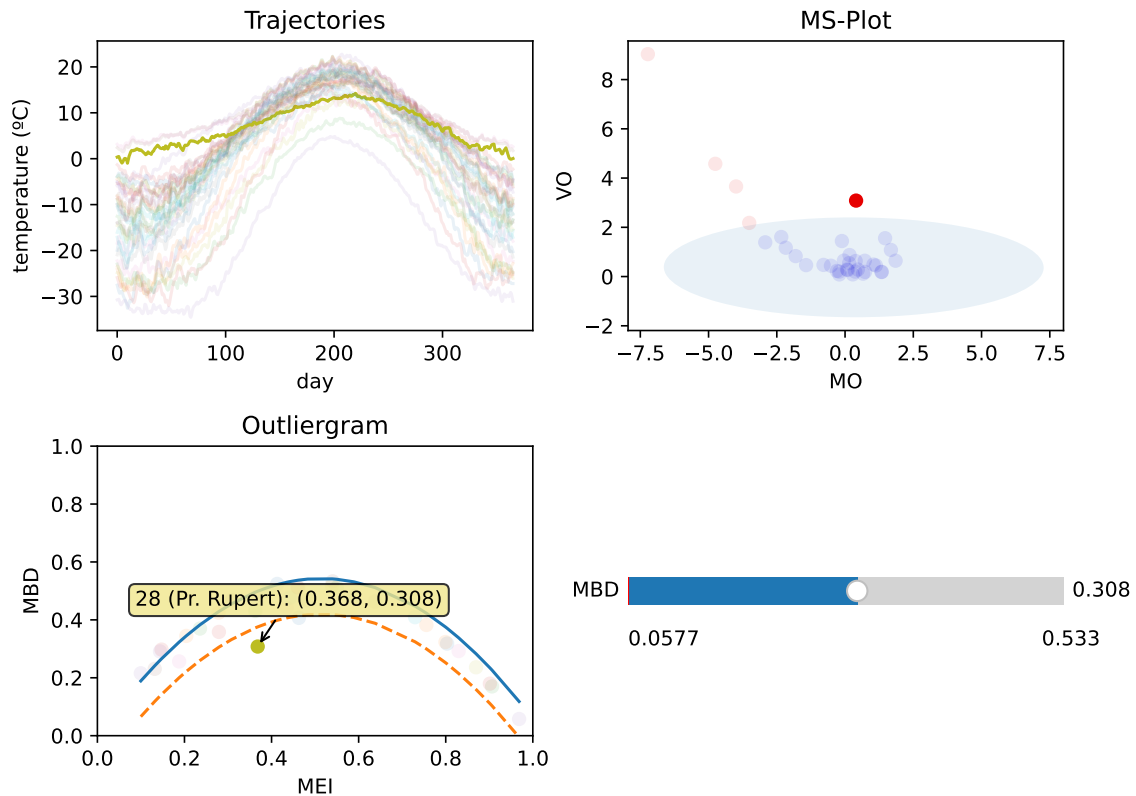


Figure 15: Interactive features in multiple plots for the *Canadian Weather* dataset.

plot. This brings up a tool-tip in which relevant information of the corresponding functional observation is displayed.

3.5. Integration with scikit-learn for machine learning

The **scikit-fda** package has been especially designed for seamless integration with **scikit-learn** (Pedregosa *et al.* 2011). Specifically, there are a number of methods that transform the functional data into a two-dimensional array so that **scikit-learn**'s machine learning algorithms can be applied. For instance, the method `transform()` of the class `EvaluationTransformer` returns an array with the values of the functions in the sample at a specified set of points. The coefficients of a functional data object represented as a basis expansion can be extracted using the methods of the class `CoefficientsTransformer`. Additionally, other **scikit-fda** methods, such as variable selection, can be used to this end.

The provided classes and methods for preprocessing conform to **scikit-learn**'s application programming interface (API) (Buitinck *et al.* 2013). An advantage of adopting this standard is that they can be employed in a **scikit-learn** pipeline (class `Pipeline`). The following code illustrates how to build such a pipeline for a classification problem with functional data:

```
>>> import skfda
>>> from sklearn.model_selection import GridSearchCV, train_test_split
>>> from sklearn.pipeline import Pipeline
>>> from sklearn.svm import SVC
```

```

>>> import skfda.preprocessing.smoothing as smoothing
>>> import skfda.preprocessing.dim_reduction as dimred
>>> from skfda.misc.hat_matrix import KNeighborsHatMatrix
>>> X, y = skfda.datasets.fetch_phoneme(return_X_y=True)
>>> X_train, X_test, y_train, y_test = train_test_split(X, y,
...     random_state=0)
>>> smoothing_step = smoothing.KernelSmoother(
...     kernel_estimator=KNeighborsHatMatrix())
>>> dimred_step = dimred.FPCA(n_components=3)
>>> classification_step = SVC()
>>> pipeline = Pipeline([('smoothing', smoothing_step),
...     ('dimred', dimred_step), ('classification', classification_step)])
>>> grid = GridSearchCV(pipeline, param_grid={
...     'smoothing_kernel_estimator__n_neighbors': [3, 5, 7],
...     'dimred__n_components': [1, 2, 3],
...     'classification__C': [0.001, 0.01, 0.1, 1, 10]})
>>> grid.fit(X_train, y_train)
>>> score = grid.score(X_test, y_test)
>>> print(f"{score:.3}")

```

0.879

In this pipeline, the following sequence of operations is applied: Nearest-neighbors kernel smoothing, dimensionality reduction by functional PCA, and, finally, a standard (multivariate) support vector machine (SVM) classifier with an RBF kernel. The hyperparameters of the whole model, including those that correspond to preprocessing, are tuned by cross-validation using **scikit-learn**'s class `'GridSearchCV'`. Specifically, the number of neighbors in the nearest-neighbors kernel smoother is selected between the values 3, 5, and 7. The number of principal components considered ranges from 1 to 3. The grid [0.001, 0.01, 0.1, 1.0, 10.0] is explored to determine the regularization parameter of the SVM. Then, the best values of the hyperparameters are used to fit the model using the complete training set. Finally, the accuracy of the classifier is computed and printed.

4. Code quality and documentation

The **scikit-fda** library has been designed to provide a powerful, self-contained, flexible, stable, and easy-to-use framework for the analysis of functional data in Python. The package is built as a **SciPy** Toolkit (SciKit, see <https://svn.scipy.org/scikits.html> for further details on SciKits). It is fully integrated in the Scientific Python software ecosystem (<https://scientific-python.org/>). Scientific Python is a collection of free and open-source Python packages for scientific and technical computing (Oliphant 2007; Millman and Aivazis 2011). Standard coding and naming practices are used throughout the project (Van Rossum, Warsaw, and Coghlan 2001; Goodger and Van Rossum 2001). This not only improves the legibility of the code and simplifies its maintenance, but also facilitates external contributions to the development of the package. Whenever appropriate, the design conforms to **scikit-learn** specifications (Pedregosa *et al.* 2011), so that the machine learning tools implemented in that

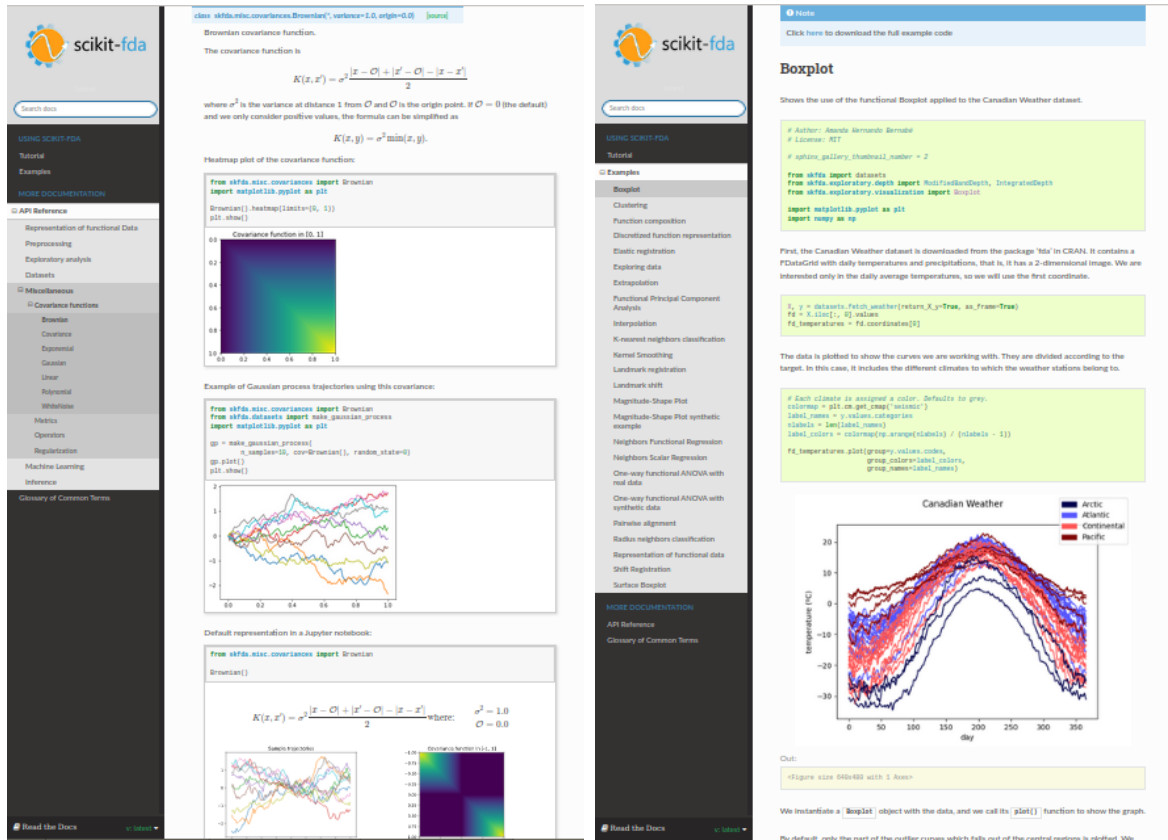


Figure 16: **scikit-fda** documentation: On the left panel, the reference page of the Brownian covariance function is shown. On the right one, an illustration of the functional boxplot functionality.

package can be readily applied to functional data. To ensure the quality and robustness of the software, a comprehensive suite of unit and integration tests is provided. These automated tests are executed regularly in a continuous integration environment. We have attempted also to minimize the number of dependencies and to provide flexible interfaces that are intuitive and consistent throughout the application.

The **scikit-fda** package is free and open-source software distributed under the OSI-approved 3-clause BSD license (<https://opensource.org/licenses/BSD-3-Clause>). The version described in this paper is 0.9.1 and was released on 2024-02-26. The GitHub page <https://github.com/GAA-UAM/scikit-fda> is the main communication channel with the developers of the package for questions, bug reports, and feature requests. Contributions from the members of the FDA community are encouraged, as are comments and suggestions to improve the quality of the software.

To facilitate the use of **scikit-fda**, exhaustive documentation, including installation instructions, tutorials, API references, and illustrative examples are provided. The documentation, which is available online at <https://fda.readthedocs.io/>, is built with the Python tool **Sphinx** (Sphinx Development Team 2020). The examples and tutorials have been devised with **Sphinx-Gallery** (Nájera *et al.* 2020). They can be viewed online or downloaded as interactive **Jupyter** notebooks (Kluyver *et al.* 2016). In Figure 16, two screenshots of the

documentation pages are shown: A reference page for the Brownian covariance function and an example of use of the functional boxplot functionality are displayed on the left and right panels of the figure, respectively.

The package `scikit-fda` is under active development, with new features being continuously incorporated. Users are encouraged to consult the documentation at <https://fda.readthedocs.io/> for an up-to-date description of the project functionality.

Acknowledgments

The authors want to express their gratitude to the developers who made contributions to the `scikit-fda` package. In particular, we thank David García Fernández, Amanda Hernando Bernabé, Yujian Hong, Pedro Martín Rodríguez-Ponga Eyriès, Pablo Pérez Manso, Elena Petrunina, Luis Alberto Rodríguez Ramírez, and Álvaro Sánchez Romero for their participation in the project. The authors acknowledge financial support from the Spanish Ministry of Education and Innovation, projects PID2019-106827GB-I00 / AEI / 10.13039/501100011033 and PID2019-109387GB-I00. This research was also supported by an FPU grant (Formación de Profesorado Universitario) from the Spanish Ministry of Science, Innovation and Universities (MICIU) with reference FPU18/00047.

References

- Aneiros G, Vieu P (2016). “Sparse Nonparametric Model for Regression with Functional Covariate.” *Journal of Nonparametric Statistics*, **28**(4), 839–859. doi:10.1080/10485252.2016.1234050.
- Arribas-Gil A, Romo J (2014). “Shape Outlier Detection and Visualization for Functional Data: The Outliergram.” *Biostatistics*, **15**(4), 603–619. doi:10.1093/biostatistics/kxu006.
- Assunção RV, Silva AC, Roy A, Boursières B, Silva CHS, TERNON JF, Araujo M, Bertrand A (2020). “3D Characterisation of the Thermohaline Structure in the Southwestern Tropical Atlantic Derived from Functional Data Analysis of In Situ Profiles.” *Progress in Oceanography*, **187**, 102399. doi:10.1016/j.pocean.2020.102399.
- Bagnall A, Dau HA, Lines J, Flynn M, Large J, Bostrom A, Southam P, Keogh E (2018). “The UEA Multivariate Time Series Classification Archive, 2018.” *arXiv 1811.00075*, arXiv.org E-Print Archive. doi:10.48550/arXiv.1811.00075.
- Bagnall A, Lines J, Vickers W, Keogh E (2021). “The UEA & UCR Time Series Classification Repository.” URL <https://www.timeseriesclassification.com/>.
- Bernard G, Lejeune C, Ferrières S, Teste O (2021). *Curve Shape Analysis*. GitHub Repository, URL https://github.com/Clej/curve_shape_analysis.
- Berrendero JR, Cuevas A, Torrecilla JL (2016a). “The mRMR Variable Selection Method: A Comparative Study for Functional Data.” *Journal of Statistical Computation and Simulation*, **86**(5), 891–907. doi:10.1080/00949655.2015.1042378.

- Berrendero JR, Cuevas A, Torrecilla JL (2016b). “Variable Selection in Functional Data Classification: A Maxima-Hunting Proposal.” *Statistica Sinica*, **26**(2), 619–638. doi:[10.5705/ss.202014.0014](https://doi.org/10.5705/ss.202014.0014).
- Berrendero JR, Cuevas A, Torrecilla JL (2018). “On the Use of Reproducing Kernel Hilbert Spaces in Functional Classification.” *Journal of the American Statistical Association*, **113**(523), 1210–1218. doi:[10.1080/01621459.2017.1320287](https://doi.org/10.1080/01621459.2017.1320287).
- Beyaztas U, Yaseen ZM (2019). “Drought Interval Simulation Using Functional Data Analysis.” *Journal of Hydrology*, **579**, 124141. doi:[10.1016/j.jhydro1.2019.124141](https://doi.org/10.1016/j.jhydro1.2019.124141).
- Bezanson J, Edelman A, Karpinski S, Shah VB (2017). “Julia: A Fresh Approach to Numerical Computing.” *SIAM Review*, **59**(1), 65–98. doi:[10.1137/141000671](https://doi.org/10.1137/141000671).
- Biau G, Bunea F, Wegkamp MH (2005). “Functional Classification in Hilbert Spaces.” *IEEE Transactions on Information Theory*, **51**(6), 2163–2172. doi:[10.1109/tit.2005.847705](https://doi.org/10.1109/tit.2005.847705).
- Bischl B, Lang M, Kotthoff L, Schiffner J, Richter J, Studerus E, Casalicchio G, Jones ZM (2016). “**mlr**: Machine Learning in R.” *Journal of Machine Learning Research*, **17**(170), 1–5. URL <https://jmlr.org/papers/v17/15-066.html>.
- Bouveyron C, Côme E, Jacques J (2015). “The Discriminative Functional Mixture Model for a Comparative Analysis of Bike Sharing Systems.” *The Annals of Applied Statistics*, **9**(4), 1726–1760. doi:[10.1214/15-aos861](https://doi.org/10.1214/15-aos861).
- Brockhaus S, Rügamer D, Greven S (2020). “Boosting Functional Regression Models with **FDboost**.” *Journal of Statistical Software*, **94**(10), 1–50. doi:[10.18637/jss.v094.i10](https://doi.org/10.18637/jss.v094.i10).
- Buitinck L, Louppe G, Blondel M, Pedregosa F, Mueller A, Grisel O, Niculae V, Prettenhofer P, Gramfort A, Grobler J, Layton R, VanderPlas J, Joly A, Holt B, Varoquaux G (2013). “API Design for Machine Learning Software: Experiences from the **scikit-learn** Project.” In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, pp. 108–122. URL https://www.ecmlpkdd2013.org/wp-content/uploads/2013/09/lml2013_api_sklearn.pdf.
- Chen D, Cremona MA, Qi Z, Mitra RD, Chiaromonte F, Makova KD (2020). “Human L1 Transposition Dynamics Unraveled with Functional Data Analysis.” *Molecular Biology and Evolution*, **37**(12), 3576–3600. doi:[10.1093/molbev/msaa194](https://doi.org/10.1093/molbev/msaa194).
- Consagra W, Venkataraman A, Qiu X (2022). “Efficient Multidimensional Functional Data Analysis Using Marginal Product Basis Systems.” *arXiv 2107.14728*, arXiv.org E-Print Archive. doi:[10.48550/arXiv.2107.14728](https://doi.org/10.48550/arXiv.2107.14728).
- Cremona MA, Xu H, Makova KD, Reimherr M, Chiaromonte F, Madrigal P (2019). “Functional Data Analysis for Computational Biology.” *Bioinformatics*, **35**(17), 3211–3213. doi:[10.1093/bioinformatics/btz045](https://doi.org/10.1093/bioinformatics/btz045).
- Cuevas A (2014). “A Partial Overview of the Theory of Statistics with Functional Data.” *Journal of Statistical Planning and Inference*, **147**, 1–23. doi:[10.1016/j.jspi.2013.04.002](https://doi.org/10.1016/j.jspi.2013.04.002).

- Dai W, Genton MG (2018). “Multivariate Functional Data Visualization and Outlier Detection.” *Journal of Computational and Graphical Statistics*, **27**(4), 923–934. doi:[10.1080/10618600.2018.1473781](https://doi.org/10.1080/10618600.2018.1473781).
- Dai W, Genton MG (2019). “Directional Outlyingness for Multivariate Functional Data.” *Computational Statistics & Data Analysis*, **131**, 50–65. doi:[10.1016/j.csda.2018.03.017](https://doi.org/10.1016/j.csda.2018.03.017).
- Dau HA, Bagnall A, Kamgar K, Yeh CCM, Zhu Y, Gharghabi S, Ratanamahatana CA, Keogh E (2019). “The UCR Time Series Archive.” *IEEE/CAA Journal of Automatica Sinica*, **6**(6), 1293–1305. doi:[10.1109/jas.2019.1911747](https://doi.org/10.1109/jas.2019.1911747).
- Delaigle A, Hall P, Bathia N (2012). “Componentwise Classification and Clustering of Functional Data.” *Biometrika*, **99**(2), 299–313. doi:[10.1093/biomet/ass003](https://doi.org/10.1093/biomet/ass003).
- Díaz-Vico D, Ramos-Carreño C (2022). “**scikit-datasets**: **scikit-learn**-Compatible Datasets.” doi:[10.5281/zenodo.6383047](https://doi.org/10.5281/zenodo.6383047).
- Ding C, Peng H (2005). “Minimum Redundancy Feature Selection from Microarray Gene Expression Data.” *Journal of Bioinformatics and Computational Biology*, **03**(02), 185–205. doi:[10.1142/s0219720005001004](https://doi.org/10.1142/s0219720005001004).
- Febrero-Bande M, Oviedo de la Fuente M (2012). “Statistical Computing in Functional Data Analysis: The R Package **fda.usc**.” *Journal of Statistical Software*, **51**(4), 1–28. doi:[10.18637/jss.v051.i04](https://doi.org/10.18637/jss.v051.i04).
- Fermanian A (2022). “Functional Linear Regression with Truncated Signatures.” *Journal of Multivariate Analysis*, **192**, 105031. doi:[10.1016/j.jmva.2022.105031](https://doi.org/10.1016/j.jmva.2022.105031).
- Ferrando L, Ventura-Campos N, Epifanio I (2020). “Detecting and Visualizing Differences in Brain Structures with SPHARM and Functional Data Analysis.” *NeuroImage*, **222**, 117209. doi:[10.1016/j.neuroimage.2020.117209](https://doi.org/10.1016/j.neuroimage.2020.117209).
- Ferraty F, Hall P, Vieu P (2010). “Most-Predictive Design Points for Functional Data Predictors.” *Biometrika*, **97**(4), 807–824. doi:[10.1093/biomet/asq058](https://doi.org/10.1093/biomet/asq058).
- Ferraty F, Vieu P (2006). *Nonparametric Functional Data Analysis: Theory and Practice*. Springer-Verlag, New York. doi:[10.1007/0-387-36620-2](https://doi.org/10.1007/0-387-36620-2).
- Fraiman R, Muniz G (2001). “Trimmed Means for Functional Data.” *Test*, **10**(2), 419–440. doi:[10.1007/bf02595706](https://doi.org/10.1007/bf02595706).
- Frois Caldeira J, Gupta R, Suleman MT, Torrent HS (2020). “Forecasting the Term Structure of Interest Rates of the BRICS: Evidence from a Nonparametric Functional Data Analysis.” *Emerging Markets Finance and Trade*, **0**(0), 1–18. doi:[10.1080/1540496x.2020.1808458](https://doi.org/10.1080/1540496x.2020.1808458).
- Goldsmith J, Scheipl F, Huang L, Wrobel J, Di C, Gellar J, Harezlak J, McLean MW, Swihart B, Xiao L, Crainiceanu C, Reiss PT (2024). **refund**: *Regression with Functional Data*. R package version 0.1-35, URL <https://CRAN.R-project.org/package=refund>.
- Golovkine S (2021). “**FDAPy**: A Python Package for Functional Data.” *arXiv 2101.11003*, arXiv E-Print Archive. doi:[10.48550/arXiv.2101.11003](https://doi.org/10.48550/arXiv.2101.11003).

- Gong X, Wang Y, Lin B (2021). “Assessing Dynamic China’s Energy Security: Based on Functional Data Analysis.” *Energy*, **217**, 119324. doi:10.1016/j.energy.2020.119324.
- Goodger D, Van Rossum G (2001). “PEP 257: Docstring Conventions.” Accessed 2020-09-28, URL <https://www.python.org/dev/peps/pep-0257/>.
- Green PJ, Silverman BW (1993). *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. Chapman & Hall/CRC. doi:10.1201/b15710.
- Happ-Kurz C (2020). “Object-Oriented Software for Functional Data.” *Journal of Statistical Software*, **93**(1), 1–38. doi:10.18637/jss.v093.i05.
- Harris CR, Millman KJ, Van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C, Oliphant TE (2020). “Array Programming with NumPy.” *Nature*, **585**(7825), 357–362. doi:10.1038/s41586-020-2649-2.
- Hastie T, Tibshirani R, Friedman J (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd edition. Springer-Verlag, New York. doi:10.1007/978-0-387-84858-7.
- Horsley KJ, Ramsay JO, Ditto B, Da Costa D (2021). “Maternal Blood Pressure Trajectories and Associations with Gestational Age at Birth: A Functional Data Analytic Approach.” *Journal of Hypertension*, **40**(2), 213–220. doi:10.1097/hjh.0000000000002995.
- Hsing T, Eubank R (2015). *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators*. John Wiley & Sons, Chichester. doi:10.1002/9781118762547.
- Hu X, Yuan Y, Zhu X, Yang H, Xie K (2019). “Behavioral Responses to Pre-Planned Road Capacity Reduction Based on Smartphone GPS Trajectory Data: A Functional Data Analysis Approach.” *Journal of Intelligent Transportation Systems*, **23**(2), 133–143. doi:10.1080/15472450.2018.1488133.
- Hunter JD (2007). “matplotlib: A 2D Graphics Environment.” *Computing in Science & Engineering*, **9**(3), 90–95. doi:10.1109/mcse.2007.55.
- Hyndman RJ, Shahid Ullah M (2007). “Robust Forecasting of Mortality and Fertility Rates: A Functional Data Approach.” *Computational Statistics & Data Analysis*, **51**(10), 4942–4956. doi:10.1016/j.csda.2006.07.028.
- Hyndman RJ, Shang HL (2010). “Rainbow Plots, Bagplots, and Boxplots for Functional Data.” *Journal of Computational and Graphical Statistics*, **19**(1), 29–45. doi:10.1198/jcgs.2009.08158.
- Ieva F, Paganoni AM, Romo J, Tarabelloni N (2019). “roahd Package: Robust Analysis of High Dimensional Data.” *The R Journal*, **11**(2), 291–307. doi:10.32614/rj-2019-032.
- Jiménez-Cordero A, Maldonado S (2021). “Automatic Feature Scaling and Selection for Support Vector Machine Classification with Functional Data.” *Applied Intelligence*, **51**(1), 161–184. doi:10.1007/s10489-020-01765-6.

- Joshi SH, Klassen EP, Srivastava A, Jermyn I (2007). “A Novel Representation for Riemannian Analysis of Elastic Curves in \mathbb{R}^n .” In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–7. doi:[10.1109/cvpr.2007.383185](https://doi.org/10.1109/cvpr.2007.383185).
- Kluyver T, Ragan-Kelley B, Pérez F, Granger B, Bussonnier M, Frederic J, Kelley K, Hamrick J, Grout J, Corlay S, Ivanov P, Avila D, Abdalla S, Willing C (2016). “**Jupyter** Notebooks – A Publishing Format for Reproducible Computational Workflows.” In F Loizides, B Schmidt (eds.), *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pp. 87–90. IOS Press. doi:[10.3233/978-1-61499-649-1-87](https://doi.org/10.3233/978-1-61499-649-1-87).
- Kneip A, Poss D, Sarda P (2016). “Functional Linear Regression with Points of Impact.” *The Annals of Statistics*, **44**(1), 1–30. doi:[10.1214/15-aos1323](https://doi.org/10.1214/15-aos1323).
- Lardin-Puech P, Cardot H, Goga C (2014). “Analysing Large Datasets of Functional Data: a Survey Sampling Point of View.” *Journal de la Société Française de Statistique*, **155**(4), 70–94. URL http://www.numdam.org/item/JSFS_2014__155_4_70_0/.
- Leng X, Müller HG (2006). “Classification Using Functional Data Analysis for Temporal Gene Expression Data.” *Bioinformatics*, **22**(1), 68–76. doi:[10.1093/bioinformatics/bti742](https://doi.org/10.1093/bioinformatics/bti742).
- López-Pintado S, Romo J (2009). “On the Concept of Depth for Functional Data.” *Journal of the American Statistical Association*, **104**(486), 718–734. doi:[10.1198/jasa.2009.0108](https://doi.org/10.1198/jasa.2009.0108).
- Millman KJ, Aivazis M (2011). “Python for Scientists and Engineers.” *Computing in Science Engineering*, **13**(2), 9–12. doi:[10.1109/mcse.2011.36](https://doi.org/10.1109/mcse.2011.36).
- Nagy S, Gijbels I, Omelka M, Hlubinka D (2016). “Integrated Depth for Functional Data: Statistical Properties and Consistency.” *ESAIM: Probability and Statistics*, **20**, 95–130. doi:[10.1051/ps/2016005](https://doi.org/10.1051/ps/2016005).
- Nájera Ó, Larson E, Estève L, Varoquaux G, Liu L, Grobler J, de Andrade ES, Holdgraf C, Gramfort A, Jas M, Nothman J, Grisel O, Varoquaux N, Gouillart E, Luessi M, Lee A, Vanderplas J, Hoffmann T, Caswell TA, Sullivan B, Batula A, jaeilepp, Robitaille T, Appelhoff S, Kunzmann P, Geier M, Lars, Sunden K, Stańczak D, Shih AY (2020). **Sphinx-Gallery: Release V0.7.0**. doi:[10.5281/zenodo.3741780](https://doi.org/10.5281/zenodo.3741780).
- Ojo OT, Lillo RE, Fernandez Anta A (2023). **fdaoutlier: Outlier Detection Tools for Functional Data Analysis**. R package version 0.2.1, URL <https://CRAN.R-project.org/package=fdaoutlier>.
- Oliphant TE (2007). “Python for Scientific Computing.” *Computing in Science Engineering*, **9**(3), 10–20. doi:[10.1109/mcse.2007.58](https://doi.org/10.1109/mcse.2007.58).
- Ordóñez C, Oviedo de la Fuente M, Roca-Pardiñas J, Rodríguez-Pérez JR (2018). “Determining Optimum Wavelengths for Leaf Water Content Estimation from Reflectance: A Distance Correlation Approach.” *Chemometrics and Intelligent Laboratory Systems*, **173**, 41–50. doi:[10.1016/j.chemolab.2017.12.001](https://doi.org/10.1016/j.chemolab.2017.12.001).
- pandas Development Team (2020). *pandas-dev/pandas: pandas*. doi:[10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134).

- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É (2011). “**scikit-learn**: Machine Learning in Python.” *Journal of Machine Learning Research*, **12**, 2825–2830. URL <https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf>.
- Pegoraro M, Beraha M (2021). “Fast PCA in 1-D Wasserstein Spaces via B-Splines Representation and Metric Projection.” *Proceedings of the AAAI Conference on Artificial Intelligence*, **35**(10), 9342–9349. doi:10.1609/aaai.v35i10.17126.
- Peng H, Long F, Ding C (2005). “Feature Selection Based on Mutual Information Criteria of Max-Dependency, Max-Relevance, and Min-Redundancy.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **27**(8), 1226–1238. doi:10.1109/tpami.2005.159.
- Peng J, Paul D (2011). **fpca**: *Restricted MLE for Functional Principal Components Analysis*. R package version 0.2-1 (archived), URL <https://CRAN.R-project.org/package=fpca>.
- Pfisterer F, Beggel L, Sun X, Scheipl F, Bischl B (2021). “Benchmarking Time Series Classification – Functional Data vs Machine Learning Approaches.” *arXiv 1911.07511*, arXiv E-Print Archive. doi:10.48550/arXiv.1911.07511.
- Poskitt DS, Sengarapillai A (2013). “Description Length and Dimensionality Reduction in Functional Data Analysis.” *Computational Statistics & Data Analysis*, **58**, 98–113. doi:10.1016/j.csda.2011.03.018.
- Ramos-Carreño C (2020). “**dcor**: Distance Correlation and Related E-Statistics in Python.” Zenodo. doi:10.5281/zenodo.3468124.
- Ramos-Carreño C, Suárez A, Torrecilla JL, Carbajo Berrocal M, Marcos Manchón P, Pérez Manso P, Hernando Bernabé A, García Fernández D, Hong Y, Rodríguez-Ponga Eyriès PM, Sánchez Romero Á, Petrunina E (2022). **scikit-fda**: *Functional Data Analysis in Python*. doi:10.5281/zenodo.3468127.
- Ramsay JO, Hooker G, Graves S (2009). *Functional Data Analysis with R and MATLAB*. Springer-Verlag, New York. doi:10.1007/978-0-387-98185-7.
- Ramsay JO, Hooker G, Graves S (2024). **fda**: *Functional Data Analysis*. R package version 6.1.8, URL <https://CRAN.R-project.org/package=fda>.
- Ramsay JO, Silverman BW (2005). *Functional Data Analysis*. Springer-Verlag Series in Statistics, 2nd edition. Springer-Verlag, New York. doi:10.1007/b98888.
- R Core Team (2024). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Scheipl F, Arnone E, Hooker G, Tucker JD, Wrobel J (2022). “CRAN Task View: Functional Data Analysis.” Version 2022-03-21, URL <https://CRAN.R-project.org/view=FunctionalData>.
- Scheipl F, Goldsmith J, Wrobel J (2024). **tidyfun**: *Tools for Tidy Functional Data*. R package version 0.0.98, URL <https://github.com/tidyfun/tidyfun>.

- Schmutz A, Jacques J, Bouveyron C, Chèze L, Martin P (2020). “Clustering Multivariate Functional Data in Group-Specific Functional Subspaces.” *Computational Statistics*, **35**, 1101–1131. doi:[10.1007/s00180-020-00958-4](https://doi.org/10.1007/s00180-020-00958-4).
- Sørensen H, Goldsmith J, Sangalli LM (2013). “An Introduction with Medical Applications to Functional Data Analysis.” *Statistics in Medicine*, **32**(30), 5222–5240. doi:[10.1002/sim.5989](https://doi.org/10.1002/sim.5989).
- Sphinx** Development Team (2020). “**Sphinx** 4.0.0+ Documentation.” URL <https://www.sphinx-doc.org/en/master/>.
- Srivastava A, Klassen EP (2016). *Functional and Shape Data Analysis*. Springer-Verlag, New York. doi:[10.1007/978-1-4939-4020-2](https://doi.org/10.1007/978-1-4939-4020-2).
- Srivastava A, Wu W, Kurtek S, Klassen EP, Marron JS (2011). “Registration of Functional Data Using Fisher-Rao Metric.” *arXiv 1103.3817*, arXiv E-Print Archive. doi:[10.48550/arXiv.1103.3817](https://doi.org/10.48550/arXiv.1103.3817).
- Sun Y, Genton MG (2011). “Functional Boxplots.” *Journal of Computational and Graphical Statistics*, **20**(2), 316–334. doi:[10.1198/jcgs.2011.09224](https://doi.org/10.1198/jcgs.2011.09224).
- Székely GJ, Rizzo ML, Bakirov NK (2007). “Measuring and Testing Dependence by Correlation of Distances.” *The Annals of Statistics*, **35**(6), 2769–2794. doi:[10.1214/009053607000000505](https://doi.org/10.1214/009053607000000505).
- Tan CW, Bergmeir C, Petitjean F, Webb GI (2020). “Monash University, UEA, UCR Time Series Extrinsic Regression Archive.” *arXiv 2006.10996*, arXiv E-Print Archive. doi:[10.48550/arXiv.2006.10996](https://doi.org/10.48550/arXiv.2006.10996).
- Tan CW, Bergmeir C, Petitjean F, Webb GI (2021). “Time Series Extrinsic Regression.” *Data Mining and Knowledge Discovery*, **35**(3), 1032–1060. doi:[10.1007/s10618-021-00745-9](https://doi.org/10.1007/s10618-021-00745-9).
- The MathWorks Inc (2021). *MATLAB – The Language of Technical Computing, Version R2021a*. Natick. URL <https://www.mathworks.com/products/matlab/>.
- Torrecilla JL, Ramos-Carreño C, Sánchez-Montañés M, Suárez A (2020). “Optimal Classification of Gaussian Processes in Homo- And Heteroscedastic Settings.” *Statistics and Computing*, **30**(4), 1091–1111. doi:[10.1007/s11222-020-09937-7](https://doi.org/10.1007/s11222-020-09937-7).
- Torrecilla JL, Suárez A (2016). “Feature Selection in Functional Data Classification with Recursive Maxima Hunting.” In DD Lee, M Sugiyama, UV Luxburg, I Guyon, R Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 4835–4843. Curran Associates. URL <https://proceedings.neurips.cc/paper/2016/file/28b60a16b55fd531047c0c958ce14b95-Paper.pdf>.
- Tucker JD (2021). “**ElasticFDA**: Julia Library for Elastic Functional Data Analysis.” URL <https://github.com/jdtuck/ElasticFDA.jl>.
- Tucker JD (2024a). *fdasrsf: Functional Data Analysis Using the Square Root Slope Framework*. Python package version 2.5.9, URL <https://pypi.org/project/fdasrsf/>.

- Tucker JD (2024b). *fdasrvf: Elastic Functional Data Analysis*. R package version 2.2.0, URL <https://CRAN.R-project.org/package=fdasrvf>.
- Tuddenham RD, Snyder MM (1954). “Physical Growth of California Boys and Girls from Birth to Eighteen Years.” *Publications in Child Development. University of California, Berkeley*, **1**(2), 183–364.
- Ullah S, Finch CF (2013). “Applications of Functional Data Analysis: A Systematic Review.” *BMC Medical Research Methodology*, **13**(1), 43. doi:10.1186/1471-2288-13-43.
- Van Rossum G, Warsaw B, Coghlan N (2001). “PEP 8: Style Guide for Python Code.” Accessed 2020-09-28, URL <https://www.python.org/dev/peps/pep-0008/>.
- Van Rossum G, *et al.* (2011). *Python Programming Language*. URL <https://www.python.org/>.
- Vieu P (2018). “On Dimension Reduction Models for Functional Data.” *Statistics & Probability Letters*, **136**, 134–138. doi:10.1016/j.spl.2018.02.032.
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey CJ, Polat İ, Feng Y, Moore EW, VanderPlas J, Laxalde D, Perktold J, Cimrman R, Henriksen I, Quintero EA, Harris CR, Archibald AM, Ribeiro AH, Pedregosa F, van Mulbregt P (2020). “**SciPy** 1.0: Fundamental Algorithms for Scientific Computing in Python.” *Nature Methods*, **17**(3), 261–272. doi:10.1038/s41592-019-0686-2.
- Wagner-Muns IM, Guardiola IG, Samaranyake VA, Kayani WI (2018). “A Functional Data Analysis Approach to Traffic Volume Forecasting.” *IEEE Transactions on Intelligent Transportation Systems*, **19**(3), 878–888. doi:10.1109/tits.2017.2706143.
- Wang JL, Chiou JM, Müller HG (2016). “Functional Data Analysis.” *Annual Review of Statistics and Its Application*, **3**(1), 257–295. doi:10.1146/annurev-statistics-041715-033624.
- Wasserman L (2006). *All of Nonparametric Statistics*. Springer-Verlag, New York. doi:10.1007/0-387-30623-4.
- Wickham H, Averick M, Bryan J, Chang W, McGowan LD, François R, Grolemund G, Hayes A, Henry L, Hester J, Kuhn M, Pedersen TL, Miller E, Bache SM, Müller K, Ooms J, Robinson D, Seidel DP, Spinu V, Takahashi K, Vaughan D, Wilke C, Woo K, Yutani H (2019). “Welcome to the Tidyverse.” *Journal of Open Source Software*, **4**(43), 1686. doi:10.21105/joss.01686.
- Wrobel J, Park SY, Staicu AM, Goldsmith J (2016). “Interactive Graphics for Functional Data Analyses.” *Stat*, **5**(1), 108–118. doi:10.1002/sta4.109.
- Yao F, Müller HG, Wang JL (2015). “**PACE** Package for Functional Data Analysis and Empirical Dynamics (MATLAB).” URL <https://www.stat.ucdavis.edu/PACE/>.

Zhou Y, Bhattacharjee S, Carroll C, Chen Y, Dai X, Fan J, Gajardo A, Hadjipantelis PZ, Han K, Ji H, Zhu C, Mueller HG, Wang JL (2022). **fdapace**: *Functional Data Analysis and Empirical Dynamics*. R package version 0.5.9, URL <https://CRAN.R-project.org/package=fdapace>.

Affiliation:

Carlos Ramos-Carreño, Alberto Suárez
Department of Computer Science
Escuela Politécnica Superior
Universidad Autónoma de Madrid
28049 Madrid, Spain
E-mail: carlos.ramos@uam.es, alberto.suarez@uam.es

José Luis Torrecilla
Department of Mathematics
Facultad de Ciencias
Universidad Autónoma de Madrid
28049 Madrid, Spain
E-mail: joseluis.torrecilla@uam.es